Faster Pattern Matching under Edit Distance

A Reduction to Dynamic Puzzle Matching and

The Seaweed Monoid of Permutation Matrices

Panagiotis Charalampopoulos [iD]
Reichman University, Herzliya, Israel
pcharalampo@gmail.com

Tomasz Kociumaka [iD]
University of California, Berkeley, U.S.
kociumaka@berkeley.edu

Philip Wellnitz [iD]
Max Planck Institute for Informatics, SIC, Saarbrücken, Germany
wellnitz@mpi-inf.mpg.de

─────────── **Abstract** ───────────

We consider the approximate pattern matching problem under the *edit distance*. Given a text T of length n, a pattern P of length m, and a threshold k, the task is to find the starting positions of all substrings of T that can be transformed to P with at most k *edits*. More than 20 years ago, Cole and Hariharan [SODA'98, J. Comput.'02] gave an $\mathcal{O}(n + k^4 \cdot n/m)$-time algorithm for this classic problem, and this runtime has not been improved since.

Here, we present an algorithm that runs in time $\mathcal{O}(n + k^{3.5}\sqrt{\log m \log k} \cdot n/m)$, thus breaking through this long-standing barrier. In the case where $n^{1/4+\varepsilon} \leq k \leq n^{2/5-\varepsilon}$ for some arbitrarily small positive constant ε, our algorithm improves over the state-of-the-art by polynomial factors: it is polynomially faster than both the algorithm of Cole and Hariharan and the classic $\mathcal{O}(kn)$-time algorithm of Landau and Vishkin [STOC'86, J. Algorithms'89].

We observe that the bottleneck case of the alternative $\mathcal{O}(n + k^4 \cdot n/m)$-time algorithm of Charalampopoulos, Kociumaka, and Wellnitz [FOCS'20] is when the text and the pattern are (almost) periodic. Our new algorithm reduces this case to a new dynamic problem (Dynamic Puzzle Matching), which we solve by building on tools developed by Tiskin [SODA'10, Algorithmica'15] for the so-called seaweed monoid of permutation matrices. Our algorithm relies only on a small set of primitive operations on strings and thus also applies to the fully-compressed setting (where text and pattern are given as straight-line programs) and to the dynamic setting (where we maintain a collection of strings under creation, splitting, and concatenation), improving over the state of the art.
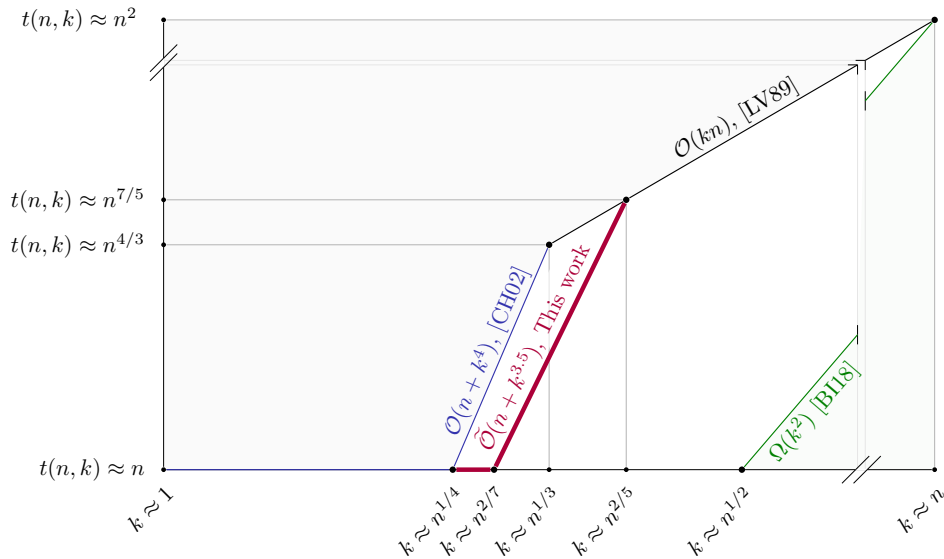
Figure 1 The running time $t(n,k)$ of algorithms for the approximate pattern matching problem under the edit distance as a function of k for the important special case where $m = \Theta(n)$. The scale is doubly logarithmic and sub-polynomial factors are hidden; running times below n and above n^2 are not relevant, neither are values of k that lie above n. Any point that lies strictly to the bottom-right of the green line segment is unattainable unless SETH fails.

Contents

1 Introduction

Almost every introductory algorithms textbook covers the *pattern matching* problem: in a given text T of length n, we wish to find all occurrences of a given pattern P of length m. As fundamental as both this problem and its solutions are by today, as apparent are their limitations: a single surplus or missing character in the pattern (or in a potential occurrence) results in (potentially all) occurrences being missed. Hence, a large body of work focuses on *approximate pattern matching*, where we want to identify substrings of the text that are *close* to the pattern. In particular, in this paper, we consider a classic variant of approximate pattern matching where we allow for up to k insertions, deletions, and substitutions (collectively: edits); that is, we consider approximate pattern matching under the edit distance.

Formally, for two strings X and Y, their edit distance (also known as the Levenshtein distance) $\delta_E(X, Y)$, is the minimum number of insertions, deletions, and substitutions of single characters required to transform X into Y. Now, in the *pattern matching with edits* problem, for a given text T, pattern P, and an integer threshold $k > 0$, the task is to find the starting positions of all *k-error* (or *k-edit*) *occurrences* of P in T. Specifically, we wish to list all positions v in T such that the edit distance between $T[v \mathinner{.\,.} w) = T[v]T[v+1]\cdots T[w-1]$ and P is at most k for some position w; we write $\mathrm{Occ}_k^E(P, T)$ to denote the set of all such positions v.

Let us highlight the main prior results for pattern matching with edits; for a thorough review of other (in particular) early results on pattern matching with edits, we refer to the extensive survey of Navarro [Nav01]. Back in 1980, Sellers [Sel80] demonstrated how the standard dynamic-programming algorithm for computing $\delta_E(P, T)$ can be adapted to an $\mathcal{O}(nm)$-time algorithm for the pattern matching with edits problem. Around the same time, Masek and Paterson [MP80] reduced the running time by a poly-logarithmic factor using the Four-Russians technique. Only several year later, Landau and Vishkin [LV88] presented an $\mathcal{O}(nk^2)$-time solution, which they could then improve to the—by now—classic "kangaroo jumping" algorithm that solves this problem in $\mathcal{O}(nk)$ time [LV89]. In search of even faster algorithms, Sahinalp and Vishkin [SV96] developed an algorithm that runs in time $\mathcal{O}(n + nk^{8+1/3}(\log^* n)^{1/3}/m^{1/3})$— this algorithm was then improved by Cole and Hariharan [CH02], who gave an $\mathcal{O}(n + k^4 n/m)$-time solution, which is asymptotically faster than the aforementioned Landau–Vishkin algorithm when $k = o(\sqrt[3]{m})$, and in that setting also the fastest known algorithm even today.

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 [BI18] recently proved that any polynomial-factor improvement would yield a major breakthrough for the satisfiability problem. For pattern matching with edits, this means that there is no hope for an algorithm running in time $\mathcal{O}(n + k^{2-\varepsilon}n/m)$ for any constant $\varepsilon > 0$: given an $\mathcal{O}(n + k^{2-\varepsilon}n/m)$-time algorithm for pattern matching with edits, we could compute the edit distance of any two given strings X and Y of total length N over an alphabet Σ in time $\mathcal{O}(N^{2-\varepsilon}\log N)$. Specifically, we pad X and Y to $P := \$^{2N}X\$^{2N}$ and $T := \$^{2N}Y\$^{2N}$, where $\$ \notin \Sigma$. Now, as $\min_{v,w}\delta_E(P, T[v \mathinner{.\,.} w)) = \delta_E(P, T) = \delta_E(X, Y)$, we can binary search for the smallest value of k such that $\mathrm{Occ}_k^E(P, T)$ is not empty.

Despite the large gap between the quadratic and bi-quadratic dependency on k, no further advancements have been made to settle the running time of the pattern matching with edits problem. In particular, there has not even been any progress on resolving the 24-year-old conjecture of Cole and Hariharan [CH02] that an $\mathcal{O}(n + k^3 n/m)$-time algorithm *should be possible*—until now. We give the first algorithm that improves over the running time achieved by Cole and Hariharan [CH02]:

⬛ **Main Theorem 1**. *Given a text T of length n, a pattern P of length m, and an integer threshold $k > 0$, we can compute the set $\mathrm{Occ}_k^E(P, T)$ in $\mathcal{O}(n + n/m \cdot k^{3.5}\sqrt{\log m \log k})$ time.* ◪

Observe that if k is roughly between $n^{1/4}$ and $n^{2/7}$, we obtain the first linear-time algorithm for the important special case where text and pattern are close in length. Further, we still obtain polynomial improvements in the running time for values of k that are roughly less than $n^{2/5}$. Consult Figure 1 for a graphical comparison of the running times of our algorithm with the previous state-of-the-art and the conditional lower bound discussed above.

The `PILLAR` Model and Faster Algorithms in Other Settings

Our approach is reasonably general and allows for an easy adaption to different settings (where the text and the pattern are not given explicitly). In particular, we follow the approach by Charalampopoulos, Kociumaka, and Wellnitz [CKW20] and implement the algorithm in the so-called `PILLAR` model. In that model, one bounds the running times of algorithms in terms of the number of calls to a small set of very common operations (the `PILLAR` operations) on strings, such as computing the length of their longest common prefix. Then, for any setting, an efficient implementation of the `PILLAR` operations yields a fast algorithm for approximate pattern matching. For pattern matching with edits, [CKW20] presented an algorithm that runs in $\mathcal{O}(n/m \cdot k^4)$ time in the `PILLAR` model. We improve upon their algorithm.

▟ **Main Theorem 2**. *Given a pattern P of length m, a text T of length n, and an integer threshold $k > 0$, we can compute a representation of the set $\mathrm{Occ}_k^E(P,T)$ as $\mathcal{O}(n/m \cdot k^3)$ arithmetic progressions with the same difference in $\mathcal{O}(n/m \cdot k^{3.5}\sqrt{\log m \log k})$ time in the `PILLAR` model.* ▟

Consistently with [CKW20], we represent the output set $\mathrm{Occ}_k^E(P,T)$ as $\mathcal{O}(k^3)$ disjoint arithmetic progressions with a common difference. Unless P is almost periodic, though, $\mathrm{Occ}_k^E(P,T)$ is of size $\mathcal{O}(k^2)$, and we can report $\mathrm{Occ}_k^E(P,T)$ explicitly; see [CKW20] for a structural characterization of $\mathrm{Occ}_k^E(P,T)$.

Now, in the standard setting, where the text and the pattern are both given explicitly, after an $\mathcal{O}(n)$-time preprocessing, we can perform each primitive `PILLAR` operation in constant time. We thus instantly obtain Main Theorem 1. The same `PILLAR` implementation remains valid in the *internal setting* introduced in [KRRW15]. Specifically, after a linear-time preprocessing of an input string X, the algorithm of Main Theorem 2 can efficiently compute $\mathrm{Occ}_k^E(P,T)$ for any two fragments P,T of the string X.

In Section 7, we show that existing implementations of the primitive operations of the `PILLAR` model allow us to also obtain efficient algorithms for pattern matching under edit distance in the fully-compressed setting (where the text and the pattern are given as straight-line programs) and in the dynamic setting (where we maintain a collection of strings under creation, splitting, and concatenation). Our algorithms improve over the state-of-the-art algorithms of [CKW20] for these settings: we trade a $\sqrt{k}$ factor for a factor that is asymptotically upper-bounded by the logarithm of the length of the considered pattern. Formally, we obtain the following results.

▟ **Main Theorem 3**. *We can maintain a collection $\mathcal{X}$ of non-empty persistent strings of total length N subject to $\mathtt{makestring}(U)$, $\mathtt{concat}(U,V)$, and $\mathtt{split}(U,i)$ operations that require $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$, and $\mathcal{O}(\log N)$ time, respectively, so that given two strings $P,T \in \mathcal{X}$, and an integer threshold $k > 0$, we can compute a representation of $\mathrm{Occ}_k^E(P,T)$ as $\mathcal{O}(|T|/|P| \cdot k^3)$ arithmetic progressions with the same difference in time $\mathcal{O}(|T|/|P| \cdot k^{3.5}\sqrt{\log|P|\log k}\log^2 N)$.*[1] ▟

▟ **Main Theorem 4**. *Let $\mathcal{G}_T$ denote a straight-line program of size n generating a string T, let $\mathcal{G}_P$ denote a straight-line program of size m generating a string P, let $k > 0$ denote an integer threshold, and set $N := |T|$ and $M := |P|$. We can compute $|\mathrm{Occ}_k^E(P,T)|$ in time $\mathcal{O}(m\log N + n\,k^{3.5}\log^2 N\sqrt{\log M \log k}\log\log N)$ and we can report the elements of $\mathrm{Occ}_k^E(P,T)$ within $\mathcal{O}(|\mathrm{Occ}_k^E(P,T)|)$ extra time.* ▟

---

[1] All running time bounds hold with high probability (that is, $1 - 1/N^{\Omega(1)}$). A deterministic version can be obtained at the cost of a $\mathrm{poly}(\log\log N)$-factor overhead.

1.1 Related Work

Pattern Matching with Mismatches. The Hamming distance of two (equal-length) strings is the number of positions where the strings differ. This metric is more restrictive than edit distance since it allows substitutions but does not support insertions or deletions.

In the *pattern matching with mismatches* problem, we are given a text T of length n, a pattern P of length m, and an integer threshold $k > 0$, and we wish 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 problem has been extensively studied since the 1980s. A long line of works [Abr87, Kos87, LV86, GG86, ALP04, CFP$^+$16, GU18, CGK$^+$20] has culminated in an $\widetilde{\mathcal{O}}(n + kn/\sqrt{m})$-time algorithm, presented by Gawrychowski and Uznański [GU18], who also 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.

As shown in [CKW20], pattern matching with mismatches admits an $\widetilde{\mathcal{O}}(k^2 \cdot n/m)$-time algorithm in the `PILLAR` model. Analogously to pattern matching with edits, this solution constitutes the basis of the state-of-the-art algorithms in the internal, fully-compressed, and dynamic settings.

Online Algorithms for Pattern Matching with Edits. The pattern matching with edits problem has also been considered in the online setting where the text arrives character by character and, by the time $T[w]$ becomes available, the algorithm needs to decide whether $\min_v \delta_E(P, T[v \mathinner{..} w]) \leq k$. Landau, Myers, and Schmidt [LMS98] provided an online algorithm that runs in $\mathcal{O}(k)$ time per character. Subsequent work focused on the streaming model, whether the main emphasis is on reducing the space complexity of an online algorithm, usually at the cost of introducing Monte-Carlo randomization. Starikovskaya [Sta17] presented an algorithm for this setting with both the space usage and the time required to process each character of the text being proportional to $\sqrt{m}(k \log m)^{\mathcal{O}(1)}$. Very recently, Kociumaka, Porat, and Starikovskaya [KPS21], improved upon this result, presenting an algorithm that uses $\widetilde{\mathcal{O}}(k^5)$ space and processes each character of the text in $\widetilde{\mathcal{O}}(k^8)$ amortized time; here, $\widetilde{\mathcal{O}}(\star)$ hides $\log^{\mathcal{O}(1)} m$ factors.

Approximating Pattern Matching with Edits. Chakraborty, Das, and Koucký [CDK19] presented an $\widetilde{\mathcal{O}}(nm^{3/4})$-time algorithm that produces, for each position w of the text, a constant factor approximation of $\min_v \delta_E(P, T[v \mathinner{..} w])$. They also provided an online algorithm with a weaker approximation guarantee.

1.2 Open Problems

The most important and obvious open problem is to close the gap between upper and lower bounds for the pattern matching with edits problem; as is depicted in Figure 1. In the quest for faster algorithms, one could try to relax the problem in scope, for instance, by considering its (easier) decision version where we only need to check whether $\mathrm{Occ}_k^E(P, T)$ is empty, or by allowing for some approximation by also reporting an arbitrary subset of the positions in $\mathrm{Occ}_{(1+\varepsilon)k}^E(P, T) \setminus \mathrm{Occ}_k^E(P, T)$ for a small $\varepsilon > 0$.

Another research direction could be to devise an algorithm with an analogous running time as the one presented here that reports all fragments of T that are at edit distance at most k from P (in appropriate batches); recall that $\mathrm{Occ}_k^E(P, T)$ is only the set of the starting positions of such fragments. While we think that the $\mathcal{O}(k^4 \cdot n/m)$-time `PILLAR` algorithm of [CKW20] can be generalized to report all such fragments, our $\widetilde{\mathcal{O}}(k^{3.5} \cdot n/m)$-time solution *does not* seem to generalize. We remark that Landau, Myers, and Schmidt [LMS98] showed that all the sought fragments can be listed in $\mathcal{O}(nk)$ time; for this, they adapted the algorithm of [LV89].

1.3 Technical Overview

For a string P (also called a *pattern*), a string T (also called a *text*), and an integer $k > 0$ (also called a *threshold*), we say that P has a *k-error occurrence* in T at position v if we have $\delta_E(P, T[v \mathinner{..} w)) \le k$ for some $w \ge v$. We write $\mathrm{Occ}_k^E(P, T)$ to denote the set of the starting positions of k-error occurrences of P in T, that is, $\mathrm{Occ}_k^E(P, T) \coloneqq \{v : \exists_{w \ge v}\, \delta_E(P, T[v \mathinner{..} w)) \le k\}$. We now formally state the *pattern matching with edits* problem.

$\mathrm{PMWITHEDITS}(P, T, k)$
Input: A pattern P of length m, a text T of length n, and a positive integer $k \le m$.
Output: The set $\mathrm{Occ}_k^E(P, T)$.

The NewPeriodicMatches Problem

Let us start with a short exposition of parts of our notation.[2] A string S is *primitive* if it cannot be expressed as U^y for a string U and an integer $y > 1$. For two strings U and V, we write $\delta_E(U, V^*) \coloneqq \min\{\delta_E(U, V^\infty[0 \mathinner{..} j)) : j \in \mathbb{Z}_{\ge 0}\}$ to denote the minimum edit distance between U and any prefix of $V^\infty = V \cdot V \cdots$. Further, we write $\delta_E(U, {}^*V^*) \coloneqq \min\{\delta_E(U, V^\infty[i \mathinner{..} j)) : i, j \in \mathbb{Z}_{\ge 0}, i \le j\}$ to denote the minimum edit distance between U and any substring of V^∞.

As we explain in Sections 2.2 and 3.1, a recent algorithm of Charalampopoulos, Kociumaka, and Wellnitz [CKW20] reduces the $\mathrm{PMWITHEDITS}$ problem to several instances of the following restricted variant; the reduction takes $\mathcal{O}(n/m \cdot k^3)$ time in the `PILLAR` model.

$\mathrm{NEWPERIODICMATCHES}(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$
Input: A pattern P of length m, an integer threshold $k \in [0 \mathinner{..} m]$, a positive integer $d \ge 2k$, a text T of length $n \in [m - k \mathinner{..} \lceil 3/2\, m \rceil + k)$, a primitive string Q of length $q \coloneqq |Q| \le m/8d$, an edit-distance alignment $\mathcal{A}_P : P \rightsquigarrow Q^\infty[0 \mathinner{..} y_P)$ of cost $d_P \coloneqq \delta_E(P, {}^*Q^*) = \delta_E(P, Q^*) \le d$, and an edit-distance alignment $\mathcal{A}_T : T \rightsquigarrow Q^\infty[x_T \mathinner{..} y_T)$ of cost $d_T \coloneqq \delta_E(T, {}^*Q^*) \le 3d$, where $x_T \in [0 \mathinner{..} q)$.
Output: The set $\mathrm{Occ}_k^E(P, T)$ represented as $\mathcal{O}(d^3)$ disjoint arithmetic progressions with difference q.

Specifically, [CKW20] implies the following reduction.

▰ **Fact 1.1.** *Let P denote a pattern of length m, let T denote a text of length n, and let $k \le m$ denote a positive integer.*

We can compute a representation of the set $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(n/m \cdot k^3)$ disjoint arithmetic progressions with the same difference in time $\mathcal{O}(n/m \cdot k^3)$ in the `PILLAR` model plus the time required for solving several instances $\mathrm{NEWPERIODICMATCHES}(P_i, T_i, k_i, d_i, Q_i, \mathcal{A}_{P_i}, \mathcal{A}_{T_i})$, where $\sum_i |P_i| = \mathcal{O}(n)$ and, for each i, we have $|P_i| \le m$ and $d_i = \lceil 8k/m \cdot |P_i| \rceil$. ◢

▰ **Remark 1.2.** In the case where Fact 1.1 is applied to an instance of the $\mathrm{PMWITHEDITS}$ problem such that the pattern P is approximately periodic, the input $(P_i, T_i, k_i, d_i, Q_i, \mathcal{A}_{P_i}, \mathcal{A}_{T_i})$ to each produced instance of $\mathrm{NEWPERIODICMATCHES}$ satisfies the following conditions: $P_i = P$, T_i is a fragment of T, $k_i = k$, and $d = \mathcal{O}(k)$. For the purposes of this technical overview, one can focus solely on that case. ◢

Using the algorithm `PeriodicMatches(P, T, k, d, Q)` of [CKW20, Lemma 6.11] to solve the $\mathrm{NEWPERIODICMATCHES}$ problem in $\mathcal{O}(d^4)$ time in the `PILLAR` model, the time (in the `PILLAR` model) required

---

[2] See also Section 2, where we provide a comprehensive exposition of the notation used throughout this paper, including those we consider standard. Further, consider the notation tables at the very end of this paper for a quick reference for the most important notations.

for solving all instances of the NEWPERIODICMATCHES problem that are generated by Fact 1.1 is

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

In particular, we can reinterpret the $\mathcal{O}(n/m \cdot k^4)$-time algorithm of [CKW20] for the PMWITHEDITS problem as a combination of Fact 1.1 and `PeriodicMatches(P, T, k, d, Q)`. Our main contribution is the following faster algorithm for the NEWPERIODICMATCHES problem.

▉ **Lemma 1.3** (`NewPeriodicMatches(P, T, k, d, Q, A_P, A_T)`). *We can solve the NEWPERIODIC-MATCHES problem in $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ time in the* `PILLAR` *model.* ▟

By combining Fact 1.1 and Lemma 1.3, we obtain Main Theorem 2.

▉ **Main Theorem 2**. *Given a pattern P of length m, a text T of length n, and an integer threshold $k > 0$, we can compute a representation of the set $\mathrm{Occ}_k^E(P,T)$ as $\mathcal{O}(n/m \cdot k^3)$ arithmetic progressions with the same difference in $\mathcal{O}(n/m \cdot k^{3.5}\sqrt{\log m \log k})$ time in the* `PILLAR` *model.*

Proof. By Fact 1.1, in $\mathcal{O}(n/m \cdot k^3)$ time, we can reduce the PMWITHEDITS problem to several instances NEWPERIODICMATCHES$(P_i, T_i, k_i, d_i, Q_i, \mathcal{A}_{P_i}, \mathcal{A}_{T_i})$, where $\sum_i |P_i| = \mathcal{O}(n)$ and, for each i, we have $|P_i| \leq m$ and $d_i = \lceil 8k/m \cdot |P_i| \rceil$. By Lemma 1.3, the time required for solving all of the obtained instances (in the `PILLAR` model) is

$$\sum_i \mathcal{O}(d_i^{3.5}\sqrt{\log |T_i| \log d_i}) = \sum_i \mathcal{O}(k^{3.5}/m^{3.5} \cdot |P_i|^{3.5}\sqrt{\log m \log k})$$
$$= \sum_i \mathcal{O}(k^{3.5}/m \cdot |P_i|\sqrt{\log m \log k})$$
$$= \mathcal{O}(n/m \cdot k^{3.5}\sqrt{\log m \log k}). \qquad ▟$$

A Fast Algorithm for the NewPeriodicMatches Problem

We continue with a high-level description of the algorithm that underlies Lemma 1.3. In what follows, for simplicity, we assume that $m/k \gg q \gg k$ and that both $Q^\infty[\,0\mathinner{\ldotp\ldotp}y_P\,)$ and $Q^\infty[\,x_T\mathinner{\ldotp\ldotp}y_T\,)$ are powers of Q.

A First Solution via the DynamicPuzzleMatching Problem. Let us first discuss how the (almost) periodicity of P and T yields a simple way to *filter out* many potential starting positions of k-error occurrences.

As an introductory example, suppose that P and T are perfectly periodic with period Q, that is, $P = Q^\infty[\,0\mathinner{\ldotp\ldotp}m\,)$ and $T = Q^\infty[\,0\mathinner{\ldotp\ldotp}n\,)$. Observe that in this special case, $\mathcal{A}_P$ and $\mathcal{A}_T$ are cost-0 alignments (that is, $d_T = d_P = 0$), and we have $m = y_P$, $0 = x_T$, and $n = y_T$. Next, we argue that all k-error occurrences of P in T start around the positions in T where an exact occurrence of Q starts, that is, in the intervals $[\,jq - k\mathinner{\ldotp\ldotp}jq + k\,]$ for $j \in \mathbb{Z}$.[3] To that end, observe that, for any alignment of cost at most k mapping P to a fragment $T[\,v\mathinner{\ldotp\ldotp}w\,)$ of T, at least one of the copies of Q that comprise P must match exactly; otherwise the edit distance would be much larger than k. Suppose that the i-th copy of Q, that is, $P[\,iq\mathinner{\ldotp\ldotp}(i+1)q\,)$, is matched exactly. As Q is primitive and hence does not match any of its non-trivial rotations, $P[\,iq\mathinner{\ldotp\ldotp}(i+1)q\,)$ must be matched with a fragment $T[\,i'q\mathinner{\ldotp\ldotp}(i'+1)q\,)$ of T. As the entire alignment makes at most k insertions and deletions, this implies that $v \in [\,(i'-i)q - k\mathinner{\ldotp\ldotp}(i'-i)q + k\,]$.

---

[3] Under our earlier assumption that $q \gg k$, this claim indeed allows for filtering out some positions where no occurrence may start as we have $jq + k \ll (j+1)q - k$ in that case.

S_4 `a t t e r n m a t c h`

S_3 `i m a t e p a t t e` `a t c h` S_5

S_2 `o x i m a t` `a t c h i n g e d i t` S_6

S_1 `a p p r o x i m` `e d i t d i s t a n c e` S_7

S `a p p r o x i m a t e p a t t e r n m a t c h i n g e d i t d i s t a n c e`
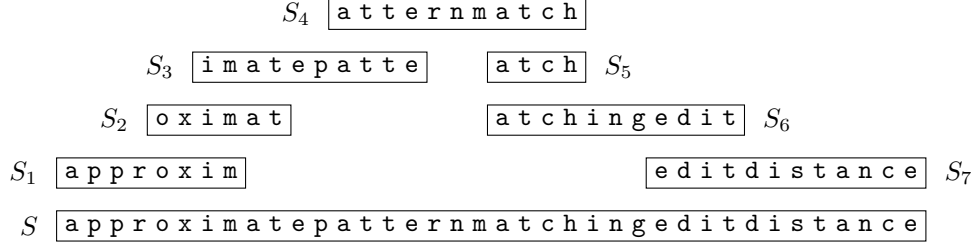
Figure 2 $S_1, \ldots, S_7$ is a 4-puzzle whose value is S.

Now, the strings P and T are only *almost* periodic—in particular, the edits in $\mathcal{A}_T$ and $\mathcal{A}_P$ may widen the intervals of potential starting positions, albeit only by a $d_T + d_P \le 4d$ additive term. Since $k < d$, we have $\mathrm{Occ}_k^E \subseteq \bigcup_{j \in \mathbb{Z}} [jq - 5d \mathinner{.\,.} jq + 5d]$. Hence, for each j, we define a fragment $R_j = T[r_j \mathinner{.\,.} r_j')$ of T that is of length $m + \mathcal{O}(d)$ and, in the considered instance, is *responsible* for capturing k-error occurrences of P in T that start in $[jq - 5d \mathinner{.\,.} jq + 5d]$; specifically, we have $r_j + \mathrm{Occ}_k^E(P, R_j) \supseteq \mathrm{Occ}_k^E(P, T) \cap [jq - 5d \mathinner{.\,.} jq + 5d]$. In addition, we identify a set $J \subseteq \mathbb{Z}$ of size $\mathcal{O}(m/q)$ such that $\mathrm{Occ}_k^E(P, T) = \bigcup_{j \in J} (r_j + \mathrm{Occ}_k^E(P, R_j))$.

Our goal is to compute occurrences of P in each R_j separately. To that end, observe that both P and all R_js essentially decompose into (possibly slightly "edited") copies of Q. In particular, for $j, j + 1 \in J$, we can obtain R_{j+1} from R_j by replacing $\mathcal{O}(d)$ such "edited" copies. As a first step toward capturing the notions of P and R_j decomposing into *pieces* and our algorithm *replacing* pieces of R_j, we define Δ-*puzzles*; consult Figure 2 for a visualization of an example of a Δ-puzzle.

Definition 1.4. *For a $\Delta \in \mathbb{Z}_{\ge 0}$, we say that $z \ge 2$ strings $S_1, \ldots, S_z$ form a Δ-puzzle if*
- $|S_i| \ge \Delta$ *for each $i \in [1 \mathinner{.\,.} z]$, and*
- $S_i[|S_i| - \Delta \mathinner{.\,.} |S_i|) = S_{i+1}[0 \mathinner{.\,.} \Delta)$ *for each $i \in [1 \mathinner{.\,.} z)$.*

The value *of the puzzle is* $\mathrm{val}_\Delta(S_1, \ldots, S_z) := S_1 \cdot S_2[\Delta \mathinner{.\,.} |S_2|) \cdot S_3[\Delta \mathinner{.\,.} |S_3|) \cdots S_z[\Delta \mathinner{.\,.} |S_z|)$.

In Section 4, we define *pieces* $P_1, \ldots, P_z$ and $T_{j,1}, \ldots, T_{j,z}$ (for each $j \in J$) that form Δ-puzzles with values P and R_j, respectively, where $\Delta := 6(d_P + d_T + k)$. Let us intuitively describe these pieces.[4] First, let us partition both P and T into *tiles*, that is, maximal fragments that are aligned to different copies of Q by $\mathcal{A}_P$ and $\mathcal{A}_T$, respectively. Observe that all but $\mathcal{O}(d)$ tiles are exact copies of Q. Further, the endpoints R_j are $\mathcal{O}(d)$ positions apart from tile boundaries. We then obtain an induced partition for R_j by extending the first and last tiles that it fully contains by $\mathcal{O}(d)$ positions. Finally, we extend all tiles of the partition of P and the induced partition of R_j, other than the trailing ones, by Δ characters to the right. Consult Figure 3 for a visualization of this setting.

We call pieces $P_2, \ldots, P_{z-1}$ and $T_{j,2}, \ldots, T_{j,z-1}$ (for $j \in J$) *internal*. Observe that, for each i, all internal pieces of the form $T_{j,i'}$ with $j + i' = i$ coincide; that is, overlapping parts of different R_js share their internal pieces. Hence, for each $i \in (\min J + 1 \mathinner{.\,.} \max J + z)$, we define $T_i := T_{j,i'}$ for any $j \in J$ and $i' \in (1 \mathinner{.\,.} z)$ with $j + i' = i$. This is an essential property for our approach to work: when moving from R_j to R_{j+1}, we exploit that we need to only shift the pieces T_i, and not recompute them altogether.

Now, suppose that we can efficiently maintain a pair of Δ-puzzles so that we can at any time efficiently query for the k-error occurrences of the value of the first puzzle in the value of the second one. Then, as a warm-up solution, we can initialize the two puzzles as $P_1, \ldots, P_z$ and $T_{\min J,1}, \ldots, T_{\min J,z}$ and then

---

[4] This description provides an oversimplified definition of pieces. In particular, as defined in Section 4, P_1 covers at least 2 tiles whereas P_z covers 17 tiles. This is due to complications arising without the assumption $q \gg k$.
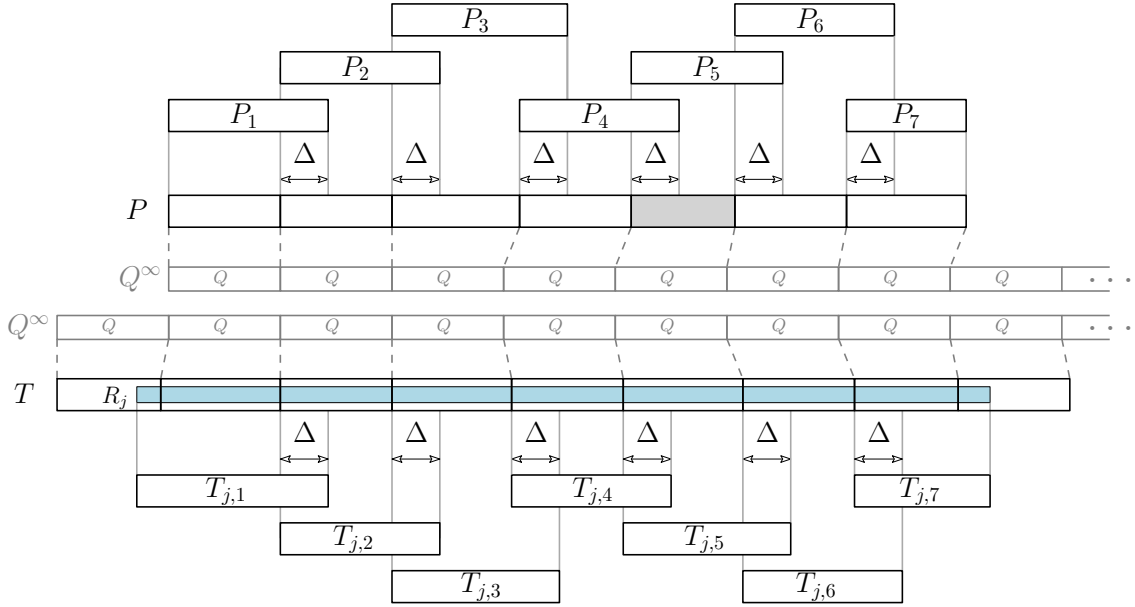
Figure 3 An alignment $\mathcal{A}_P : P \rightsquigarrow Q^7$ and an alignment $\mathcal{A}_T : P \rightsquigarrow Q^9$ are shown. Both P and T are partitioned into tiles. Specifically, dashed lines indicate the copy of Q to which a tile of P (or T) is aligned by $\mathcal{A}_P$ (or $\mathcal{A}_T$). For example, $\mathcal{A}_P$ aligns the shaded tile of P with the fifth copy of Q. The fragment R_j starts $\mathcal{O}(d)$ positions prior to the start of the second tile of T and ends $\mathcal{O}(d)$ positions after the end of the eighth tile. The pieces $P_1, \ldots, P_7$ and $T_{j,1}, \ldots, T_{j,7}$ form Δ-puzzles with values P and R_j, respectively.

replace pieces of the second puzzle as necessary in order to iterate over puzzles $T_{j,1}, \ldots, T_{j,z}$ for all $j \in J$. In fact, our final algorithm iterates over carefully trimmed versions of such puzzles, where we omit *plain* pieces that do not contribute to the solution set in an interesting manner. Formally, we capture the problem of maintaining such a pair of puzzles with the DYNAMICPUZZLEMATCHING problem.

DYNAMICPUZZLEMATCHING$(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$

Input: Positive integers k and Δ, as well as string families $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, and $\mathcal{S}_\varphi$ of *leading*, *internal*, and *trailing* pieces, respectively.

Maintained object: A sequence $\mathcal{I} = (U_1, V_1)(U_2, V_2) \cdots (U_z, V_z)$ of ordered pairs of strings (a *DPM-sequence*), that additionally satisfies the following two conditions:

(a) $U_1, V_1 \in \mathcal{S}_\beta$, $U_z, V_z \in \mathcal{S}_\varphi$, and, for all $i \in (1 \mathinner{.\,.} z)$, $U_i, V_i \in \mathcal{S}_\mu$,

(b) The torsion $\mathrm{tor}(\mathcal{I}) := \sum_{i=1}^z \big| |U_i| - |V_i| \big|$ satisfies $\mathrm{tor}(\mathcal{I}) \leq \Delta/2 - k$.

Update operations:

- DPM-Delete(i): Delete the i-th pair of strings.
- DPM-Insert((U', V'), i): Insert the pair of strings (U', V') after the i-th pair of strings.
- DPM-Substitute((U', V'), i): Substitute the i-th pair of strings with the pair of strings (U', V').

It is assumed that $\mathcal{I}$ satisfies conditions (a) and (b) at initialization time and after each update.

Query (DPM-Query): Return $\mathrm{Occ}_k^E(\mathcal{I}) := \mathrm{Occ}_k^E(\mathsf{val}_\Delta(U_1, \ldots, U_z), \mathsf{val}_\Delta(V_1, \ldots, V_z))$ under a promise that $U_1, \ldots, U_z$ and $V_1, \ldots, V_z$ are Δ-puzzles.

We move on to our main result for the DYNAMICPUZZLEMATCHING problem. For a precise statement, we need to be able to quantify the complexity of the input families of strings; formally we define the *median edit distance* of a family $\mathcal{S}$ of strings over an alphabet Σ as $\delta_E(\mathcal{S}) := \min_{\hat{S} \in \Sigma^*} \sum_{S \in \mathcal{S}} \delta_E(S, \hat{S})$. Now, our result reads as follows.

▪ **Theorem 1.5**.   *There is a data structure for* $\text{DYNAMICPUZZLEMATCHING}(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ *with* $\mathcal{O}(\Delta \log z \log \Delta)$*-time updates and queries,* $\mathcal{O}(\Delta z \log \Delta)$*-time initialization, and* $\mathcal{O}((d^3 + \Delta^2 d) \log^2(d + \Delta))$*-time preprocessing, where* $d = \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi)$.[5]   ◣

Let us defer a detailed discussion of Theorem 1.5 (proved in Section 9) to the end of this overview. Here, we discuss its application to the NEWPERIODICMATCHES problem with the following string families.

$$\mathcal{S}_\beta := \{P_1\} \cup \{T_{j,1} : j \in J\},$$
$$\mathcal{S}_\mu := \{P_i : i \in (1 .. z)\} \cup \{T_i : i \in (\min J + 1 .. \max J + z)\},$$
$$\mathcal{S}_\varphi := \{P_z\} \cup \{T_{j,z} : j \in J\}.$$

Next, we define multisets $\textsf{Special}(P)$, $\textsf{Special}(T)$, and $\textsf{Special}_{\beta\varphi}(T)$ of *special pieces*—in this overview, we focus on the two former multisets. In our fixed instance, we have

$$\textsf{Special}(P) = \{P_i : i \in (1 .. z) \text{ and } P_i \neq Q^\infty[0 .. q + \Delta)\},$$
$$\textsf{Special}(T) = \{T_i : i \in (1 + \min J .. z + \max J) \text{ and } T_i \neq Q^\infty[0 .. q + \Delta)\}.$$

As mentioned earlier, there are only very few special pieces—crucially, we show the following lemma.

▪ **Lemma 1.6**.   *The median edit distance of each of the families* $\mathcal{S}_\beta$*,* $\mathcal{S}_\mu$*, and* $\mathcal{S}_\varphi$ *is bounded by* $\mathcal{O}(d)$*.*
Further, each of the multisets $\textsf{Special}(P)$*,* $\textsf{Special}(T)$*, and* $\textsf{Special}_{\beta\varphi}(T)$ *is of size* $\mathcal{O}(d)$ *and can be computed in* $\mathcal{O}(d)$ *time in the* `PILLAR` *model.*   ◣

For $j \in J$, let $\mathcal{I}_j$ denote the DPM-sequence $(P_1, T_{j,1})(P_2, T_{j,2}) \cdots (P_z, T_{j,z})$. Now for $j \in J \setminus \max J$, each of $\mathcal{I}_j$ and $\mathcal{I}_{j+1}$ contains $\mathcal{O}(d)$ *special* pairs, that is, pairs that contain special pieces. Hence, we can naively iterate over all $\mathcal{I}_j$s in an instance of the DYNAMICPUZZLEMATCHING problem using $\mathcal{O}(d \cdot |J|)$ updates; in the considered instance we have $d \cdot |J| = \mathcal{O}(d \cdot m/q)$. See Section 4.2 for details of this reduction in the general case. As a preliminary improvement, in Section 4.4, we show how to reduce the number of updates to $\mathcal{O}(d^3)$. Let us give a brief sketch of this reduction.

We call a pair of pieces $(P_i, T_{j,i})$ *plain* if $i \in (1 .. z)$ and neither P_i nor $T_{j,i}$ is special; in the restricted case that we are considering here, the second condition is equivalent to $P_i = T_{j,i} = Q^\infty[0 .. q + \Delta)$. For $j \in J$, let $\mathcal{I}'_j$ denote the DPM-sequence obtained from $\mathcal{I}_j$ by trimming each run of plain pairs (that is, maximal contiguous subsequences that consist of plain pairs) in $\mathcal{I}_j$ to length $k + 1$ by deleting excess pairs.

The main idea is that we do not gain or lose any k-error occurrences by trimming the DPM-sequences, that is, we have $\text{Occ}_k^E(\mathcal{I}_j) = \text{Occ}_k^E(\mathcal{I}'_j)$. One direction is easy: removing the same substring from two strings P and T may only decrease the edit distance between P and T; this naturally translates to DPM-sequences. For the other direction, observe that if a DPM-sequence contains a run of at least $k + 1$ plain pairs, then any cost-k alignment between the corresponding strings has to perfectly match at least one copy of Q pair in such a run—we can hence duplicate said copy by adding more plain pairs in the DPM-sequence without increasing the cost of the alignment. Induction then yields the claim.

With the aim of obtaining an $\mathcal{O}(d^3)$ upper bound on the number of required updates for iterating over the $\mathcal{I}'_j$s, let us think of the process of shifting P along T. For each j, each run of plain pairs in $\mathcal{I}_j$ can be attributed to a run of plain pieces in $P_2, \ldots, P_{z-1}$ that overlap a run of plain pieces in $T_{j,2} \ldots T_{j,z-1}$. As we shift P, in the most general case, the length of the overlap first increases, then it remains static, and, finally, it decreases. Overall, as j gets incremented, a run of plain pairs that is attributed to a specific pair of runs of plain pieces may change length $\omega(d)$ times. However, after trimming the lengths of all runs

---

[5] Recall that z is the length of the DPM-sequence that we maintain in the data structure.

of plain pairs to $k + 1$, the length of such a run gets incremented/decremented $\mathcal{O}(d)$ times. As we have $\mathcal{O}(d)$ special pieces in each of P and T, we have $\mathcal{O}(d^2)$ pairs of runs of plain pieces, and hence we get the desired $\mathcal{O}(d^3)$ upper bound, as we can bound the number of updates other than insertions/deletions of plain pairs by $\mathcal{O}(d^2)$.

Note that we cannot always iterate explicitly over all $\mathcal{I}'_j$s as this would require $\Omega(m/q)$ calls to `DPM-Query`. We circumvent this problem by observing that if we have $\mathcal{I}'_{j-1} = \mathcal{I}'_j$ (for some $j \in J \setminus \min J$), then $r_j + \mathrm{Occ}^E_k(P, R_j) = q + r_{j-1} + \mathrm{Occ}^E_k(P, R_{j-1})$. Consequently, for any maximal interval $[\, j_1 \mathinner{.\,.} j_2 \,] \subseteq J$ where $\mathcal{I}'_{j_1} = \cdots = \mathcal{I}'_{j_2}$, we only process $\mathcal{I}'_{j_1}$; then, for each position $u \in \mathrm{Occ}^E_k(P, R_{j_1})$, we report an arithmetic progression $\{r_{j_1} + u + iq : i \in [\, 0 \mathinner{.\,.} j_2 - j_1 \,]\}$ of k-error occurrences of P in T. On a high level, we are offloading the computation of $\bigcup^{j_2}_{j=j_1} \left(r_j + \mathrm{Occ}^E_k(P, R_j) \right)$ to the computation of $\mathrm{Occ}^E_k(\mathcal{I}'_{j_1})$.

A Faster Solution.   To obtain a faster solution for the NEWPERIODICMATCHES problem, we intend to trim runs of plain pairs even further, to a length of roughly $\widetilde{\mathcal{O}}(\sqrt{d})$. Now, naively processing the obtained DPM-sequences, we may obtain "false-positive" occurrences, but—as we can prove—not too many. In particular, we can extend existing tools to filter such "false-positive" occurrences.

For a slightly more detailed overview, for any two positions $v < w$ of T, let us write $Q^\infty[\, \rho(v) \mathinner{.\,.} \rho(w) \,)$ for the fragment of Q^∞ that $\mathcal{A}_T$ aligns with $T[\, v \mathinner{.\,.} w \,)$. Suppose that we have

$$\delta_E(P, {}^*Q^*) = \delta_E(P, Q^\infty[\, \rho(v) \mathinner{.\,.} \rho(w) \,)) \leq \delta_E(T[\, v \mathinner{.\,.} w \,), {}^*Q^*) = \delta_E(T[\, v \mathinner{.\,.} w \,), Q^\infty[\, \rho(v) \mathinner{.\,.} \rho(w) \,)).$$

Then, the triangle inequality yields

$$\Lambda := \delta_E(P, {}^*Q^*) + \delta_E(T[\, v \mathinner{.\,.} w \,), {}^*Q^*) \geq \delta_E(P, T[\, v \mathinner{.\,.} w \,)) \geq \delta_E(T[\, v \mathinner{.\,.} w \,), {}^*Q^*) - \delta_E(P, {}^*Q^*).$$

We see that, intuitively, the best case is when all the errors of P with $Q^\infty[\, \rho(v) \mathinner{.\,.} \rho(w) \,)$ cancel out with errors of $T[\, v \mathinner{.\,.} w \,)$ with $Q^\infty[\, \rho(v) \mathinner{.\,.} \rho(w) \,)$. Now, roughly speaking, for each position v of T, we quantify the "potential savings" that an alignment $P \rightsquigarrow T[\, v \mathinner{.\,.} w \,)$ of cost at most k may yield compared to $\min_x \delta_E(P, Q^\infty[\, \rho(v) \mathinner{.\,.} x \,)) + \min_y \delta_E(T[\, v \mathinner{.\,.} w \,), Q^\infty[\, \rho(v) \mathinner{.\,.} y \,)))$. To this end, we use the notion of *locked fragments* from [CKW20] to mark each position of the text with a number of marks proportional to said "potential savings". (A similar notion was used in [CH02].) Based on a threshold $\eta = \widetilde{\Theta}(\sqrt{d})$ on the number of marks (and a few technical conditions), we then classify each position as either *heavy* or *light*. Details on locked fragments and our marking scheme can be found in Section 5.

We then present our solution for NEWPERIODICMATCHES in Section 6. First, we show that the set of heavy positions intersects $\widetilde{\mathcal{O}}(\sqrt{d})$ ranges, each of size $\mathcal{O}(d)$, where a k-error occurrences of P may start (recall that $\mathrm{Occ}^E_k \subseteq \bigcup_{j \in \mathbb{Z}}[\, jq - 5d \mathinner{.\,.} jq + 5d \,]$). We can then compute the intersection of $\mathrm{Occ}^E_k(P, T)$ with heavy positions efficiently, that is, in $\widetilde{\mathcal{O}}(d^{3.5})$ time, using known tools.

Having taken care of the heavy positions, we can return to DYNAMICPUZZLEMATCHING for the light positions. To that end, consider again $T[\, v \mathinner{.\,.} w \,)$, supposing that v is a light position of T. We then have that $\Lambda \geq \delta_E(P, T[\, v \mathinner{.\,.} w \,)) \geq \Lambda - \eta$. Now, the optimal alignment $\mathcal{A}$ from P to $T[\, v \mathinner{.\,.} w \,)$ has to make $\Lambda - \eta$ edit operations just to align the locked fragments of the text and the pattern. This means that the number of edit operations that $\mathcal{A}$ makes in aligning portions of P disjoint from the locked fragments of P to portions of T disjoint from the locked fragments of T is at most η.

Now, we define a set $\mathsf{Red}(P) \supseteq \mathsf{Special}(P)$ that additionally contains all pieces of P that overlap some locked fragment of P; we similarly define a set $\mathsf{Red}(T)$ of pieces of T. Importantly, both $\mathsf{Red}(P)$ and $\mathsf{Red}(T)$ are of size $\mathcal{O}(d)$. Redefining plain pairs to be those that contain no red piece, we show that we can trim each run of plain pairs to have a length of $\mathcal{O}(\eta) = \widetilde{\Theta}(\sqrt{d})$. This allows us to reduce our problem to an instance of the DYNAMICPUZZLEMATCHING problem with $\widetilde{\Theta}(d^{2.5})$ updates in total; as before, we can essentially charge all but $\mathcal{O}(d^2)$ updates to $\mathcal{O}(d^2)$ pairs of runs of plain pairs, so that each such pair gets charged with $\widetilde{\Theta}(\sqrt{d})$ updates.

A Solution for the DynamicPuzzleMatching Problem

For our solution to the DYNAMICPUZZLEMATCHING problem (which we present in Section 9), we rely on a framework of Tiskin [Tis07, Tis08, Tis15] (which we recall and extend in Section 8). A key observation behind this framework is that *semi-local* alignments between strings U and V can be represented as paths between boundary vertices of a certain *alignment graph*: a grid on vertices $[0 \mathinner{.\,.} |V|] \times [0 \mathinner{.\,.} |U|]$, augmented with diagonal edges. All horizontal and vertical edges have weight 1 (they represent insertions and deletions), whereas each diagonal edge $(u,v) \leftrightarrow (u+1, v+1)$ has weight 0 (for a match) or 1 (for a substitution). Then, $\delta_E(V[v \mathinner{.\,.} w), U)$ corresponds to the distance from $(v, 0)$ to $(w, |U|)$. As observed in [Tis08], even though there are quadratically many such distances, they can be encoded in linear space using a certain *permutation matrix* that we denote by $P_{V,U}$. Moreover, we can *stitch* alignment graphs by computing a certain *seaweed product* of permutation matrices. For example, $P_{V,UU'}$ can be expressed as the seaweed product of $P_{V,U}$ and $P_{V,U'}$ (shifted appropriately so that the characters of U' are indexed from $|U|$ rather than from 0). Tiskin [Tis15] provided an $\mathcal{O}(n \log n)$-time algorithm for computing the seaweed product of two $n \times n$ permutation matrices, but we cannot hope to compute $P_{V,U}$ in truly subquadratic time because it encodes $\delta_E(U, V)$.

In our setting, though, the strings U and V are of similar length (that is, $\big||V| - |U|\big| \le \mathrm{tor}(\mathcal{I}) \le \Delta/2 - k$) and we only care about alignments of cost at most k. The underlying paths corresponding to such alignments are fully contained within a narrow *diagonal band* of the alignment graph: all of their vertices (u, v) satisfy $u - v \in I := [-k \mathinner{.\,.} |V| - |U| + k]$ (in short, they belong to band I of the alignment graph); see Figure 4 for an illustration. In order to capture this scenario, we restrict the alignment graph to band I, which corresponds to zeroing out the costs of all diagonal edges outside band I. We prove that the permutation matrix $P_{U,V}|_I$ of the restricted graph can be encoded in $\mathcal{O}(|I|)$ space and computed in $\widetilde{\mathcal{O}}(|I|^2)$ time (in the `PILLAR` model). Moreover, we show that $P_{U,V}|_I$ can be expressed solely in terms of $P_{U,V}$, which leads to a new operation of *restricting* a permutation matrix P to a given interval I. We write $P|_I$ for the result of said operation and we present a linear-time algorithm that computes $P|_I$ directly from P and I.

Let us now explain how these techniques are helpful in solving the DYNAMICPUZZLEMATCHING problem. Our high-level idea is to express $P_{V,U}|_I$ as the seaweed product of z smaller permutation matrices $P_1, \ldots, P_z$, with P_i depending only on the i-th pair (U_i, V_i). For a first attempt, we could use P_{V_i, U_i}, but the corresponding parts of the alignment overlap and thus cannot be stitched easily. Thus, we trim each piece U_i to U_i' so that $U = \mathsf{val}_\Delta(U_1, \ldots, U_z) = U_1' \cdots U_z'$. Now, the seaweed product of matrices $P_{V_i, U_i'}$ (shifted appropriately), restricted *a posteriori* to interval I, yields $P_{V,U}|_I$. However, the individual matrices $P_{V_i, U_i'}$ are still too large, so we need to restrict them *a priori* as well. Thus, we actually use $P_{V_i, U_i'}|_{I_i}$, for appropriate intervals I_i of size at most Δ; see Figure 5 for an illustration. We build a balanced binary tree on top of the permutation matrices $P_{V_i, U_i'}|_{I_i}$ in order to maintain their seaweed product (so that every update requires recomputing $\mathcal{O}(\log z)$ partial products). For each query, we retrieve $P_{V,U}|_I$ and apply the SMAWK algorithm [AKM+87] in order to check, for every $v \in [0 \mathinner{.\,.} |V| - |U| + k]$, whether $\delta_E(U, V[v \mathinner{.\,.} w)) \le k$ holds for some $w \in [|U| - k \mathinner{.\,.} |W|]$ in $\widetilde{\mathcal{O}}(\Delta)$ time in total.

The remaining challenge is to build the matrices $P_{V_i, U_i'}|_{I_i}$. For this, we exploit the small median edit distance of the families $\mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi$ to show that all such matrices can be precomputed in $\widetilde{\mathcal{O}}(d^3 + \Delta^2 d)$ time. If the puzzle pieces were of size $\mathcal{O}(d + \Delta)$, we could simply use an algorithm of Charalampopoulos, Kociumaka, and Mozes [CKM20] that maintains $P_{X,Y}$ subject to edits of X, Y. In general, though, we decompose each piece into $\mathcal{O}(d)$ parts: perfect parts, which can be arbitrarily long but are kept intact among all the puzzle pieces, and imperfect parts, which can contain edits but are of size $\mathcal{O}(\Delta)$. For each perfect part, we compute a single restricted permutation matrix in $\widetilde{\mathcal{O}}(\Delta^2)$ time. For imperfect parts, we use the dynamic algorithm of [CKM20]. Finally, the restricted permutation matrix of a pair of pieces is obtained by stitching the matrices for pairs of parts similarly to how we obtain $P_{V,U}|_I$ from $P_{V_i, U_i'}|_{I_i}$s.
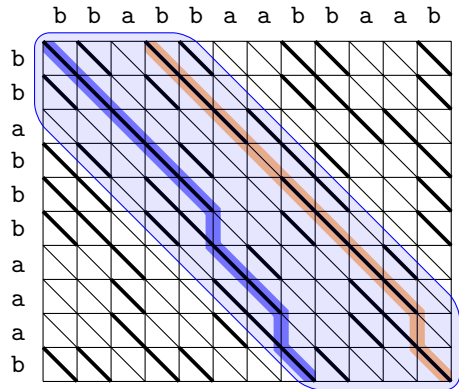
Figure 4 The alignment graph for $U = \texttt{bbabbbaaab}$ and $V = \texttt{bbabbaabbaab}$. Thin edges have cost 1 whereas thick edges have cost 0. The blue and orange path represent cost-2 alignments $U \rightsquigarrow V[\,0\mathbin{.\,.}8\,)$ and $U \rightsquigarrow V[\,3\mathbin{.\,.}12\,)$, respectively. The diagonal band $I = [\,-2\mathbin{.\,.}4\,] = [\,-k\mathbin{.\,.}|V| - |U| + k\,]$ corresponding to $k = 2$ is shaded in blue.
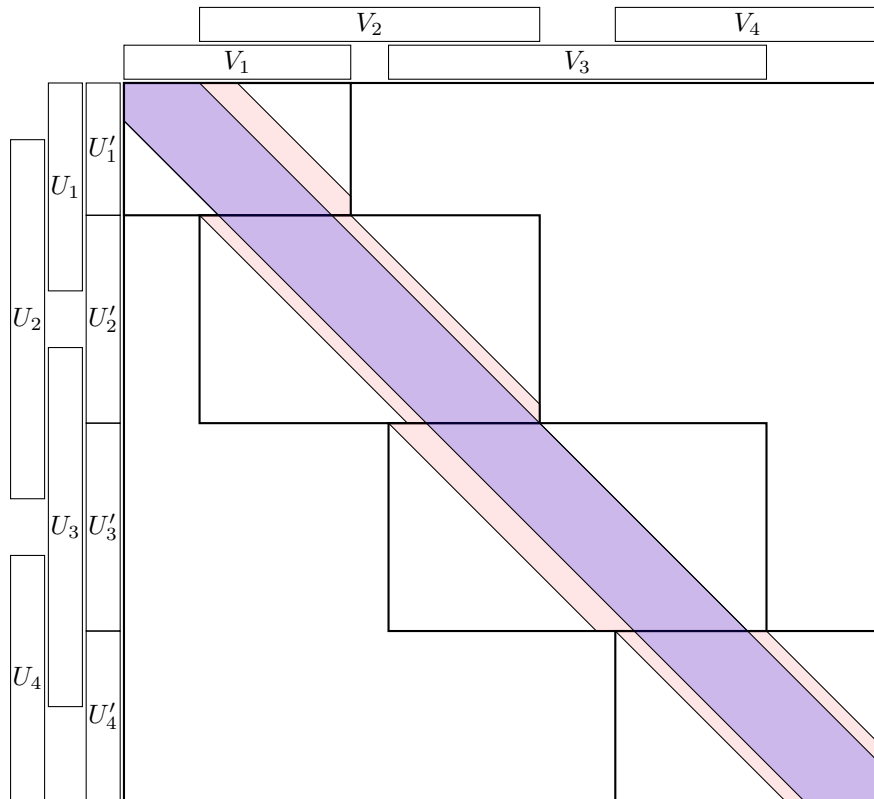
Figure 5 A schematic illustration explaining why $P_{V,U}|_I$, which corresponds to the purple band, can be obtained by the seaweed products of matrices $P_{V_i,U_i'}|_{I_i}$, which correspond to the pink bands within the rectangles representing the alignment graphs of U_i' and V_i (subgraphs of the alignment graph of U and V).

2   Preliminaries

Sets and Arithmetic Progressions

For $i, j \in \mathbb{Z}$, we write $[i..j]$ to denote $\{i, \ldots, j\}$ and $[i..j)$ to denote $\{i, \ldots, j-1\}$; the sets $(i..j]$ and $(i..j)$ are defined similarly.

For integers a, d, and $\ell > 0$, the set $\{a + j \cdot d : j \in [0..\ell)\}$ 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 triples (a, d, ℓ) consisting of their first value, their difference, and their length.

For a set $X \subseteq \mathbb{Z}$ and an integer $s \in \mathbb{Z}$, we write $s + X$ and $X + s$ to denote the set $\{s + x : x \in X\}$ containing all elements of X incremented by s.

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*.

A string P is a *substring* of a string T (denoted by $P \preccurlyeq T$) if for some integers i, j with $0 \leq i \leq j \leq |T|$, we have $P = T[i] \cdots T[j-1]$. 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. We write $T[i..j)$ for this particular occurrence of P in T, which is formally a *fragment* of T specified by the two endpoints i, j. For notational convenience, we may also refer to this fragment as $T[i..j-1]$, $T(i-1..j-1]$, or $T(i-1..j)$. Two fragments (perhaps of different strings) *match* if they are occurrences of the same strings.

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 $j \in [0..|T|)$). 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 $i \in [0..|T|)$). We write $\mathsf{lcp}(U, V)$ for the length of the *longest common prefix* of two strings U and V, that is $\mathsf{lcp}(U, V)$ is the length of the longest string that occurs as a prefix of both U and V. Similarly, we write $\mathsf{lcp}^R(U, V)$ for the length of the *longest common suffix* of two strings U and V, that is $\mathsf{lcp}(U, V)$ is the length of the longest string that occurs as a suffix of both U and V.

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, we write U^∞ to denote an infinite string obtained by concatenating infinitely many copies of U. To simplify our exposition at certain points, we may access such an infinite repetition of U also at negative positions; hence for an integer $j \in [0..|Q|)$ and a (possibly negative) integer i, we formally set $U^\infty[i \cdot |U| + j] := U[j]$. A string T is *primitive* if it cannot be expressed as $T = U^k$ for any string U and any integer $k > 1$.

A positive integer p is 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* $\mathsf{per}(T)$ of the string. Further, we call $T[0..\mathsf{per}(T))$ the *string period* of T. A string is *periodic* if its period is at most half of its length.

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

For a string T, the *reverse string* of T is $T[n-1]T[n-2] \cdots T[0]$.

Edit Distance and Pattern Matching with Edits

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

For a formal definition, we first define an *alignment* between strings.

▶ **Definition 2.1.** *A sequence $\mathcal{A} = (x_i, y_i)_{i=0}^m$ is an* alignment *of $X[x \mathinner{.\,.} x')$ onto $Y[y \mathinner{.\,.} y')$, denoted by $\mathcal{A} : X[x \mathinner{.\,.} x') \rightsquigarrow Y[y \mathinner{.\,.} y')$, if we have*

$$(x_0, y_0) = (x, y); \quad and$$
$$for\ all\ i \in [0 \mathinner{.\,.} m): \quad (x_{i+1}, y_{i+1}) \in \{(x_i + 1, y_i + 1), (x_i + 1, y_i), (x_i, y_i + 1)\}; \quad and$$
$$(x_m, y_m) = (x', y').$$

- *If $(x_{i+1}, y_{i+1}) = (x_i + 1, y_i)$, we say that $\mathcal{A}$ deletes $X[x_i]$,*
- *If $(x_{i+1}, y_{i+1}) = (x_i, y_i + 1)$, we say that $\mathcal{A}$ inserts $Y[y_i]$,*
- *If $(x_{i+1}, y_{i+1}) = (x_i + 1, y_i + 1)$, we say that $\mathcal{A}$ aligns $X[x_i]$ and $Y[y_i]$. If additionally $X[x_i] = Y[y_i]$, we say that $\mathcal{A}$ matches $X[x_i]$ and $Y[y_i]$. Otherwise, we say that $\mathcal{A}$ substitutes $Y[y_i]$ for $X[x_i]$.* ◀

Further, for an alignment $\mathcal{A} : X[x \mathinner{.\,.} x') \rightsquigarrow Y[y \mathinner{.\,.} y')$ with $\mathcal{A} = (x_i, y_i)_{i=0}^m$, we define the *inverse alignment* $\mathcal{A}^{-1} : Y[y \mathinner{.\,.} y') \rightsquigarrow X[x \mathinner{.\,.} x')$ as $\mathcal{A}^{-1} := (y_i, x_i)_{i=0}^m$. The *cost* of an alignment $\mathcal{A}$ of $X[x \mathinner{.\,.} x')$ onto $Y[y \mathinner{.\,.} y')$, denoted by $\delta_E^{\mathcal{A}}(X[x \mathinner{.\,.} x'), Y[y \mathinner{.\,.} y'))$, is the total number of characters that $\mathcal{A}$ inserts, deletes, or substitutes. Now, we define the edit distance $\delta_E(X, Y)$ as the minimum cost of an alignment of $X[0 \mathinner{.\,.} |X|)$ onto $Y[0 \mathinner{.\,.} |Y|)$. An alignment of X onto Y is *optimal* if its cost is equal to $\delta_E(X, Y)$. The deletion distance $\delta_D(X, Y)$ is defined as the minimum cost of an alignment $\mathcal{A} : X \rightsquigarrow Y$ that aligns $X[x]$ to $Y[y]$ only if the two characters match.

Given an alignment $\mathcal{A} : X[x \mathinner{.\,.} x') \rightsquigarrow Y[y \mathinner{.\,.} y')$ and a fragment $X[\bar{x} \mathinner{.\,.} \bar{x}') \preccurlyeq X[x \mathinner{.\,.} x')$, we write $\mathcal{A}(X[\bar{x} \mathinner{.\,.} \bar{x}'))$ for the fragment $Y[\bar{y} \mathinner{.\,.} \bar{y}') \preccurlyeq Y[y \mathinner{.\,.} y')$ that $\mathcal{A}$ aligns against $X[\bar{x} \mathinner{.\,.} \bar{x}')$. Given that insertions and deletions may render this definition ambiguous, we formally set

$$\bar{y} := \min\{\hat{y} : (\bar{x}, \hat{y}) \in \mathcal{A}\} \quad and \quad \bar{y}' := \begin{cases} y' & if\ \bar{x}' = x', \\ \min\{\hat{y}' : (\bar{x}', \hat{y}') \in \mathcal{A}\} & otherwise. \end{cases}$$

This particular choice satisfies the following decomposition property.

▶ **Fact 2.2.** *For any alignment $\mathcal{A}$ of X onto Y and a decomposition $X = X_1 \cdots X_t$ into t fragments, $Y = \mathcal{A}(X_1) \cdots \mathcal{A}(X_t)$ is a decomposition into t fragments with $\delta_E^{\mathcal{A}}(X, Y) \geq \sum_{i=1}^t \delta_E(X_i, \mathcal{A}(X_i))$.*
Further, if $\mathcal{A}$ is an optimal alignment, then have equality: $\delta_E^{\mathcal{A}}(X, Y) = \sum_{i=1}^t \delta_E(X_i, \mathcal{A}(X_i))$. ◀

Consider Figure 6 for a visualization of an example.

▶ **Definition 2.3.** *For an alignment $\mathcal{A} : X[x \mathinner{.\,.} x') \rightsquigarrow Y[y \mathinner{.\,.} y')$, and a pair $(x_i, y_i) \in \mathcal{A}$ where $\mathcal{A}$ does not match x_i and y_i, we define the corresponding* augmented breakpoint B *as*

$$B := \begin{cases} (x_i, y_i, \mathtt{INS}), & if\ \mathcal{A}\ inserts\ Y[y_i], \\ (x_i, y_i, \mathtt{DEL}), & if\ \mathcal{A}\ deletes\ X[x_i], \\ (x_i, y_i, \mathtt{SUB}), & if\ \mathcal{A}\ substitutes\ X[x_i],\ and \\ (x_i, y_i, \bot), & if\ x_i = x'\ and\ y_i = y'\ or\ if\ x_i = x - 1\ and\ y_i = y - 1. \end{cases}$$

We write $\mathsf{B}_{X,Y}(\mathcal{A})$ for the set of augmented breakpoints. ◀

$X =$		a		a		c		$-$		b		c		d
		$*$		$\parallel$		$\parallel$		$*$		$\parallel$		$\parallel$		$*$
$Y =$		b		a		c		d		b		c		$-$
$\mathcal{A} =$	$(0,0)$		$(1,1)$		$(2,2)$		$(3,3)$		$(3,4)$		$(4,5)$		$(5,6)$		$(6,6)$

Figure 6 Consider strings $X = $ `aacbcd` and $Y = $ `bacdbc`. We have $\delta_E(X,Y) = 3$, witnessed by an optimal alignment $\mathcal{A} : X \rightsquigarrow Y$. The alignment $\mathcal{A}$ inserts, deletes, or substitutes three characters; these edit operations are denoted by red asterisks. The set of (augmented) breakpoints of $\mathcal{A}$ is $\mathsf{B}_{X,Y}(\mathcal{A}) = \{(-1,-1,\bot), (0,0,\mathtt{SUB}), (3,3,\mathtt{DEL}), (5,6,\mathtt{INS}), (6,6,\bot)\}$. Further, we have $\mathcal{A}(X[0..3)) = Y[0..3) = $ `bac`, $\mathcal{A}(X[3..5)) = Y[3..5) = $ `dbc`, and $\mathcal{A}(X[5..6)) = Y[6..6) = \varepsilon$.

Algorithm 1 Pseudo code for Lemma 2.4.

---

1   $\mathtt{Reverse}(\mathsf{B}_{X,Y}(\mathcal{A}) = \{(x-1, y-1, \bot), \ldots, (x', y', \bot)\})$

2      $px \leftarrow x;\ py \leftarrow y;$

3   **foreach** $B \in \mathsf{B}_{X,Y}(\mathcal{A})$ **do**

4   **switch** B **do**

5   **case** $(a, b, \mathtt{INS})$ **do**

6   **yield** $[\sigma_{px}^{\mathcal{A}} .. \sigma_a^{\mathcal{A}}) \leftarrow [px .. a);$

7   **yield** $[\lambda_{py}^{\mathcal{A}} .. \lambda_b^{\mathcal{A}}] \leftarrow [py .. b];$

8               $px \leftarrow a;\ py \leftarrow b + 1;$

9   **case** $(a, b, \mathtt{DEL})$ **do**

10   **yield** $[\sigma_{px}^{\mathcal{A}} .. \sigma_a^{\mathcal{A}}] \leftarrow [px .. a];$

11   **yield** $[\lambda_{py}^{\mathcal{A}} .. \lambda_b^{\mathcal{A}}) \leftarrow [py .. b);$

12               $px \leftarrow a + 1;\ py \leftarrow b;$

13   **otherwise do**

14   **continue**;

15   **yield** $[\sigma_{px}^{\mathcal{A}} .. \sigma_{x'}^{\mathcal{A}}) \leftarrow [px .. x');$

16   **yield** $[\lambda_{py}^{\mathcal{A}} .. \lambda_{y'}^{\mathcal{A}}) \leftarrow [py .. y');$

Observe that $\delta_E^{\mathcal{A}}(X,Y) = |\mathsf{B}_{X,Y}(\mathcal{A})| - 2$, with the -2 term corresponding to $(x', y', \bot) \in \mathsf{B}_{X,Y}(\mathcal{A})$ and $(x - 1, y - 1, \bot) \in \mathsf{B}_{X,Y}(\mathcal{A})$.

Moreover, observe that we can uniquely reconstruct $\mathcal{A}$ from $\mathsf{B}_{X,Y}(\mathcal{A})$. Our algorithms use $\mathsf{B}_{X,Y}(\mathcal{A})$ (with elements stored in a sorted array) to represent the alignment $\mathcal{A}$; if the cost of $\mathcal{A}$ is d, then this *breakpoint representation* takes $\mathcal{O}(d + 1)$ space.[6]

▸ **Lemma 2.4.** *For an alignment $\mathcal{A} : X[x..x') \rightsquigarrow Y[y..y')$ and a position $\bar{y} \in [y..y']$, write $[\sigma_{\bar{y}}^{\mathcal{A}} .. \lambda_{\bar{y}}^{\mathcal{A}}] := \{\bar{x} \in [x..x'] : (\bar{x}, \bar{y}) \in \mathcal{A}\}$ for the corresponding positions in X under $\mathcal{A}$.*

Given the breakpoint representation of a cost-d alignment $\mathcal{A}$, in $\mathcal{O}(d + 1)$ time, we can compute the sequences $(\sigma_{\bar{y}}^{\mathcal{A}})_{\bar{y}=y}^{y'}$ and $(\lambda_{\bar{y}}^{\mathcal{A}})_{\bar{y}=y}^{y'}$, represented as concatenations of $d + 1$ arithmetic progressions with difference 1.

Proof. First, observe that for a fixed $\bar{y} \in [y..y']$, we have $\sigma_{\bar{y}}^{\mathcal{A}} \neq \lambda_{\bar{y}}^{\mathcal{A}}$ if and only if $\mathcal{A}$ deletes $X[\sigma_{\bar{y}}^{\mathcal{A}} .. \lambda_{\bar{y}}^{\mathcal{A}})$. Similarly, multiple positions $y_i, y_{i+1}, \ldots, y_{i+a}$ share a common value of $\sigma_{y_\star}^{\mathcal{A}} = \lambda_{y_\star}^{\mathcal{A}}$ if and only if $\mathcal{A}$ inserts

---

[6] Observe that the breakpoint representation requires constant (non-zero) space for alignments with cost $d = 0$.

the fragment $X[\,y_i\mathbin{.\,.}y_{i+a-1}\,)$. Hence, to compute the values $\sigma_{\bar{y}}^{\mathcal{A}}$ and $\lambda_{\bar{y}}^{\mathcal{A}}$ for all $\bar{y}$, it suffices to do a linear pass over the (augmented) breakpoints: we keep a pointer p_X in X and a pointer p_y in Y, each pointing to the character after the last insertion or deletion of $\mathcal{A}$ (or to $X[\,x\,]$ and $Y[\,y\,]$ at the beginning). Now, when we encounter an insertion $(a, b, \texttt{INS})$, deletion $(a, b, \texttt{DEL})$ or the last breakpoint $(a = x', b = y', \perp)$ in $\mathsf{B}_{X,Y}(\mathcal{A})$, we set

$$[\,\sigma_{p_x}^{\mathcal{A}}\mathbin{.\,.}\sigma_a^{\mathcal{A}}\,] := [\,p_x\mathbin{.\,.}a\,] \quad \text{and} \quad [\,\lambda_{p_y}^{\mathcal{A}}\mathbin{.\,.}\lambda_b^{\mathcal{A}}\,] := [\,p_y\mathbin{.\,.}b\,].$$

Afterward, we update x_p and y_p accordingly. Consult Algorithm 1 for a detailed pseudo-code.

The correctness follows immediately from the preceding discussion; for the running time, observe that at each of the $d+1$ events, our algorithm requires constant time. In total, we hence obtain the claim. ▪

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

Next, it is easy to verify that the edit distance satisfies the triangle inequality.

�------ **Fact 2.5** (Triangle Inequality). *Any strings A, B, and C satisfy*

$$\delta_E(A, C) + \delta_E(C, B) \geq \delta_E(A, B) \geq |\delta_E(A, C) - \delta_E(C, B)|.$$ ▪

Equally useful is the fact that we can easily remove prefixes (or suffixes) of the same length.

�------ **Fact 2.6**. *For any non-empty strings A and B, we have*

$$\delta_E(A[\,1\mathbin{.\,.}|A|\,), B[\,1\mathbin{.\,.}|B|\,)) \leq \delta_E(A, B) \quad \text{and} \quad \delta_E(A[\,0\mathbin{.\,.}|A|-1\,), B[\,0\mathbin{.\,.}|B|-1\,)) \leq \delta_E(A, B).$$ ▪

Combining Facts 2.2 and 2.6, we see that we can "cut" fragments of equal length out of an optimal alignment.

�------ **Corollary 2.7**. *Let $A = A_p \cdot A_s$ and B denote non-empty strings and write $\mathcal{A} : A \rightsquigarrow B$ for an optimal alignment of A onto B. For any $u \in [\,0\mathbin{.\,.}\min\{|A_s|, |\mathcal{A}(A_s)|\}\,)$, we have*

$$\delta_E(A_p, \mathcal{A}(A_p)) + \delta_E(A_s[\,u\mathbin{.\,.}|A_s|\,), \mathcal{A}(A_s)[\,u\mathbin{.\,.}|\mathcal{A}(A_s)|\,)) \leq \delta_E(A, B).$$ ▪

We conclude with the easy, but useful, observation that no k-error occurrence of a pattern may start in the final part of the text.

�------ **Fact 2.8**. *For any text T of length n and any pattern P of length m we have $\mathrm{Occ}_k^E(P, T) \cap (\,n - m + k\mathbin{.\,.}n\,) = \varnothing$.* ▪

2.1 The `PILLAR` Model

To unify the implementations of algorithms for (approximate) pattern matching problems in different settings, [CKW20] introduced the `PILLAR` model. We use the same `PILLAR` model in this work. In particular, we bound the running times of the algorithms in this work in terms of the number of calls to a small set of very common operations (the primitive `PILLAR` operations) on strings.[7] Together with

---

[7] If our algorithms require (asymptotically) significant extra computations, we also specify the required extra running time.

the implementations of said primitive operations (presented in detail in [CKW20]), we then obtain (fast) algorithms for various different settings at once. This includes the standard setting, the fully compressed setting, and a dynamic setting. We provide more details in Section 7.

To keep the PILLAR model flexible, we do not directly work on specific representations of (the input) strings. Instead, in the PILLAR model, we maintain a collection of strings $\mathcal{X}$; the operations in the PILLAR model work on fragments $X[\ell \mathinner{.\,.} r)$ of $X \in \mathcal{X}$, which are represented via a *handle*.[8] At the start of the computation, the PILLAR model provides a handle to each $X \in \mathcal{X}$, which represents $X[0 \mathinner{.\,.} |X|)$. Using an Extract operation, we can obtain handles to other fragments of the strings in $\mathcal{X}$ [CKW20]:

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

The other primitive PILLAR model operations read as follows [CKW20]:

- 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 $|T| \le 2|P|$, compute $\mathrm{Occ}(P, T)$ (represented as an arithmetic progression with difference $\mathrm{per}(P)$).
- Access(S, i): Assuming $i \in [0 \mathinner{.\,.} |S|)$, retrieve the character $S[i]$.
- Length(S): Retrieve the length $|S|$ of the string S.

As working just with the primitive PILLAR model operations is rather bothersome, [CKW20] also provides a useful toolbox of operations already implemented. For brevity, we only list these operations here and refer to [CKW20] for (pointers to) their implementation.

⬚ **Fact 2.9** (PILLAR Toolbox, [CKW20]).   *The PILLAR model supports all of the following operations.*

- *Equality [Koc18, Fact 2.5.2]: For strings S and T, we can check whether S and T are equal in $\mathcal{O}(1)$ time in the PILLAR model.*
- Period(S), *[KRRW15, Koc18]: For a string S, we can compute $\mathrm{per}(S)$ or declare that $\mathrm{per}(S) > |S|/2$ in $\mathcal{O}(1)$ time in the PILLAR model.*
- Rotations(S, T), *[KRRW15, Koc18]: For strings S and T, we can find all integers j such that $T = \mathrm{rot}^j(S)$ in $\mathcal{O}(1)$ time in the PILLAR model. The output is represented as an arithmetic progression.*
- LCP$(S, Q^\infty[\ell \mathinner{.\,.} r))$, *[CKW20], see also [Koc18, Fact 2.5.2] and [BGK$^+$16]: For strings S and Q and integers $0 \le \ell \le r$, we can compute LCP$(S, Q^\infty[\ell \mathinner{.\,.} r))$ in $\mathcal{O}(1)$ time in the PILLAR model.*
- LCP$^R(S, Q^\infty[\ell \mathinner{.\,.} r))$, *[CKW20]: For strings S and Q and integers $0 \le \ell \le r$, we can compute LCP$^R(S, Q^\infty[\ell \mathinner{.\,.} r))$ in $\mathcal{O}(1)$ time in the PILLAR model.*
- ExactMatches(P, T), *[CKW20]: Let T denote a string of length n and let P denote a string of length m. We can compute the set $\mathrm{Occ}(P, T)$ using $\mathcal{O}(n/\mathrm{per}(P))$ time and $\mathcal{O}(n/m)$ PILLAR operations.*
- Verify(P, T, k, I), *[CH02, Section 5]; see also [CKW20]: Let P denote a string of length m, let T denote a string, and let $k \le m$ denote a positive integer. Further, let I denote an interval of positive integers. Using $\mathcal{O}(k(k + |I|))$ PILLAR operations, we can compute $\{(\ell, \min_r \delta_E(P, T[\ell \mathinner{.\,.} r))) : \ell \in \mathrm{Occ}_k^E(P, T) \cap I\}$.* ⬛

Finally, we observe that the adaptation of the Landau–Vishkin algorithm [LV89] already provided in [CKW20] lets us efficiently compute an optimal alignment of a string S onto a substring of Q^∞ starting at a given position $x \in \mathbb{Z}$.

---

[8] The implementation details depend on the specific setting. In the standard setting, a fragment $X[\ell \mathinner{.\,.} r)$ is represented by a reference to X and the endpoints ℓ, r.

▟ **Lemma 2.10** (Alignment(S, Q, x)). *Consider non-empty strings S, Q and an integer $x \in \mathbb{Z}$. Using* $\mathcal{O}(1 + d^2)$ PILLAR *operations, we can construct (the breakpoint representation of) an alignment $\mathcal{A} :$* $S \rightsquigarrow Q^\infty[\, x \mathinner{\ldotp\ldotp} y\,)$ *of optimal cost $d = \delta_E(S, \mathrm{rot}^{-x}(Q)^*)$.*

Proof. Set $\hat{Q} := \mathrm{rot}^{-x}(Q)$. [CKW20, Lemma 6.1] provides an online algorithm that, for subsequent integers $k \in \mathbb{Z}_{\geq 0}$, computes the longest prefix S_k of S such that $\delta_E(S_k, \hat{Q}^*) \leq k$, as well as a witness length $\ell_k \in \mathbb{Z}_{\geq 0}$ and alignment $\mathcal{A}_k : S_k \rightsquigarrow \hat{Q}[\, 0 \mathinner{\ldotp\ldotp} \ell_k\,)$ of cost at most k. We run this algorithm until $S_k = S$ holds for the first time, which indicates that $k = \delta_E(S, \hat{Q}^*) = d$, and return the final alignment $\mathcal{A}_d$, reinterpreted as an alignment $\mathcal{A} : S \rightsquigarrow Q[\, x \mathinner{\ldotp\ldotp} y\,)$ for $y := x + \ell_d$.

A minor subtlety is that [CKW20, Lemma 6.1] provides a slightly weaker representation of the alignment $\mathcal{A}_d$, with (i, j) for an substitution of $S[\, i\,]$ to $\hat{Q}[\, j\,]$, $(i, \perp)$ for a deletion of $S[\, i\,]$, and $(\perp, j)$ for an insertion of $\hat{Q}[\, j\,]$. Nevertheless, a left-to-right scan of this representation, keeping track of the shift δ, equal to the number insertions processed so far minus number of deletions processed so far, lets us produce the breakpoints of $\mathcal{A}$, that is, $(-1, x - 1, \perp)$ at the beginning, $(i, j + x, \mathtt{SUB})$ for each substitution (i, j), $(i, i + \delta + x, \mathtt{DEL})$ for each deletion $(i, \perp)$, $(j - \delta, j + x, \mathtt{INS})$ for each insertion $(\perp, j)$, and $(|S|, y, \perp)$ at the end.

In the PILLAR model, the algorithm of [CKW20, Lemma 6.1] takes $\mathcal{O}(1)$ preprocessing time and $\mathcal{O}(1 + k)$ time for the kth step (for $k \in \mathbb{Z}_{\geq 0}$). Thus, running it for all $k \in [\, 0 \mathinner{\ldotp\ldotp} d\,]$ costs $\mathcal{O}(1 + d^2)$ time in total. The post-processing of the alignment takes $\mathcal{O}(1 + d)$ extra time. ▟

2.2 An Overview of an $\mathcal{O}(k^4)$-Time Algorithm for Pattern Matching with Edits in the PILLAR Model

Before we discuss the new algorithm for pattern matching with edits, we briefly recall the algorithm presented in [CKW20] that runs in $\mathcal{O}(k^4)$ time in the PILLAR model when $n = \mathcal{O}(m)$.[9] Thereby, we can naturally introduce required notations and concepts. Further, we highlight the bottleneck in the aforementioned algorithm; all other parts run in $\mathcal{O}(k^3)$ time in the PILLAR model, which translates to a running time of $\mathcal{O}(n + k^3)$ in the standard setting.

Toward obtaining algorithms for PMwithEdits, first observe that it suffices to focus on a bounded-ratio version of said problem, where $n < \frac{3}{2} m + k$ (consult [CKW20] for a rigorous proof): Given a text of arbitrary length n, we can reduce the problem to $\lfloor 2n/m \rfloor$ instances of the bounded-ratio version of the problem using the so-called *standard trick* [Abr87]. That is, we consider the overlapping fragments $T_i := T[\, \lfloor i \cdot m/2 \rfloor \mathinner{\ldotp\ldotp} \min\{n, \lfloor (i + 3) \cdot m/2 \rfloor + k - 1\}\,)$, for $i \in [\, 0 \mathinner{\ldotp\ldotp} \lfloor 2n/m \rfloor - 1\,)$, and compute the approximate occurrences of P in each of $T_0, \ldots, T_{\lfloor 2n/m \rfloor - 1}$. Afterward, we straightforwardly merge the obtained partial results. Thus, an algorithm that solves the bounded-ratio version of the approximate pattern matching problem in time $f(k)$ in the PILLAR model, directly yields an algorithm that solves an arbitrary instance of the approximate pattern matching problem in time $\mathcal{O}(n/m) \cdot f(k)$ in the PILLAR model.

Next, as a first step in solving the bounded-ratio version of PMwithEdits, we analyze the pattern according to [CKW20, Lemma 6.4].[10]

▟ **Lemma 2.11** (Analyze$(P,\ k)$, [CKW20, Lemma 6.4]). *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 $B_1, \ldots, B_{2k} \preccurlyeq P$, each having period $\mathrm{per}(B_i) > m/128k$ and length $|B_i| = \lfloor m/8k \rfloor$.*

---

[9] This translates to a running time of $\mathcal{O}(n + k^4)$ in the standard setting (under the assumption that $n = \mathcal{O}(m)$).

[10] When citing specific statements of [CKW20], we do so using their numbering in the full (arXiv) version of the paper.

(b) *Disjoint repetitive regions $H_1, \ldots, H_r \preccurlyeq P$ of total length $\sum_{i=1}^{r} |H_i| \geq 3/8 \cdot m$ such that each region H_i satisfies $|H_i| \geq m/8k$ and is constructed along with a primitive approximate period Q_i such that $|Q_i| \leq m/128k$ and $\delta_E(H_i, {}^*Q_i^*) = \lceil 8k/m \cdot |H_i| \rceil$.*[11]

(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.*   ◿

For the case where the analysis of the pattern using Lemma 2.11 yields an approximate period Q, we use the algorithm encapsulated in the following lemma.

⚑ **Fact 2.12** (`PeriodicMatches(P, T, k, d, Q)`, [CKW20, Lemma 6.11]).   *Let P denote a pattern of length m and let T denote a text of length n. Further, let $k \in [0 \mathinner{.\,.} m]$ 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 $\delta_E(P, {}^*Q^*) \leq d$.*

Then, we can compute a representation of the set $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(n/m \cdot d^3)$ disjoint arithmetic progressions with difference $|Q|$ using $\mathcal{O}(n/m \cdot d^4)$ time in the `PILLAR` *model.*[12]   ◿

As the main contribution of this work, we reduce the running time in Fact 2.12 to $\widetilde{\mathcal{O}}(n/m \cdot d^{3.5})$.

For the case where the analysis of the pattern using Lemma 2.11 yields $2k$ disjoint breaks, we use the efficient algorithm encapsulated in the following lemma.

⚑ **Fact 2.13** (`BreakMatches(P, T, {B₁,…,B₂ₖ}, k)`, [CKW20, Lemmas 5.21 and 6.12]).   *Let k denote a threshold and let P denote a pattern of length m having $2k$ disjoint breaks $B_1, \ldots, B_{2k} \preccurlyeq P$ each satisfying $\mathrm{per}(B_i) \geq m/128k$. Further, let T denote a text of length $n < \frac{3}{2}m + k$. Then, we can compute the set $\mathrm{Occ}_k^E(P, T)$, which is of size $\mathcal{O}(k^2)$, using $\mathcal{O}(k^3)$ time in the* `PILLAR` *model.*   ◿

Finally, let us consider the case where the analysis of the pattern using Lemma 2.11 returns disjoint repetitive regions $H_1, \ldots, H_r \preccurlyeq P$ of total length $\sum_{i=1}^{r} |H_i| \geq \frac{3}{8}m$ such that each region H_i satisfies $|H_i| \geq m/8k$ and is constructed along with a primitive approximate period Q_i such that $|Q_i| \leq m/128k$ and $\delta_E(H_i, {}^*Q_i^*) = \lceil 8k/m \cdot |H_i| \rceil$.

In this case, we give a brief overview of the proof of [CKW20, Lemmas 5.24 and 6.13], stated below, showing that the problem in scope reduces to several calls to the `PeriodicMatches` procedure of Fact 2.12.

⚑ **Fact 2.14** (`RepetitiveMatches(P,T,{(H₁,Q₁),…,(Hᵣ,Qᵣ)},k)`, [CKW20, Lemmas 5.24 and 6.13]). *Let P denote a pattern of length m and let $k \leq m$ denote a threshold. Further, let T denote a string of length $n < \frac{3}{2}m + k$. Suppose that P contains disjoint repetitive regions $H_1, \ldots, H_r$ of total length at least $\sum_{i=1}^{r} |H_i| \geq \frac{3}{8}m$ such that each region H_i satisfies $|H_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $\delta_E(H_i, Q_i^*) = \lceil 8k/m \cdot |H_i| \rceil$.*

Then, we can compute the set $\mathrm{Occ}_k^E(P, T)$, which is of size $\mathcal{O}(k^2)$, using $\mathcal{O}(k^4)$ time in the `PILLAR` *model.*   ◿

Proof sketch. For each repetitive region H_i, set $k_i := \lfloor 4 \cdot k/m \cdot |H_i| \rfloor$ and $d_i := \lceil 8 \cdot k/m \cdot |H_i| \rceil$. We can compute $\mathrm{Occ}_{k_i}^E(H_i, T)$ using a call `PeriodicMatches(Hᵢ, T, kᵢ, dᵢ, Q)`; consult [CKW20] for a rigorous proof that the conditions in the statement of Fact 2.12 are indeed satisfied.

Next, using the sets $\mathrm{Occ}_{k_i}^E(H_i, T)$ we can identify in $\mathcal{O}(k^2 \log \log k)$ time $\mathcal{O}(k)$ length-k intervals whose union is a superset of $\mathrm{Occ}_k^E(P, T)$. Finally, we use `Verify(P, T, k, J)` for each such interval J to filter out false-positive positions; these calls to `Verify` take $\mathcal{O}(k^3)$ time in total in the `PILLAR` model.   ◿

All in all, we see that the pattern matching with edits problem reduces to several calls to the `PeriodicMatches` procedure:

---

[11] Observe that we renamed the repetitive regions to $H_\star$, as we use $R_\star$ later with a different meaning.

[12] The representation of $\mathrm{Occ}_k^E(P, T)$ is not stated explicitly in [CKW20, Lemma 6.11].

- If we can detect breaks in the pattern, no calls to `PeriodicMatches` are needed.
- If the pattern is close to being periodic, a single call `PeriodicMatches(P, T, k, d, Q)` suffices. In this case, we have $d = \mathcal{O}(k)$ and hence we obtain a representation of $\mathrm{Occ}_k^E(P, T)$ that consists of $\mathcal{O}(k^3)$ disjoint arithmetic progressions with difference $|Q|$.
- If P contains disjoint repetitive regions $H_1, \ldots, H_r$ of total length at least $\sum_{i=1}^r |H_i| \geq 3/8 \cdot m$ such that each region H_i satisfies $|H_i| \geq m/8k$ and has a primitive approximate period Q_i with $|Q_i| \leq m/128k$ and $d_i := \delta_E(H_i, Q_i^*) = \lceil 8k/m \cdot |H_i| \rceil$. Then, we make r calls `PeriodicMatches`(H_i, T, k_i, d_i, Q_i).

Not that if the pattern is not close to being periodic, the set Occ_k^E is of size $\mathcal{O}(k^2)$ due to [CKW20, Lemmas 5.21 and 5.24]. Thus, in all cases, we obtain a representation of $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(k^3)$ arithmetic progressions with the same difference. Observe that the common difference of arithmetic progressions is solely dependent on P and k: it is $|Q|$ if Lemma 2.11 returns an approximate period Q, while it can be set to an arbitrary positive integer otherwise as the $\mathcal{O}(k^2)$ positions of $\mathrm{Occ}_k^E(P, T)$ are interpreted as arithmetic progressions of length one in this case. We summarize the above discussion in the following statement.

Fact 2.15. *Let P denote a pattern of length m and let $k \leq m$ denote a threshold. Further, let T denote a string of length $n < {}^3\!/_2\, m + k$.*

We can compute a representation of $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(k^3)$ disjoint arithmetic progressions with the same difference in $\mathcal{O}(k^3)$ time in the `PILLAR` model and additionally several calls `PeriodicMatches`(P_i, T, k_i, d_i, Q_i) such that $\sum_i |P_i| \leq m$ and $d_i = \lceil 8k/m \cdot |P_i| \rceil$. The common difference of said arithmetic progressions can be computed given only P and k.

Part I: From NewPeriodicMatches to DynamicPuzzleMatching

3   The NewPeriodicMatches Problem

As a first step toward an improved implementation of the `PeriodicMatches` procedure, we recall useful tools from [CKW20]—this also allows us to slightly simplify the problem statement of the problem we need to solve. Then, we also give a first algorithm that solves said problem fast in a (very restricted) special case; however, this algorithm turns out to be useful in general as well.

3.1   Computing Occurrences in the Periodic Case: Preprocessing and Simplifications

In this (sub-)section, we exploit Fact 2.15 to reduce an instance of the PMWITHEDITS to several instances of the NEWPERIODICMATCHES problem defined in Section 1.3, which we restate here for convenience.

NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$
Input: A pattern P of length m, an integer threshold $k \in [\,0 \mathinner{.\,.} m\,]$, a positive integer $d \geq 2k$, a text T of length $n \in [\,m - k \mathinner{.\,.} \lceil 3/2\, m \rceil + k\,)$, a primitive string Q of length $q := |Q| \leq m/8d$, an edit-distance alignment $\mathcal{A}_P : P \rightsquigarrow Q^\infty[\,0 \mathinner{.\,.} y_P\,)$ of cost $d_P := \delta_E(P, {}^*Q^*) = \delta_E(P, Q^*) \leq d$, and an edit-distance alignment $\mathcal{A}_T : T \rightsquigarrow Q^\infty[\,x_T \mathinner{.\,.} y_T\,)$ of cost $d_T := \delta_E(T, {}^*Q^*) \leq 3d$, where $x_T \in [\,0 \mathinner{.\,.} q\,)$.
Output: The set $\mathrm{Occ}_k^E(P, T)$ represented as $\mathcal{O}(d^3)$ disjoint arithmetic progressions with difference q.

Sepcifically, we prove the following reduction.

⬛ **Fact 1.1.** *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.*

We can compute a representation of the set $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(n/m \cdot k^3)$ disjoint arithmetic progressions with the same difference in time $\mathcal{O}(n/m \cdot k^3)$ in the `PILLAR` model plus the time required for solving several instances NEWPERIODICMATCHES$(P_i, T_i, k_i, d_i, Q_i, \mathcal{A}_{P_i}, \mathcal{A}_{T_i})$, where $\sum_i |P_i| = \mathcal{O}(n)$ and, for each i, we have $|P_i| \leq m$ and $d_i = \lceil 8k/m \cdot |P_i| \rceil$. ⬛

As before, we use the standard trick to reduce the problem to $\mathcal{O}(n/m)$ bounded-ratio instances: specifically, for each $j \in [\,0 \mathinner{.\,.} \lfloor 2n/m \rfloor - 1\,)$, an instance PMWITHEDITS(P, T_j, k) with $|T_j| < 3/2\, m + k$ and, without loss of generality, $|T_j| \geq m - k$ (otherwise, $\mathrm{Occ}_k^E(P, T_j) = \varnothing$).

Let us fix some $j \in [\,0 \mathinner{.\,.} \lfloor 2n/m \rfloor - 1\,)$. We apply Fact 2.15 to reduce PMWITHEDITS(P, T_j, k), in $\mathcal{O}(k^3)$ time in the `PILLAR` model, to several calls `PeriodicMatches`$(P_{j,i}, T_j, k_{j,i}, d_{j,i}, Q_{j,i})$ with $\sum_i |P_{j,i}| \leq m$ and $d_{j,i} = \lceil 8k/m \cdot |P_{j,i}| \rceil$. Over all j, the total length of the patterns in the constructed instances of NEWPERIODICMATCHES is thus $\sum_{j=0}^{\lfloor 2n/m \rfloor - 2} \sum_i |P_{j,i}| \leq \sum_{j=0}^{\lfloor 2n/m \rfloor - 2} m = \mathcal{O}(n)$. As shown in the remainder of Section 3.1, answering a call `PeriodicMatches`$(P_{j,i}, T_j, k_{j,i}, d_{j,i}, Q_{j,i})$ reduces, in $\mathcal{O}(d_{j,i}^2)$ time in the `PILLAR` model, to solving an instance NEWPERIODICMATCHES$(P_{j,i}, T_j, k_{j,i}, d_{j,i}, Q_{j,i}, \mathcal{A}_{P_{j,i}}, \mathcal{A}_{T_j})$. Observe that then, for a fixed $j \in [\,0 \mathinner{.\,.} \lfloor 2n/m \rfloor - 1\,)$, we have

$$\sum_i \mathcal{O}(d_{j,i}^2) = \sum_i \mathcal{O}(k^2/m^2 \cdot |P_{j,i}|^2) = \sum_i \mathcal{O}(k^2/m \cdot |P_{j,i}|) = \mathcal{O}(k^2).$$

Hence, over all $j \in [\,0 \mathinner{.\,.} \lfloor 2n/m \rfloor - 1\,)$, the total time required for the reduction in the `PILLAR` model is $\mathcal{O}(n/m \cdot k^3)$, dominated by the time required for the calls to the algorithm underlying Fact 2.15. Finally, observe that for each T_j, we obtain a representation of $\mathrm{Occ}_k^E(P, T_j)$ as $\mathcal{O}(k^3)$ arithmetic progressions with the same difference, which is computable from P and k. As P and k are common in all instances PMWITHEDITS(P, T_j, k), all computed arithmetic progressions have the same difference. In linear time in

their number, we can linearly scan them in order to merge overlapping ones, thus ensuring that the final output consists of $\mathcal{O}(n/m \cdot k^3)$ disjoint arithmetic progressions with the same difference.

In the remainder of Section 3.1 we show that answering a call `PeriodicMatches`(P, T, k, d, Q) reduces, in $\mathcal{O}(d^2)$ time in the `PILLAR` model, to solving an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$.

We first use the following (simplified) version of [CKW20, Lemma 6.8]; it is readily verified that all conditions of Fact 3.1 are satisfied in this call.

Fact 3.1 (`FindRelevantFragment`$(P,\ T,\ k,\ d,\ Q)$, Compare [CKW20, Lemma 6.8]). *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' of T such that $\delta_E(T', {}^*Q^*) \leq 3d$ and $|\mathrm{Occ}_k^E(P, T)| = |\mathrm{Occ}_k^E(P, T')|$. The algorithm runs in $\mathcal{O}(d^2)$ time in the `PILLAR` model.

A simplified version of [CKW20, Lemma 6.5] lets us pick Q so that $\delta_E(P, {}^*Q^*) = \delta_E(P, Q^*)$.

Fact 3.2 (`FindAWitness`$(k,\ Q,\ S)$, Compare [CKW20, Lemma 6.5]). *Let k denote a positive integer, let S denote a string, and let Q denote a primitive string that satisfies $|S| \geq (2k+1)|Q|$.*

Then, we can compute a witness $x \in \mathbb{Z}_{\geq 0}$ such that $\delta_E(S, \mathrm{rot}^{-x}(Q)^*) = \delta_E(S, {}^*Q^*) \leq k$, or report that $\delta_E(S, {}^*Q^*) > k$. The algorithm takes $\mathcal{O}(k^2)$ time in the `PILLAR` model.

We use Fact 3.2 as follows. First, we make a call `FindAWitness`(d, Q, P) to derive $x \in \mathbb{Z}_{\geq 0}$ such that $\delta_E(S, \mathrm{rot}^{-x}(Q)^*) = \delta_E(P, {}^*Q^*)$. Then, we make a call `Alignment`(P, Q, x) to the function of Lemma 2.10, which yields an optimum alignment $\mathcal{A} : P \rightsquigarrow Q^\infty [\,x \mathinner{\ldotp\ldotp} y\,)$ of cost $\delta_E(P, {}^*Q^*)$. Next, we extract a fragment Q' of P that $\mathcal{A}$ matches without any edits against $Q^\infty [\,x' \mathinner{\ldotp\ldotp} x' + |Q|\,)$ for some $x' \equiv x \bmod |Q|$. Finally, we replace $Q := Q'$.

We conclude that it suffices to implement the `PeriodicMatches` procedure under the following assumptions:

- $m - k \leq n < \frac{3}{2}\,m + k$,
- $\delta_E(T, {}^*Q^*) \leq 3d$,
- $\delta_E(P, {}^*Q^*) = \delta_E(P, Q^*)$.

These are exactly the conditions of the NEWPERIODICMATCHES problem.

Next, we show how to efficiently obtain optimal alignments from P and T to Q^∞, thus completing the proof that of the fact that answering a call `PeriodicMatches`(P, T, k, d, Q) reduces, in $\mathcal{O}(d^2)$ time in the `PILLAR` model, to solving an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$.

Lemma 3.3. *For any instance of the NEWPERIODICMATCHES problem, we can construct optimal alignments $\mathcal{A}_P : P \rightsquigarrow Q^\infty [\,0 \mathinner{\ldotp\ldotp} y_P\,)$ of cost d_P (for some $y_P \in \mathbb{Z}_{\geq 0}$) and $\mathcal{A}_T : T \rightsquigarrow Q^\infty [\,x_T \mathinner{\ldotp\ldotp} y_T\,)$ of cost d_T (for some $x_T \in [\,0 \mathinner{\ldotp\ldotp} |Q|\,)$ and $y_T \in \mathbb{Z}_{\geq 0}$) in $\mathcal{O}(d^2)$ time in the `PILLAR` model.*

Proof. As for $\mathcal{A}_P$, we call `Alignment`$(P, Q, 0)$ of Lemma 2.10, which directly yields the alignment $\mathcal{A}_P : P \rightsquigarrow Q^\infty [\,0 \mathinner{\ldotp\ldotp} y_P\,)$ of cost $d_P = \delta_E(P, Q^*)$.

As for $\mathcal{A}_T$, we first call `FindAWitnes`$(3d, Q, T)$ of Fact 3.2 to obtain a positive integer x_T such that $\delta_E(T, \mathrm{rot}^{-x_T}(Q)^*) = d_T$. This call is valid as the properties of the NEWPERIODICMATCHES instance yield

$$n \geq m - k \geq 8d|Q| - k \geq 7d|Q| \geq (6d+1)|Q|.$$

Moreover, due to $\mathrm{rot}^{-x_T}(Q) = \mathrm{rot}^{-x_T \bmod |Q|}(Q)$, we may set $x_T := x_T \bmod |Q|$ and thus assume $x_T \in [\,0 \mathinner{\ldotp\ldotp} |Q|\,)$ without loss of generality. Next, we call `Alignment`(T, Q, x_T) to obtain an alignment $T \rightsquigarrow Q^\infty [\,x_T \mathinner{\ldotp\ldotp} y_T\,)$ of cost $\delta_E(T, \mathrm{rot}^{-x_T}(Q)^*) = d_T$.

The overall running time in the `PILLAR` model is $\mathcal{O}(1 + d_P^2 + d^2 + d_T^2) = \mathcal{O}(d^2)$.

3.2 A First Algorithm for NewPeriodicMatches

For the remainder of this section, we fix an instance of the $\textsc{NewPeriodicMatches}(P, T, d, k, Q, \mathcal{A}_P, \mathcal{A}_T)$ problem and set $\kappa := k + d_P + d_T$, and $\tau := q \lceil \kappa/2q \rceil$.

In a first step toward algorithms for $\textsc{NewPeriodicMatches}$, we discuss how the (almost-)periodicity of P and T yield simple ways to *filter out* many potential starting positions of occurrences. In particular, this allows us to obtain a fast algorithm for $\textsc{NewPeriodicMatches}$ when q is very large.

Let us briefly recall an example from Section 1.3. Suppose that P and T are perfectly periodic with period Q, that is $P = Q^\infty[0 \mathbin{.\,.} m]$ and $T = Q^\infty[0 \mathbin{.\,.} n]$. In particular, in this special case $\mathcal{A}_P$ and $\mathcal{A}_T$ are identity maps and we have $m = y_P$, $0 = x_T$, $n = y_T$, and $d_T = d_P = 0$. Now, clearly all occurrences start around the positions in T where an exact occurrence of Q starts, that is, in the intervals $[jq - k \mathbin{.\,.} jq + k]$ for $j \in \mathbb{Z}$. Now, if we have $q > 2k$, then indeed $jq + k < (j+1)q - k$ and we can thus filter out positions where no occurrence may start.

Now, as we are dealing with *almost* periodic strings P and T, potential edits in $\mathcal{A}_T$ and $\mathcal{A}_P$ widen the intervals of potential starting positions. Returning to the more general T and P fixed at the beginning of this (sub-)section and given an alignment $\mathcal{A} : P \rightsquigarrow T[v \mathbin{.\,.} w)$ of cost at most k, observe that $\mathcal{A}_T \circ \mathcal{A} \circ \mathcal{A}_P^{-1}$ induces an alignment $Q^\infty[0 \mathbin{.\,.} y_P) \rightsquigarrow U := \mathcal{A}_T^{-1}(T[v \mathbin{.\,.} w))$ of cost at most $\kappa = k + d_T + d_P$. Now, as Q is primitive and $Q^\infty[0 \mathbin{.\,.} y_P)$ is a long substring of Q^∞, we can conclude that U must be a prefix of $(\mathrm{rot}^r(Q))^\infty$ for $r \in [-\kappa \mathbin{.\,.} \kappa]$. Hence, U starts at some position $u \in \bigcup_{j \in \mathbb{Z}}[jq - x_T - \kappa \mathbin{.\,.} jq - x_T + \kappa]$ of $Q^\infty[x_T \mathbin{.\,.} y_T)$. The fact that $\mathcal{A}_T$ is of cost d_T, yields that $v \in \bigcup_{j \in \mathbb{Z}}[jq - x_T - \kappa - d_T \mathbin{.\,.} jq - x_T + \kappa + d_T]$. Now, if we have $q > 2\kappa + 2d_T$, then indeed $jq - x_T + \kappa + d_T < (j+1)q - x_T - \kappa - d_T$ and we can thus filter out positions where no occurrence may start.

▮ **Lemma 3.4.** *Consider an alignment $\mathcal{A} : P \rightsquigarrow T[v \mathbin{.\,.} w)$ of cost at most k and write $Q^\infty[v' \mathbin{.\,.} w') := \mathcal{A}_T(T[v \mathbin{.\,.} w))$. Further, for an integer j, we have*

$$v \in [jq - x_T - \kappa - d_T \mathbin{.\,.} jq - x_T + \kappa + d_T],$$
$$v' \in [\max\{x_T, jq - \kappa\} \mathbin{.\,.} jq + \kappa], \quad and$$
$$w' \in [jq + y_P - \kappa \mathbin{.\,.} \min\{y_T, jq + y_P + \kappa\}].$$

Further, we have $|Q^\infty[v' \mathbin{.\,.} w')| \in [y_P - \kappa \mathbin{.\,.} y_P + \kappa]$.

Proof. We formalize the example from before. To that end, write $\mathcal{B} := \mathcal{A}_T \circ \mathcal{A} \circ \mathcal{A}_P^{-1}$ for the alignment of $Q^\infty[0 \mathbin{.\,.} y_P)$ to $Q^\infty[v' \mathbin{.\,.} w')$ induced by $\mathcal{A}$; observe that $\mathcal{B}$ has a cost of at most κ. Observe that this immediately yields the claimed bound on the size of $Q^\infty[v' \mathbin{.\,.} w')$.

Now, observe that the fragments of Q^∞ that $\mathcal{B}$ aligns to each other are both long; hence at least one full occurrence of Q is matched exactly under $\mathcal{B}$.

▱ **Claim 3.5.** There is a fragment $Q_i := Q^\infty[iq \mathbin{.\,.} (i+1)q)$ of $Q^\infty[0 \mathbin{.\,.} y_P)$ with $\mathcal{B}(Q_i) = Q_i = Q$.

Proof. As P, T, d, k, and Q stem from our fixed instance of $\textsc{NewPeriodicMatches}$, we have

$$8dq \le m \quad and \quad d_P \le d \quad and \quad \kappa \le 5d \quad and \quad 0 < d.$$

Combined, we hence also have

$$|Q^\infty[0 \mathbin{.\,.} y_P)| = y_P \ge m - d_P \ge 8dq - d \ge 7dq \ge (\kappa + 2)q.$$

In particular, the fragment $Q^\infty[0 \mathbin{.\,.} y_P)$ contains at least $(\kappa + 1)$ full repetitions of Q. As the alignment $\mathcal{B}$ makes at most κ edits, at least one full occurrence of Q is hence aligned without edits. ⌟

Now, observe that the fragment $\mathcal{B}(Q_i)$ from Claim 3.5 starts at some position $i'q$ in Q^∞—as $\mathcal{B}$ allows for at most κ insertions and deletions, we hence obtain that $v' \in [\,(i'-i)q - \kappa \mathinner{..} (i'-i)q + \kappa\,]$. Symmetrically, we obtain $w' \in [\,(i'-i)q + y_P - \kappa \mathinner{..} (i'-i)q + y_p + \kappa\,]$. Further, we observe that $\mathcal{B}$ cannot extend beyond $\mathcal{A}_T$. Finally, as $\mathcal{A}_T(T[\,v \mathinner{..} w\,)) = Q^\infty[\,v' \mathinner{..} w'\,)$ and $\mathcal{A}_T$ is of cost d_T, we also have that $v \in [\,(i'-i)q - \kappa - d_T \mathinner{..} (i'-i)q + \kappa + d_T\,]$, thus completing the proof. ◾

Observe that expanding the minimum and maximum expressions from Lemma 3.4 easily yields bounds on the possible values of j; we have

$$j \in \bar{J} := [\,\lceil(x_T - \kappa)/q\rceil \mathinner{..} \lfloor(y_T - y_P + \kappa)/q\rfloor\,] = \{j \in \mathbb{Z} : x_T \le jq + \kappa \text{ and } jq + y_P - \kappa \le y_T\}.$$

Now, intuitively, we wish to process each fragment $T[\,jq - x_T - \kappa - d_T \mathinner{..} jq - x_T + \kappa + d_T\,]$ separately (each corresponding to a single $j \in \bar{J}$). However, if q is very small (that is, if $q < \kappa/2$), this results in too many fragments to process. Hence we group them into longer fragments of length at least $\tau := q\,\lceil\kappa/2q\rceil \ge q(\kappa/2q) = \kappa/2$; observe that τ is an integer multiple of q and that we have

$$\tau = q\,\lceil\kappa/2q\rceil = q, \qquad \text{if } q \ge \kappa/2 \text{ and} \tag{1}$$
$$\tau \le q(1 + \kappa/2q) < \kappa, \quad \text{if } q < \kappa/2. \tag{2}$$

As discussed in the introductory example, for two consecutive integers j and $j + 1$, the intervals $[\,jq - \kappa \mathinner{..} jq + \kappa\,]$ and $[\,(j+1)q - \kappa \mathinner{..} (j+1)q + \kappa\,]$ are disjoint only if $q > 2\kappa$—in which case $\tau = q$. Hence, the definition of τ agrees with Lemma 3.4.

In total, this leads to the following definition.

�макед **Definition 3.6.** *Partition Q^∞ into blocks of length τ and number them starting from 0. For the j-th block $Q_j := Q^\infty[\,j\tau \mathinner{..} (j+1)\tau\,)$, we define the* interesting region *$\bar{Q}_j$ of Q_j as the fragment $\bar{Q}_j := Q^\infty[\,j\tau - \kappa \mathinner{..} j\tau + \kappa\,]$. Now, write J for the set of all numbers of blocks whose interesting regions overlap $Q^\infty[\,x_T \mathinner{..} y_T - y_P + \kappa\,)$. Formally, we set*

$$J := [\,\lceil(x_T - \kappa)/\tau\rceil \mathinner{..} \lfloor(y_T - y_P + \kappa)/\tau\rfloor\,] = \{j \in \mathbb{Z} : x_T \le j\tau + \kappa \text{ and } j\tau + y_P - \kappa \le y_T\}.$$

Now, for each $j \in J$, we write $R_j := T[\,r_j \mathinner{..} r_j'\,)$ for the largest fragment of T that corresponds to the interesting region $\bar{Q}_j$; that is, we set

$$r_j := \min\{a_T : (a_T, a_Q) \in \mathcal{A}_T \text{ for } j\tau - \kappa \le a_Q\} \quad \text{and}$$
$$r_j' := \max\{a_T : (a_T, a_Q) \in \mathcal{A}_T \text{ for } a_Q \le j\tau + y_P + \kappa\}.$$

For convenience, we also set $r_{\max J + 1} := n$. ◾

▭ **Remark 3.7.** For convenience, we write

$$\mathcal{A}_T(r_j) := \min\{a_Q : (r_j, a_Q) \in \mathcal{A}_T \text{ and } j\tau - \kappa \le a_Q\} \quad \text{and}$$
$$\mathcal{A}_T(r_j') := \max\{a_Q : (r_j', a_Q) \in \mathcal{A}_T \text{ and } a_Q \le j\tau + y_P + \kappa\}.$$

Observe that for *most* of $j \in J$, we have $\mathcal{A}_T(r_j) = j\tau - \kappa$ and $\mathcal{A}_T(r_j') = j\tau + y_P + \kappa$. In particular, whenever $\mathcal{A}_T(r_j) > j\tau - \kappa$, that is, if $j \in [\,\lceil(x_T - \kappa)/\tau\rceil \mathinner{..} \lfloor(x_T + \kappa)/\tau\rfloor\,]$, then $j\tau - \kappa$ is a position before the first position that $\mathcal{A}_T$ maps to: we have $j\tau - \kappa \le x_T$ and hence $\mathcal{A}_T(r_j) = x_T$, which in turn implies $r_j = 0$.

Similarly, whenever $\mathcal{A}_T(r_j') < j\tau + y_P + \kappa$, that is, if $j \in [\,\lceil(y_T - y_P - \kappa)/\tau\rceil \mathinner{..} \lfloor(y_T - y_P + \kappa)/\tau\rfloor\,]$, we have $y_T < j\tau + y_P + \kappa$ and hence $\mathcal{A}_T(r_j') = y_T$, which in turn implies $r_j' = n$. ◾

◫ **Remark 3.8.** It is useful to number the blocks before the zeroth block with negative indices: observe that for $x_T < \kappa$, we may have $(\min J) < 0$. In particular, we see that $(\min J) \geq \lceil -\kappa/\tau \rceil \geq -2$, so we may indeed need to address a -1-st and -2-nd block.

Similarly, $(\max J)$ may exceed the natural barrier of $(y_T - y_P)/\tau$. In particular, we see that $(\max J) \leq \lceil y_T/\tau \rceil - \lceil y_P/\tau \rceil + \kappa/\tau \leq \lceil y_T/\tau \rceil - \lceil y_P/\tau \rceil + 2$. ◪

◫ **Remark 3.9.** It is useful to upper bound the length of each fragment R_j: fix a $j \in J$ and consider the fragment $\mathcal{A}_T(R_j)$ of Q^∞. By construction, $\mathcal{A}_T(R_j)$ has a length of

$$|\mathcal{A}_T(R_j)| = (y_P + j\tau + \kappa) - (j\tau - \kappa) = y_P + 2\kappa.$$

Now, as $\mathcal{A}_T$ has a cost of d_T, we conclude that the fragment R_j has a length of at most

$$|R_j| = r_j' - r_j \leq |\mathcal{A}_T(R_j)| + d_T = y_P + 2\kappa + d_T \leq m + 3\kappa - k.$$ ◪

Intuitively, we want to think of R_j as the fragment of T that contains all k-error occurrences of P that start at a position in T that corresponds to a position in $\bar{Q}_j$. Formally, this requires a proof as a fragment R_j could (in theory) be too short to fully contain such an occurrence.

◫ **Lemma 3.10.** *Fix a $j \in J$ and a position $v \in [\, r_j \mathinner{\ldotp\ldotp} r_{j+1})$. For any position $w > r_j'$, any alignment $\mathcal{A} : P \rightsquigarrow T[\, v \mathinner{\ldotp\ldotp} w \,)$ has a cost of at least $k + 1$.*

Proof. We prove the contraposition. To that end, fix an alignment $\mathcal{A} : P \rightsquigarrow T[\, v \mathinner{\ldotp\ldotp} w \,)$ with a cost of at most k. We intend to show that $w \leq r_j'$.

To that end, write $Q^\infty[\, \mathcal{A}_T(v) \mathinner{\ldotp\ldotp} \mathcal{A}_T(w) \,) := \mathcal{A}_T(T[\, v \mathinner{\ldotp\ldotp} w \,))$ and consider the position $\mathcal{A}_T(v)$. First, suppose that $q \geq \kappa/2$ and, in particular, $\tau = q$. Now, the definition of r_j yields[13]

$$\mathcal{A}_T(v) \in [\, j\tau - \kappa \mathinner{\ldotp\ldotp} \min\{y_T, (j+1)\tau - \kappa\} \,) \subseteq [\, jq - \kappa \mathinner{\ldotp\ldotp} (j+1)q - \kappa \,)$$

Applying Lemma 3.4, we see that $\mathcal{A}_T(v) \in [\, jq - \kappa \mathinner{\ldotp\ldotp} jq + \kappa \,]$ and

$$\mathcal{A}_T(w) \leq \min\{y_T, jq + y_P + \kappa\} \leq jq + y_P + \kappa = j\tau + y_P + \kappa$$

Hence, by definition $w \leq r_j'$.

Next, suppose that $q < \kappa/2$. Now, using Equation (1), we see that[14]

$$\mathcal{A}_T(v) \leq (j+1)\tau - \kappa \leq j\tau + \kappa - \kappa.$$

Now, using Lemma 3.4, we see that $|\mathcal{A}_T(T[\, v \mathinner{\ldotp\ldotp} w \,))| \leq y_P + \kappa$ and hence

$$\mathcal{A}_T(w) \leq \mathcal{A}_T(v) + |\mathcal{A}_T(T[\, v \mathinner{\ldotp\ldotp} w \,))| \leq j\tau + y_P + \kappa.$$

Thus, we have $w \leq r_j'$; completing the proof. ◪

In particular, Lemma 3.10 confirms our earlier intuition: indeed, the fragment R_j fully contains all k-error occurrences of P that start at a position in T that corresponds to a position in $\bar{Q}_j$.

As another consequence of Lemma 3.10, we see that the set $\mathrm{Occ}_k^E(P, T)$ decomposes over J and R_j.

---

[13] Observe that if j is equal to $\max J$, then by Lemma 3.4, for $\mathcal{A}_T(v) \in [\, (j+1)\tau - \kappa \mathinner{\ldotp\ldotp} y_T \,]$ the interval of possible positions for w' is empty. Hence, also in this corner case, we can safely assume $\mathcal{A}_T(v) \leq (j+1)\tau$.

[14] Again, if j is $\max J$, then by Lemma 3.4 with $\tau \geq q$, we may assume $\mathcal{A}_T(v) < (j+1)\tau$.

▪ **Corollary 3.11**. *For any position $v \in \mathrm{Occ}_k^E(P, T)$ and a corresponding alignment $\mathcal{A} : P \rightsquigarrow T[v \mathbin{..} w)$, there is a (unique) $j \in J$ such that $v \in [r_j \mathbin{..} r_{j+1})$ and $w \le r_j'$; that is, we have $\mathrm{Occ}_k^E(P, T) = \bigcup_{j \in J} \left(\mathrm{Occ}_k^E(P, R_j) + r_j\right)$.*

Proof. Recall that $R_j = T[r_j \mathbin{..} r_j')$. Now, using Lemma 3.10, we see that for each $j \in J$, the set $\mathrm{Occ}_k^E(P, R_j)$ (shifted by r_j) contains all k-error occurrences of P that start in $T[r_j \mathbin{..} r_{j+1})$.

Finally, we need to argue that no k-error occurrence may start at a position in T not covered by the fragments R_j. To that end, first observe that $r_{\max J+1} = n$, so we may lose an occurrence only if it starts before $r_{\min J}$ in T. Set $j^* \coloneqq \min J$ and observe that we are done if $j^* = 0$, as $r_0 = 0$. Now, for a positive j^*, if $q \ge \kappa/2$ (and hence $\tau = q$), we see that for any positive integer α, we have $(j^* - \alpha)q + \kappa < x_T$—hence (by Lemma 3.4) no k-error occurrence starts in T before r_{j^*}.

Finally, if $q < \kappa/2$ (and hence $\tau < \kappa$ by Equation (2)); we have

$$j^*\tau - \kappa \le \left(1 + \frac{x_T - \kappa}{\tau}\right)\tau - \kappa < x_T - \kappa \le x_T.$$

Hence, we have $\mathcal{A}_T(r_{j^*}) = x_T$ and hence $r_{j^*} = 0$; completing the proof. ◾

Next, we discuss how to compute the sequences (r_j) and (r_j').

▪ **Lemma 3.12**. *Given the alignment $\mathcal{A}_T$, in $\mathcal{O}(d_T + 1)$ time, we can construct the sequences $(r_j)_{j \in J}$ and $(r_j')_{j \in J}$, represented as concatenations of $\mathcal{O}(d_T + 1)$ arithmetic progressions with difference τ.*

Proof. We intend to employ Lemma 2.4 to compute the first positions in T that $\mathcal{A}_T$ maps to a position $j\tau - \kappa$ and the last positions in T that $\mathcal{A}_T$ maps to a position $j\tau + y_P + \kappa$. Thereby, we immediately obtain the desired entries of r_j when $\mathcal{A}_T(r_j) = j\tau - \kappa$ (by setting $r_j \coloneqq \sigma_{j\tau-\kappa}^{\mathcal{A}_T}$) and the desired entries of r_j' when $\mathcal{A}_T(r_j') = j\tau + y_P - \kappa$ (by setting $r_j' \coloneqq \lambda_{j\tau+y_P-\kappa}^{\mathcal{A}_T}$). Recalling Remark 3.7, we see that the remaining entries of (r_j) are those that correspond to a (short) prefix of T (for which $r_j = 0$) and we see that the remaining entries of (r_j') are those that correspond to a (short) suffix of T (for which $r_j' = n$).

Formally, we use Lemma 2.4 to compute the sequences $(\sigma_v^{\mathcal{A}_T})_{v=x_T}^{y_T}$ and $(\lambda_w^{\mathcal{A}_T})_{w=x_T}^{y_T}$—which we obtain represented as $\mathcal{O}(d_T + 1)$ arithmetic progressions with difference 1 (that is intervals) each. Now, for the sequence $(r_j)_{j \in J}$; for $j \in [\lceil (x_T - \kappa)/\tau \rceil \mathbin{..} \lfloor (x_T + \kappa)/\tau \rfloor]$, we set $r_j = 0$ (which we output as single-element arithmetic progressions).

For the remaining values of (r_j), that is for $j \in [\lceil (x_T + \kappa)/\tau \rceil \mathbin{..} \lfloor (y_T - y_P + \kappa)/\tau \rfloor]$, we process the intervals that represent $(\sigma_v^{\mathcal{A}_T})$ as follows: for each interval I representing $(\sigma_v^{\mathcal{A}_T})_{v=x_T}^{x_T+|I|-1}$, we compute the intersection $[x_T \mathbin{..} x_T + |I|) \cap (j\tau - \kappa)_j$ and return the corresponding entries of I as an arithmetic progression with difference τ. For $(r_j')_{j \in J}$ we proceed analogously.

For the running time, computing $(\sigma_v^{\mathcal{A}_T})$ and $(\sigma_w^{\mathcal{A}_T})$ takes time $\mathcal{O}(d_T + 1)$; post-processing the obtained intervals takes constant time per interval for a total running time of $\mathcal{O}(d_T + 1)$. Next, observe that as $\tau \ge \kappa/2$, we have

$$|[\lceil (x_T - \kappa)/\tau \rceil \mathbin{..} \lfloor (x_T + \kappa)/\tau \rfloor]| \le 2\kappa/\tau \le 4.$$

In particular we can handle the initial values of (r_j) in constant time and we output a constant number of extra arithmetic progressions. Similarly, we have

$$|[\lceil (y_T - y_P - \kappa)/\tau \rceil \mathbin{..} \lfloor (y_T - y_P + \kappa)/\tau \rfloor]| \le 2\kappa/\tau \le 4.$$

In particular we can handle the final values of (r'_j) in constant time and we output a constant number of extra arithmetic progressions. Hence in total, the algorithm runs in time $\mathcal{O}(d_T + 1)$ and produces $\mathcal{O}(d_T + 1)$ arithmetic progressions, each with difference τ;[15] thus completing the proof. ◾

Finally, we present a first algorithm for NewPeriodicMatches, which is particularly useful for small sets J, that is, when $q \gg k$. Let us note that this algorithm is also used in the general case for the efficient computation of k-error occurrences that start in heavy positions; see Section 6.1.

◫ **Lemma 3.13**. *For each $j \in J$, we can compute the set $\mathrm{Occ}_k^E(P, R_j)$ in $\mathcal{O}(d^2)$ time in the* PILLAR *model. In particular, we can solve the NewPeriodicMatches problem in $\mathcal{O}(d^2|J|)$ time in the* PILLAR *model.*

Proof. Recall from Remark 3.9 that for each $j \in J$ we have $|R_j| \leq m + 3\kappa - k$. Thus, by Fact 2.9 and Lemma 3.12, the construction of $\mathrm{Occ}_k^E(P, R_j)$ using $\mathtt{Verify}(P, R_j, k, [\,0\mathinner{.\,.}|R_j| - m + k\,))$ takes $\mathcal{O}(\log(1 + d_T) + k(k + \kappa)) = \mathcal{O}(d^2)$ time in the PILLAR model. Corollary 3.11 shows how the sets $\mathrm{Occ}_k^E(P, R_j)$ can be combined into $\mathrm{Occ}_k^E(P, T)$; this takes time $\mathcal{O}(\kappa|J|) = \mathcal{O}(d|J|)$. In total, we see that the overall running time in the PILLAR model is dominated by the calls to $\mathtt{Verify}$, which take time $\mathcal{O}(d^2|J|)$ in the PILLAR model. ◾

4 Using DynamicPuzzleMatching for Algorithms for NewPeriodicMatches

In this section, we formalize the connection between the DynamicPuzzleMatching problem and the NewPeriodicMatches problem. In particular, we define suitable puzzles for the strings P and T; then we show that in the setting of NewPeriodicMatches, we can compute $\mathrm{Occ}_k^E(P, T)$ with using very few DPM-∗ operations. To that end, we first define partitions of P and T (resulting in *tiles*) from which we then construct appropriate puzzles.

◫ **Definition 4.1** (τ-tile partition). *Consider a string S, a primitive string Q of length q, an integer $\tau \in \mathbb{Z}_{>0}$ divisible by q, and an alignment $\mathcal{A}_S : S \rightsquigarrow Q^\infty[\,x_S \mathinner{.\,.} y_S\,)$, where $x_S \in [\,0\mathinner{.\,.}q\,)$ and $y_S \geq x_S$.*

Partition Q^∞ into blocks of length τ and number them starting from 1. For the j-th block $Q_j \coloneqq Q^\infty[\,\max\{x_S, (j-1)\tau\} \mathinner{.\,.} \min\{y_S, j\tau\}\,)$, we define the j-th tile of S (with respect to $\mathcal{A}_S$) as

$$S[\,s_{i-1} \mathinner{.\,.} s_i\,) \coloneqq \mathcal{A}_S^{-1}(Q_j).$$

Further, we define the τ-tile partition of S with respect to $\mathcal{A}_S$ as the partition

$$S = \bigodot_{i \in \mathbb{Z}_{>0}} S[\,s_{i-1} \mathinner{.\,.} s_i\,).$$ ◾

◫ **Remark 4.2**. Observe that the $\lceil y_S/\tau \rceil$-th tile is the last non-empty tile of S. Hence, writing $\beta_S \coloneqq \lceil y_S/\tau \rceil$, we have

$$S = \bigodot_{i=1}^{\beta_S} S[\,s_{i-1} \mathinner{.\,.} s_i\,).$$ ◾

◫ **Remark 4.3**. We observe that $s_0 = 0$ and $s_{\beta_S} = |S|$. We also see that that for all other $j \in (\,0\mathinner{.\,.}\beta_S\,)$, we have

$$s_i = \min\{a_S \in [\,0\mathinner{.\,.}|S|\,] : (a_S, i\tau) \in \mathcal{A}_S\}.$$

Recalling Lemma 2.4, we hence have $s_i = \sigma_{i\tau}^{\mathcal{A}_S}$. ◾

---

[15] Observe that this is also true for single-element arithmetic progressions.

It is easy to verify that we can indeed efficiently compute the τ-tile partition of a string.

Lemma 4.4. *Let $S = \bigodot_{i=1}^{\beta} S[\, s_{i-1} \mathinner{..} s_i \,)$ denote the τ-tile partition of a string S with respect to a cost-d alignment $\mathcal{A}_S : S \rightsquigarrow Q^\infty[\, x_S \mathinner{..} y_S \,)$.*

Given (the breakpoint representation of) $\mathcal{A}_S$, in $\mathcal{O}(d+1)$ time, we can construct the sequence $(s_i)_{i=0}^{\beta}$, represented as $\mathcal{O}(d+1)$ arithmetic progressions with difference τ.

Proof. We use Lemma 2.4 to obtain the sequence $(\sigma_v^{\mathcal{A}_S})_{v=x_S}^{y_S}$ as a concatenation of $\mathcal{O}(d+1)$ arithmetic progressions with difference 1. Next, in view of Remark 4.3, we extract and return the subsequence $(\sigma_{j\tau}^{\mathcal{A}_S})_{j=1}^{\beta_S-1} = s_j$; together with the boundary values of $s_0 = 0$ and $s_{\beta_S} = |S|$.

For the running time, observe that it suffices to process each arithmetic progression in $(\sigma_v^{\mathcal{A}_S})_{v=x_S}^{y_S}$ in $\mathcal{O}(1)$ time. Further, each such arithmetic progression contributes a single (perhaps empty) arithmetic progression of difference τ to the output; to which we add exactly two further arithmetic progressions. In total, this yields the claim. ⌟

As in Section 3.2, we fix an instance of the NEWPERIODICMATCHES($P, T, d, k, Q, \mathcal{A}_P, \mathcal{A}_T$) problem and we set $\kappa := k + d_P + d_T$, and $\tau := q\lceil \kappa/2q \rceil$.

Now additionally, let $P = \bigodot_{i=1}^{\beta_P} P[\, p_{i-1} \mathinner{..} p_i \,)$ denote the τ-tile partition of P with respect to $\mathcal{A}_P$, and let $T = \bigodot_{i=1}^{\beta_T} T[\, t_{i-1} \mathinner{..} t_i \,)$ denote the τ-tile partition of T with respect to $\mathcal{A}_T$.

Toward our goal of creating a suitable puzzle for DYNAMICPUZZLEMATCHING, we first observe that if β_P is very small (that is if the tiles are very long), then we can already efficiently solve the NEWPERIODICMATCHES problem.

Lemma 4.5. *We can solve NEWPERIODICMATCHES in time $\mathcal{O}(d^2\beta_P)$ in the* `PILLAR` *model.*

Proof. In light of Lemma 3.13, it suffices to prove that $|J| = \mathcal{O}(\beta_P)$.

To that end, observe that by the definition of J (Definition 3.6), we have

$$|J| - 1 = \left\lfloor \frac{y_T - y_P + \kappa}{\tau} \right\rfloor - \left\lceil \frac{x_T - \kappa}{\tau} \right\rceil \le \frac{y_T - y_P - x_T + 2\kappa}{\tau} = \frac{(y_T - x_T) - (y_P - 0) + 2\kappa}{\tau}$$

Rephrasing in terms of n and m yields

$$|J| - 1 \le \frac{n + d_T - m + d_P + 2\kappa}{\tau}.$$

Using $n - m \le m/2 + k$ (from using the Standard Trick) and, thereafter, $\kappa := k + d_P + d_T$ yield

$$|J| - 1 \le \frac{m/2 + k + d_P + d_T + 2\kappa}{\tau} = \frac{m + 6\kappa}{2\tau}.$$

Rephrasing in terms of y_P and, thereafter, using $\tau \ge \kappa/2$, and the definition of β_P yield

$$|J| - 1 \le \frac{y_P + 7\kappa}{2\tau} \le \frac{y_P}{2\tau} + 7 \le \beta_P/2 + 7.$$

Rearranging yields $|J| \le \beta_P/2 + 8 \le {}^{17}\!/_2\, \beta_P$, completing the proof. ⌟

In particular, Lemma 4.5 allows us to assume (without loss of generality) that $\beta_P \ge 20$. Further, we set $\Delta := 6\kappa$ and $z := \beta_P - 17$.

Now, we (essentially) extend the tiles from the τ-tile partition of P by an additional Δ characters to obtain a Δ-puzzle with value P.

Lemma 4.6. *The following sequence $P_1, \ldots, P_z$ forms a Δ-puzzle with value P:*

- $P_1 := P[\, p_0 \mathinner{\ldotp\ldotp} p_2 + \Delta\,)$;
- $P_i := P[\, p_i \mathinner{\ldotp\ldotp} p_{i+1} + \Delta\,)$ *for* $i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$;
- $P_z := P[\, p_z \mathinner{\ldotp\ldotp} |P|\,)$.

Proof. The only non-trivial claim is that each internal piece is well-defined, that is that $0 \le p_i$ and that $p_{i+1} + \Delta \le |P|$. To verify this claim, observe that we have for every $i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$

$$0 = p_0 \le p_i \quad \text{and}$$

$$p_{i+1} + \Delta \le p_z + \Delta = p_z + 6\kappa$$

Using $\tau \ge \kappa/2$ and, thereafter, $\bigl|(p_{z+14} - p_z) - 14\tau\bigr| \le \delta_E(P[\, p_z \mathinner{\ldotp\ldotp} p_{z+14}\,), Q^\infty[\, z\tau \mathinner{\ldotp\ldotp} (z+14)\tau\,]) \le d_P$, we obtain

$$p_{i+1} + \Delta \le p_z + 14\tau - d_P \le p_{z+14} \le |P|. \hfill \blacksquare$$

Similarly, we obtain Δ-puzzles for T. In particular, we use the τ-tile partition of T in order to represent the fragments R_j as Δ-puzzles.

⌗ Lemma 4.7. *For each $j \in J$, the following sequence $T_{j,1}, \ldots, T_{j,z}$ forms a Δ-puzzle with value R_j:*
- $T_{j,1} := T[\, r_j \mathinner{\ldotp\ldotp} t_{j+2} + \Delta\,)$;
- $T_{j,i} := T[\, t_{j+i} \mathinner{\ldotp\ldotp} t_{j+i+1} + \Delta\,)$ *for* $i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$;
- $T_{j,z} := T[\, t_{j+z} \mathinner{\ldotp\ldotp} r'_j\,)$.

Proof. The only non-trivial claim is that each internal piece is contained within R_j, that is that $r_j \le t_{j+i}$ and that $t_{j+i+1} + \Delta \le r'_j$. To verify this claim, first observe that $\tau \ge \kappa/2$ yields for every $j \in J = [\, \lceil (x_T - \kappa)/\tau \rceil \mathinner{\ldotp\ldotp} \lfloor (y_T - y_P + \kappa)/\tau \rfloor\,]$ and every $i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$ that

$$x_T \le x_T - \kappa + 2\tau \le \left(\left\lceil \frac{x_T - \kappa}{\tau} \right\rceil + 2 \right) \tau \le (j+i)\tau.$$

In particular, we have $(t_{j+i}, (j+i)\tau) \in \mathcal{A}_T$. Similarly, we obtain $j\tau - \kappa \le (j+2)\tau \le (j+i)\tau$, so $r_j \le t_{j+i}$ holds by definition of r_j.

For the other bound, we first observe that for any $j \in J$ and $i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$, we have

$$(j+i+15)\tau \le \left(\left\lfloor \frac{y_T - y_P + \kappa}{\tau} \right\rfloor + \beta_P - 3 \right) \tau \le y_T - y_P + \kappa + (\beta_P - 3)\tau \le y_T - y_P + \kappa + y_P - 2\tau \le y_T.$$

In particular, we have $(t_{j+i+15}, (j+i+15)\tau) \in \mathcal{A}_T$. Similarly, we obtain $(j+i+15)\tau \le (j+\beta_P - 3)\tau \le y_P + j\tau - \kappa$, so $t_{j+i+15} \le r'_j$ holds by definition of r'_j. Now, as $\Delta = 7\kappa - \kappa \le 14\tau - d_T \le t_{j+i+15} - t_{j+i+1}$, we obtain $t_{j+i+1} + \Delta \le r'_j$; completing the proof. $\hfill \blacksquare$

Taken together, we obtain the families of puzzle pieces that we use in the remainder of this work.

⌗ Definition 4.8. — *We write $\mathcal{S}_\beta := \{P_1\} \cup \{T_{j,1} : j \in J\}$ for the family of* leading *puzzle pieces.*
- *We write $\mathcal{S}_\mu := \{P_i : i \in (\, 1 \mathinner{\ldotp\ldotp} z\,)\} \cup \{T_i : i \in (\, \min J + 1 \mathinner{\ldotp\ldotp} \max J + z\,)\}$ for the family of* internal *puzzle pieces.*
- *We write $\mathcal{S}_\varphi := \{P_z\} \cup \{T_{j,z} : j \in J\}$ for the family of* trailing *puzzle pieces.* $\hfill \blacksquare$

⌗ Remark 4.9. For convenience, for $i \in (\, \min J + 1 \mathinner{\ldotp\ldotp} \max J + z\,)$, we write $T_i := T[\, t_i \mathinner{\ldotp\ldotp} t_{i+1} + \Delta\,)$. Observe that by construction, we have $T_i = T_{j,i'}$ for all $j + i' = i$ with $i' \in (\, 1 \mathinner{\ldotp\ldotp} z\,)$; that is overlapping parts of different R_j's share their internal pieces. Observe further that this is an essential property for our approach to work: when moving from R_j to R_{j+1}, we exploit that we need to only shift the pieces T_i, and not recompute them altogether. $\hfill \blacksquare$

Finally, we wish to use the sets from Definition 4.8 for DYNAMICPUZZLEMATCHING. To that end, we need to convince ourselves that pairs of pieces P_i and $T_{j,i}$ are indeed roughly of the same length on average. In particular, we verify that our choice of $\Delta := 6\kappa$ is indeed sufficient.

▰ **Lemma 4.10.** *For each $j \in J$, we have $\sum_{i=1}^{z} \big||T_{j,i}| - |P_i|\big| \leq 3\kappa - k$.*

Proof. Fix a $j \in J$. In a first step, we obtain bounds on the length differences of each individual pair of pieces. To that end, consider a pair of internal pieces P_i and $T_{j,i}$ for an $i \in (\,1 \mathinner{.\,.} z\,)$; the triangle inequality yields

$$
\begin{aligned}
\big||T_{j,i}| - |P_i|\big| &\leq \delta_E(T[\,t_{j+i} \mathinner{.\,.} t_{j+i+1}\,), P[\,p_i \mathinner{.\,.} p_{i+1}\,)) \\
&\leq \delta_E(T[\,t_{j+i} \mathinner{.\,.} t_{j+i+1}\,), Q^\infty[\,(j+i)\tau \mathinner{.\,.} (j+i+1)\tau\,)) \\
&\quad + \delta_E(P[\,p_i \mathinner{.\,.} p_{i+1}\,), Q^\infty[\,i\tau \mathinner{.\,.} (i+1)\tau\,)).
\end{aligned}
\tag{3}
$$

As for $i = 1$, recall from Remark 3.7 that $\mathcal{A}_T(r_j) = \max\{x_T, j\tau - \kappa\} \leq j\tau + \kappa$. Hence, the triangle inequality yields

$$
\begin{aligned}
\big||T_{j,1}| - |P_1|\big| &\leq \delta_E(T[\,r_j \mathinner{.\,.} t_{j+2}\,), P[\,0 \mathinner{.\,.} p_2\,)) \\
&\leq \delta_E(T[\,r_j \mathinner{.\,.} t_{j+2}\,), Q^\infty[\,\mathcal{A}_T(r_j) \mathinner{.\,.} (j+2)\tau\,)) \\
&\quad + \delta_E(Q^\infty[\,\mathcal{A}_T(r_j) \mathinner{.\,.} (j+2)\tau\,), Q^\infty[\,0 \mathinner{.\,.} 2\tau\,)) \\
&\quad + \delta_E(P[\,0 \mathinner{.\,.} p_2\,), Q^\infty[\,0 \mathinner{.\,.} 2\tau\,)) \\
&\leq \delta_E(T[\,0 \mathinner{.\,.} t_{j+2}\,), Q^\infty[\,x_T \mathinner{.\,.} (j+2)\tau\,)) + \kappa + \delta_E(P[\,0 \mathinner{.\,.} p_2\,), Q^\infty[\,0 \mathinner{.\,.} 2\tau\,)).
\end{aligned}
\tag{4}
$$

As for $i = z$, recall from Remark 3.7 that $\mathcal{A}_T(r'_j) = \min\{y_T, y_P + j\tau + \kappa\} \geq y_P + j\tau - \kappa$. Hence, the triangle inequality yields

$$
\begin{aligned}
\big||T_{j,z}| - |P_z|\big| &\leq \delta_E(T[\,t_{j+z} \mathinner{.\,.} r'_j\,), P[\,p_z \mathinner{.\,.} |P|\,)) \\
&\leq \delta_E(T[\,t_{j+z} \mathinner{.\,.} r'_j\,), Q^\infty[\,(j+z)\tau \mathinner{.\,.} \mathcal{A}_T(r'_j)\,)) \\
&\quad + \delta_E(Q^\infty[\,(j+z)\tau \mathinner{.\,.} \mathcal{A}_T(r'_j)\,), Q^\infty[\,z\tau \mathinner{.\,.} y_P\,)) \\
&\quad + \delta_E(P[\,p_z \mathinner{.\,.} |P|\,), Q^\infty[\,z\tau \mathinner{.\,.} y_P\,)) \\
&\leq \delta_E(T[\,t_{j+z} \mathinner{.\,.} |T|\,), Q^\infty[\,(j+z)\tau \mathinner{.\,.} y_T\,)) + \kappa + \delta_E(P[\,p_z \mathinner{.\,.} |P|\,), Q^\infty[\,z\tau \mathinner{.\,.} y_P\,)).
\end{aligned}
\tag{5}
$$

Now, observe that we have

$$
\begin{aligned}
d_P &\geq \delta_E(P[\,0 \mathinner{.\,.} p_2\,), Q^\infty[\,0 \mathinner{.\,.} 2\tau\,)) \\
&\quad + \sum_{i=2}^{z-1} \delta_E(P[\,p_i \mathinner{.\,.} p_{i+1}\,), Q^\infty[\,i\tau \mathinner{.\,.} (i+1)\tau\,)) \\
&\quad + \delta_E(P[\,p_z \mathinner{.\,.} |P|\,), Q^\infty[\,z\tau \mathinner{.\,.} y_P\,)).
\end{aligned}
$$

Symmetrically,

$$
\begin{aligned}
d_T &\geq \delta_E(T[\,0 \mathinner{.\,.} t_{j+2}\,), Q^\infty[\,x_T \mathinner{.\,.} (j+2)\tau\,)) \\
&\quad + \sum_{i=2}^{z-1} \delta_E(T[\,t_{j+i} \mathinner{.\,.} t_{j+i+1}\,), Q^\infty[\,(j+i)\tau \mathinner{.\,.} (j+i+1)\tau\,)) \\
&\quad + \delta_E(T[\,t_{j+z} \mathinner{.\,.} |T|\,), Q^\infty[\,(j+z)\tau \mathinner{.\,.} y_T\,)).
\end{aligned}
$$

Adding Equations (3) to (5) hence yields the claimed $\sum_{i=1}^{z} \big||T_{j,i}| - |P_i|\big| \leq 2\kappa + d_P + d_T = 3\kappa - k$. ◾

4.1 Special Puzzle Pieces and How to Compute Them Efficiently

In order for puzzle pieces to be useful to us, we need to be able to efficiently compute the families $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, and $\mathcal{S}_\varphi$ from Definition 4.8. As it turns out, when considering instances of NEWPERIODICMATCHES, most puzzle pieces are substrings of Q^∞—this in turn means that it suffices to locate and compute the *special* pieces (that are different from some specific substrings of Q^∞). We start with a formal definition of special puzzle pieces.

Definition 4.11. *We say that an internal piece is* special *if and only if it is different from $Q^\infty[\,0\,.\,.\,\tau + \Delta\,)$. We write* $\mathsf{Special}(P)$ *for the set of* special internal pieces *of P and we* $\mathsf{Special}(T)$ *for set of* special internal pieces *of T; that is, we set*

$$\mathsf{Special}(P) := \{P_i : i \in (\,1\,.\,.\,z\,) \text{ and } P_i \neq Q^\infty[\,0\,.\,.\,\tau + \Delta\,)\} \quad \text{and}$$
$$\mathsf{Special}(T) := \{T_i : i \in (\,1 + \min J\,.\,.\,z + \max J\,) \text{ and } T_i \neq Q^\infty[\,0\,.\,.\,\tau + \Delta\,)\}.$$

Further, we say that a leading piece $T_{j,1}$ is special *if and only if it is different from $Q^\infty[\,-\kappa\,.\,.\,2\tau + \Delta\,)$ and that a trailing piece $T_{j,z}$ is* special *if and only if it is different from $Q^\infty[\,z\tau\,.\,.\,y_P + \kappa\,)$. Similar to before, we write* $\mathsf{Special}_{\beta\varphi}(T)$ *for the set of special leading and trailing pieces of T.*

Remark 4.12. When we store a special internal piece P_i (or T_i), we store it as a pair (i, P_i) (or (i, T_i)) together with its index i in the sequence of pieces; special leading and trailing pieces are stored similarly.

As a main result of this (sub-)section, we prove that there are only very few special pieces and that we can compute them efficiently.

Lemma 1.6. *The median edit distance of each of the families $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, and $\mathcal{S}_\varphi$ is bounded by $\mathcal{O}(d)$.*
Further, each of the multisets $\mathsf{Special}(P)$, $\mathsf{Special}(T)$, *and* $\mathsf{Special}_{\beta\varphi}(T)$ *is of size $\mathcal{O}(d)$ and can be computed in $\mathcal{O}(d)$ time in the* `PILLAR` *model.*

We proceed to prove Lemma 1.6 by first bounding the edit distances of the puzzle pieces to a common substring of Q^∞; we start with the internal pieces.

Lemma 4.13. *For any $i \in (\,1\,.\,.\,z\,)$, the internal piece P_i from Lemma 4.6 satisfies*

$$\delta_E(P_i, Q^\infty[\,0\,.\,.\,\tau + \Delta\,)) \leq \sum_{j=1}^{13} 2\delta_E(P[\,p_{i+j-1}\,.\,.\,p_{i+j}\,), Q^\infty[\,(i+j-1)\tau\,.\,.\,(i+j)\tau\,)).$$

In particular, the internal piece P_i is special only if any of the fragments $P[\,p_{i+j-1}\,.\,.\,p_{i+j}\,)$ for $j \in [\,1\,.\,.\,13\,]$ differs from $Q^\infty[\,0\,.\,.\,\tau\,)$.
For any $i \in (\,\min J + 1\,.\,.\,\max J + z\,)$, the internal piece T_i from Lemma 4.7 satisfies

$$\delta_E(T_i, Q^\infty[\,0\,.\,.\,\tau + \Delta\,)) \leq \sum_{j=1}^{13} 2\delta_E(T[\,t_{i+j-1}\,.\,.\,t_{i+j}\,), Q^\infty[\,(i+j-1)\tau\,.\,.\,(i+j)\tau\,)).$$

In particular, the internal piece T_i is special only if any of the fragments $T[\,t_{i+j-1}\,.\,.\,t_{i+j}\,)$ for $j \in [\,1\,.\,.\,13\,]$ differs from $Q^\infty[\,0\,.\,.\,\tau\,)$.

Proof. Fix an $i \in (\,1\,.\,.\,z\,)$. Recall that we defined $P_i = P[\,p_i\,.\,.\,p_{i+1} + \Delta\,)$ and $\Delta = 6\kappa \leq 12\tau$. Observe that we have $0 \leq i\tau$ and $(i+15)\tau \leq (\beta_P - 3)\tau \leq y_P$; and thus indeed $(p_{i+j}, (i+j)\tau) \in \mathcal{A}_P$ for all $j \in [\,0\,.\,.\,15\,]$.

Now, exploiting $d_P \leq \kappa \leq 2\tau$ and $(i+15)\tau \leq y_P$, we see that $p_{i+1} + 12\tau \leq p_{i+1} + 14\tau - d_P \leq p_{i+15} \leq |P|$. Hence, we obtain (with the additional help of Fact 2.6)

$$\delta_E(P_i, Q^\infty[0..\tau+\Delta))$$
$$= \delta_E(P[p_i..p_{i+1}+\Delta), Q^\infty[i\tau..(i+1)\tau+\Delta))$$
$$\leq \delta_E(P[p_i..p_{i+1}+12\tau), Q^\infty[i\tau..(i+13)\tau))$$

Extending the fragments $P[p_i..p_{i+1}+12\tau)$ and $Q^\infty[i\tau..(i+13)\tau)$ and correcting for potential length differences, we obtain

$$\delta_E(P_i, Q^\infty[0..\tau+\Delta))$$
$$\leq \delta_E(P[p_i..p_{i+13}), Q^\infty[i\tau..(i+13)\tau)) + |p_{i+13} - (p_{i+1}+12\tau)|$$
$$= \delta_E(P[p_i..p_{i+13}), Q^\infty[i\tau..(i+13)\tau)) + |(p_{i+13} - p_{i+1}) - 12\tau|$$
$$\leq \delta_E(P[p_i..p_{i+13}), Q^\infty[i\tau..(i+13)\tau)) + \delta_E(P[p_{i+1}..p_{i+13}), Q^\infty[(i+1)\tau..(i+13)\tau))$$
$$\leq \sum_{j=1}^{13} 2\delta_E(P[p_{i+j-1}..p_{i+j}), Q^\infty[(i+j-1)\tau..(i+j)\tau)).$$

We proceed similarly for the internal pieces of T. To that end, fix an $i \in (\min J + 1..\max J + z)$ and recall that we set $T_i = T[t_i..t_{i+1}+\Delta)$. As before, observe that we have

$$x_T \leq x_T - \kappa + 2\tau \leq (\lceil (x_T - \kappa)/\tau \rceil + 2)\tau \leq i\tau \quad \text{and}$$
$$(i+15)\tau \leq (\lfloor (y_T - y_P + \kappa)/\tau \rfloor + z + 14)\tau$$
$$\leq y_T - y_P + \kappa + (\beta_P - 3)\tau \leq y_T + \kappa - 2\tau \leq y_T.$$

Thus, indeed, we have $(t_{i+j}, (i+j)\tau) \in \mathcal{A}_T$ for all $j \in [0..15]$. Now, exploiting $d_T \leq \kappa \leq 2\tau$ and $(i+15)\tau \leq y_T$, we see that $t_{i+1} + \Delta \leq t_{i+1} + 14\tau - d_T \leq t_{i+15} \leq |T|$. Hence, analogously to before, we conclude that

$$\delta_E(T_i, Q^\infty[0..\tau+\Delta)) \leq \sum_{j=1}^{13} 2\delta_E(T[t_{i+j-1}..t_{i+j}), Q^\infty[(i+j-1)\tau..(i+j)\tau));$$

completing the proof. ∎

As an immediate corollary, we obtain the desired bound for $\delta_E(\mathcal{S}_\mu)$.

Corollary 4.14. *We have $\delta_E(\mathcal{S}_\mu) \leq 26\kappa$ and $|\mathsf{Special}(T)| + |\mathsf{Special}(P)| \leq 26\kappa$.*

Proof. Recall that $\delta_E(\mathcal{S}_\mu) = \min_{\hat{S}} \sum_{S \in \mathcal{S}_\mu} \delta_E(S, \hat{S})$. Now, by Lemma 4.13, choosing $Q^\infty[0..\tau+\Delta)$ yields the upper bound

$$\delta_E(\mathcal{S}_\mu) \leq \sum_{S \in \mathcal{S}_\mu} \delta_E(S, Q^\infty[0..\tau+\Delta))$$
$$= \sum_{i=2}^{z-1} \delta_E(P_i, Q^\infty[0..\tau+\Delta)) + \sum_{i=\min J+2}^{\max J+z-1} \delta_E(T_i, Q^\infty[0..\tau+\Delta))$$
$$\leq 2 \cdot \sum_{j=1}^{13} \sum_{i=2}^{z-1} \delta_E(P[p_{i+j-1}..p_{i+j}), Q^\infty[(i+j-1)\tau..(i+j)\tau))$$
$$+ \sum_{i=\min J+2}^{\max J+z-1} \delta_E(T[t_{i+j-1}..t_{i+j}), Q^\infty[(i+j-1)\tau..(i+j)\tau))$$
$$\leq 2 \cdot 13 \cdot (d_P + d_T) = 26\kappa.$$

Finally observe that each special piece has an edit distance of at least 1 to $Q^\infty[\,0\mathbin{..}\tau+\Delta\,)$; hence there may be at most 26κ special pieces in total.

Finally, Lemma 4.13 easily yields an algorithm for computing the special internal pieces of P and T.

Corollary 4.15. *Given the alignment $\mathcal{A}_P$, we can compute* $\mathsf{Special}(P)$ *in $\mathcal{O}(d_P+1)$ time in the* `PILLAR` *model; given the alignment $\mathcal{A}_T$, we can compute* $\mathsf{Special}(T)$ *in $\mathcal{O}(d_T+1)$ time in the* `PILLAR` *model.*

Proof. By Lemma 4.13, for each special internal piece P_i (or T_i), there is a $j \in [\,1\mathbin{..}15\,]$ such that $P[\,p_{i+j-1}\mathbin{..}p_{i+j}\,) \neq Q^\infty[\,(i+j-1)\tau\mathbin{..}(i+j)\tau\,)$ (or $T[\,t_{i+j-1}\mathbin{..}t_{i+j}\,) \neq Q^\infty[\,(i+j-1)\tau\mathbin{..}(i+j)\tau\,))$. This, in turn, requires a breakpoint $(a_P, a_Q) \in \mathcal{A}_P$ (or $(a_T, a_Q) \in \mathcal{A}_T$) with $a_Q \in [\,i\tau\mathbin{..}(i+13)\tau\,)$.

Hence, for each breakpoint $(a_P, a_Q) \in \mathcal{A}_P$, we identify pieces P_i with $i \in (\,\lfloor a_Q/\tau\rfloor - 15\mathbin{..}\lfloor a_Q/\tau\rfloor\,] \cap (\,1\mathbin{..}z\,)$ as *candidates*; for each breakpoint $(a_T, a_Q) \in \mathcal{A}_T$, we identify pieces T_i with $i \in (\,\lfloor a_Q/\tau\rfloor - 15\mathbin{..}\lfloor a_Q/\tau\rfloor\,] \cap (\,\min J + 1\mathbin{..}\max J + z\,)$ as *candidates*. For each of these candidates, we check $P_i = Q^\infty[\,0\mathbin{..}\tau+\Delta\,)$ (or $T_i = Q^\infty[\,0\mathbin{..}\tau+\Delta\,)$ using an `LCP` operation of the `PILLAR` model.

By processing the breakpoints in left-to-right order and keeping track of the rightmost candidate generated so far, we can list the indices i of the candidates P_i (or T_i) in left-to-right order without duplicates. Then, a linear-time scan over the representation of $(p_i)_{i=0}^{\beta_P}$ (or $(p_i)_{i=0}^{\beta_T}$) obtained from Lemma 4.4 lets us determine the endpoints of each candidate P_i (or T_i). Each breakpoint yields up to 15 candidates, so the total running time in the `PILLAR` model is $\mathcal{O}(d_P+1)$ (or $\mathcal{O}(d_T+1)$).

We proceed to special leading and trailing pieces—as there is at most one special leading piece of P and at most one special trailing piece of P, we can trivially compute them fast. Hence, we first focus on special leading and trailing peaces of T. To that end, we proceed as with the internal pieces; however, the calculations become slightly more involved.

Lemma 4.16. *For any $j \in J$, the leading piece $T_{j,1}$ from Lemma 4.7 satisfies*

$$\delta_E(T_{j,1}, Q^\infty[\,-\kappa\mathbin{..}2\tau+\Delta\,)) \leq 2\kappa + 2d_T.$$

Further, for any $j \in [\,3\mathbin{..}\max J\,]$, the leading piece $T_{j,1}$ additionally satisfies

$$\delta_E(T_{j,1}, Q^\infty[\,-\kappa\mathbin{..}2\tau+\Delta\,)) \leq 2\sum_{i=-1}^{14} \delta_E(T[\,t_{i+j-1}\mathbin{..}t_{i+j}\,), Q^\infty[\,(i+j-1)\tau\mathbin{..}(i+j)\tau\,)).$$

In particular, the leading piece $T_{j,1}$ is special only if either $j < 3$ or any of the fragments $T[\,t_{i+j-1}\mathbin{..}t_{i+j}\,)$ differs from $Q^\infty[\,(i+j-1)\tau\mathbin{..}(i+j)\tau\,)$ for $i \in [\,-1\mathbin{..}14\,]$.

Proof. Fix a $j \in J$ and recall from Remark 3.7 that $\mathcal{A}_T(r_j) = \max\{x_T, j\tau - \kappa\} \leq j\tau + \kappa$. Observe that we have

$$x_T \leq x_T - \kappa + 2\tau\left(\left\lceil \frac{x_T - \kappa}{\tau}\right\rceil + 2\right)\tau \leq (j+2)\tau \quad \text{and}$$

$$(j+16)\tau \leq \left(\left\lfloor \frac{y_T - y_P + \kappa}{\tau}\right\rfloor + 16\right)\tau \leq y_T - y_P + \kappa + 16\tau$$

$$\leq y_T - (\beta_P - 1)\tau + 2\tau + 16\tau = y_T - (\beta_P - 19)\tau \leq y_T.$$

Thus, indeed, we have $(t_{j+i}, (j+i)\tau) \in \mathcal{A}_T$ for all $i \in [\,2\mathbin{..}16\,]$. Now exploiting $d_T \leq \kappa \leq 2\tau$ and $(j+16)\tau \leq y_T$, we see that $t_{j+2} + \Delta \leq t_{j+2} + 12\tau \leq t_{j+2} + 14\tau - d_T \leq t_{j+16} \leq |T|$. Hence, we conclude

(with the additional help of Fact 2.6)

$$
\begin{aligned}
&\delta_E(T_{j,1}, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta)) \\
&= \delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+2} + \Delta), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+2)\tau + \Delta)) \\
&\leq \delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+2} + 12\tau), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+14)\tau))
\end{aligned}
$$

Replacing strings by superstrings and accounting for potential length differences, we obtain

$$
\begin{aligned}
&\delta_E(T_{j,1}, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta)) \\
&\quad \leq \delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+14}), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+14)\tau)) + |t_{j+14} - (t_{j+2} + 12\tau)| \\
&\quad \leq \delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+2}), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+2)\tau)) + 2\delta_E(T[t_{j+2} \mathinner{\ldotp\ldotp} t_{j+14}), Q^\infty[(j+2)\tau \mathinner{\ldotp\ldotp} (j+14)\tau)) \\
&\quad = \delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+2}), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+2)\tau)) \\
&\qquad + 2\sum_{i=3}^{14} \delta_E(T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau)) \\
&\quad \leq (\mathcal{A}_T(r_j) - (j\tau - \kappa)) + 2d_T \leq 2\kappa + 2d_T.
\end{aligned}
$$

Now, for $j \geq 3$, we have $x_T \leq \tau \leq 3\tau - \kappa \leq j\tau - \kappa$ and hence $\mathcal{A}_T(r_j) = j\tau - \kappa$. Moreover, we have $x_T \leq \tau \leq (j-2)\tau$. Thus, in particular, we have $(t_{j+i}, (j+i)\tau) \in \mathcal{A}_T$ even for all $i \in [-2 \mathinner{\ldotp\ldotp} 16]$. Hence, we conclude similarly to before

$$
\begin{aligned}
&\delta_E(T[r_j \mathinner{\ldotp\ldotp} t_{j+2}), Q^\infty[j\tau - \kappa \mathinner{\ldotp\ldotp} (j+2)\tau)) \\
&\quad \leq \delta_E(T[t_{j-2} \mathinner{\ldotp\ldotp} t_{j+2}), Q^\infty[(j-2)\tau \mathinner{\ldotp\ldotp} (j+2)\tau)) + |(r_j - t_{j-2}) - ((j\tau - \kappa) - (j-2)\tau)| \\
&\quad \leq 2\delta_E(T[t_{j-2} \mathinner{\ldotp\ldotp} t_{j+2}), Q^\infty[(j-2)\tau \mathinner{\ldotp\ldotp} (j+2)\tau)) \\
&\quad = 2\sum_{i=-1}^{2} \delta_E(T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau)).
\end{aligned}
$$

In total, we obtain the claimed

$$
\delta_E(T_{j,1}, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta)) \leq 2\sum_{i=-1}^{14} \delta_E(T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau)). \qquad \blacksquare
$$

Similarly to Lemma 4.16, we analyze the (special) trailing pieces of T.

Lemma 4.17. *For any $j \in J$, the trailing piece $T_{j,z}$ from Lemma 4.7 satisfies*

$$
\delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa)) \leq 2\kappa + d_T.
$$

Further, for any $j \in [\min J \mathinner{\ldotp\ldotp} \beta_T - \beta_P - 3]$, the trailing piece $T_{j,z}$ additionally satisfies

$$
\delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa)) \leq \sum_{i=z+1}^{z+19} \delta_E(T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau)).
$$

In particular, the trailing piece $T_{j,z}$ is special only if either $j > \beta_T - \beta_P - 3$ or any of the fragments $T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j})$ differs from $Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau)$ for $i \in [1 \mathinner{\ldotp\ldotp} z + 19]$.

Proof. Fix a $j \in J$ and recall from Remark 3.7 that $\mathcal{A}_T(r'_j) = \min\{y_T, y_P + j\tau + \kappa\} \geq y_p + j\tau - \kappa$. Observe that we have

$$x_T \leq x_T - \kappa + 2\tau \leq x_T - \kappa + z\tau \leq \left(\left\lceil \frac{x_T - \kappa}{\tau} \right\rceil + z\right)\tau \leq (j+z)\tau) \quad \text{and}$$

$$(j+z)\tau \leq \left(\left\lfloor \frac{y_T - y_P + \kappa}{\tau} \right\rfloor + z\right)\tau \leq y_T - y_P + \kappa + (\beta_P - 17)\tau$$

$$\leq y_T - (\beta_P - 1)\tau + 2\tau + (\beta_P - 17)\tau \leq y_T - 14\tau \leq y_T$$

Thus, indeed, we have $(t_{j+z}, (j+z)\tau) \in \mathcal{A}_T$. Hence, we conclude

$$\delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{.\,.} y_P + \kappa)) = \delta_E(T[t_{j+z} \mathinner{.\,.} r'_j), Q^\infty[(j+z)\tau \mathinner{.\,.} j\tau + y_P + \kappa))$$

$$\leq 2\kappa + \delta_E(T[t_{j+z} \mathinner{.\,.} r'_j), Q^\infty[(j+z)\tau \mathinner{.\,.} \mathcal{A}_T(r'_j)))$$

$$\leq 2\kappa + d_T.$$

Now, for $j \leq \beta_T - \beta_P - 3$, we have $y_P + j\tau + \kappa \leq (\beta_P + j + 2)\tau \leq (\beta_T - 1)\tau \leq y_T$ and hence $\mathcal{A}_T(r'_j) = y_P + j\tau + \kappa$. Moreover, we have $y_P + j\tau + \kappa < (\beta_P + j + 2)\tau = (j + z + 19)\tau \leq y_T$. Thus, in particular, we have $(t_{j+i}, (j+i)\tau) \in \mathcal{A}_T$ even for all $i \in [z \mathinner{.\,.} z + 19]$. Hence,

$$\delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{.\,.} y_P + \kappa)) = \delta_E(T[t_{j+z} \mathinner{.\,.} r'_j), Q^\infty[(j+z)\tau \mathinner{.\,.} j\tau + y_P + \kappa))$$

$$\leq \delta_E(T[t_{j+z} \mathinner{.\,.} t_{j+z+19}), Q^\infty[(j+z)\tau \mathinner{.\,.} (j + z + 19)\tau))$$

$$= \sum_{i=z+1}^{z+19} \delta_E(T[t_{i+j-1} \mathinner{.\,.} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{.\,.} (i+j)\tau));$$

completing the proof. ∎

As with the internal pieces, we proceed to prove the desired bounds on $\delta_E(\mathcal{S}_\beta)$ and $\delta_E(\mathcal{S}_\varphi)$.

Corollary 4.18. *We have $\delta_E(\mathcal{S}_\beta) \leq 53\kappa$, $\delta_E(\mathcal{S}_\varphi) \leq 35\kappa$, and $|\mathsf{Special}_{\beta\varphi}(T)| \leq 88\kappa$.*

Proof. We start with the leading pieces. To that end, first observe that

$$\delta_E(P_1, Q^\infty[-\kappa \mathinner{.\,.} 2\tau + \Delta)) \leq \kappa + \delta_E(P[0 \mathinner{.\,.} p_2 + \Delta), Q^\infty[0 \mathinner{.\,.} 2\tau + \Delta))$$

$$\leq \kappa + \delta_E(P[0 \mathinner{.\,.} p_2 + 12\tau), Q^\infty[0 \mathinner{.\,.} 14\tau))$$

Replacing strings by superstrings and accounting for potential length differences, we obtain

$$\delta_E(P_1, Q^\infty[-\kappa \mathinner{.\,.} 2\tau + \Delta)) \leq \kappa + \delta_E(P[0 \mathinner{.\,.} p_{14}), Q^\infty[0 \mathinner{.\,.} 14\tau)) + |p_{14} - (p_2 - 12\tau)|$$

$$\leq \kappa + \delta_E(P[0 \mathinner{.\,.} p_{14}), Q^\infty[0 \mathinner{.\,.} 14\tau)) + \delta_E(P[p_2 \mathinner{.\,.} p_{14}), Q^\infty[2\tau \mathinner{.\,.} 14\tau))$$

$$\leq \kappa + 2\delta_E(P, Q^\infty[0 \mathinner{.\,.} y_P))$$

$$\leq \kappa + 2d_P.$$

Next, we recall Remark 3.8 and in particular that $\min J \geq -2$. Therefore, Lemma 4.16 yields

$$\sum_{j \in J} \delta_E(T_{j,1}, Q^\infty[-\kappa \mathinner{.\,.} 2\tau + \Delta))$$

$$\leq 5(2\kappa + 2d_T) + 2\sum_{i=-1}^{14} \sum_{j=3}^{\max J} \delta_E(T[t_{i+j-1} \mathinner{.\,.} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{.\,.} (i+j)\tau))$$

$$\leq 10\kappa + 10d_T + 32d_T = 10\kappa + 42d_T.$$

Combined, we obtain the desired bound

$$\delta_E(\mathcal{S}_\beta) \leq \sum_{S \in \mathcal{S}_\beta} \delta_E(S, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta))$$

$$\leq \delta_E(P_1, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta)) + \sum_{j \in J} \delta_E(T_{j,1}, Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta))$$

$$\leq \kappa + 2d_P + 10\kappa + 42d_T \leq 53\kappa.$$

Now observe that each special leading piece has an edit distance of at least 1 to $Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta))$; hence there may be at most 53κ special leading pieces in total.

We proceed to the trailing pieces. To that end, first observe that

$$\delta_E(P_z, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa)) \leq \kappa + \delta_E(P_z, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P)) \leq \kappa + d_P.$$

Next, we recall Remark 3.8 and in particular that $\max J \leq \lceil y_T/\tau \rceil - \lceil y_P/\tau \rceil + 2 = \beta_T - \beta_P + 2$. Therefore, Lemma 4.17 yields

$$\sum_{j \in J} \delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P))$$

$$\leq 5(2\kappa + d_T) + \sum_{i=z+1}^{z+19} \sum_{j=\min J}^{\beta_T - \beta_P - 3} \delta_E(T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}), Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau))$$

$$\leq 10\kappa + 5d_T + 19d_T = 10\kappa + 24d_T.$$

Combined, we obtain the desired bound

$$\delta_E(\mathcal{S}_\varphi) \leq \sum_{S \in \mathcal{S}_\varphi} \delta_E(S, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa))$$

$$\leq \delta_E(P_z, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa)) + \sum_{j \in J} \delta_E(T_{j,z}, Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa))$$

$$\leq \kappa + d_P + 10\kappa + 24d_T \leq 35\kappa.$$

Finally observe that each special trailing piece has an edit distance of at least 1 to $Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa))$; hence there may be at most 35κ special trailing pieces in total. In total, we obtain the claimed bound of $|\mathsf{Special}_{\beta\varphi}(T)| \leq 88\kappa$, completing the proof. ◾

Finally, we discuss how to compute the special leading and trailing pieces of T—again, the algorithm strongly resembles Corollary 4.15 for the special internal pieces.

� **Corollary 4.19.** *Given the alignment $\mathcal{A}_T$, we can compute $\mathsf{Special}_{\beta\varphi}(T)$ in $\mathcal{O}(d_T + 1)$ time in the* PILLAR *model.*

Proof. For computing the special leading pieces, Remark 3.8 and Lemma 4.16 let us focus on the following candidates: at most 5 pieces $T_{j,1}$ with $j \leq 2$, as well as all pieces $T_{j,1}$ such that $T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}) \neq Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau))$ holds for some $i \in [-1 \mathinner{\ldotp\ldotp} 14]$. Similarly for computing the special trailing pieces, Remark 3.8 and Lemma 4.17 let us focus on the following candidates: at most 5 pieces $T_{j,z}$ with $j \geq \beta_T - \beta_P - 2$, as well as all pieces $T_{j,z}$ such that $T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}) \neq Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau))$ holds for some $i \in [z+1 \mathinner{\ldotp\ldotp} z+19]$.

To satisfy $T[t_{i+j-1} \mathinner{\ldotp\ldotp} t_{i+j}) \neq Q^\infty[(i+j-1)\tau \mathinner{\ldotp\ldotp} (i+j)\tau))$, a breakpoint $(a_T, a_Q) \in \mathcal{A}_T$ with $a_Q \in [(i-2)\tau \mathinner{\ldotp\ldotp} (i+14)\tau)$ is needed. Hence, by scanning the breakpoints in the left-to-right order we can

generate all candidates in $\mathcal{O}(d_T + 1)$ time. Next, we retrieve their endpoints using Lemmas 4.4 and 3.12, and we verify each candidate using an `LCP` operation of the `PILLAR` model. The overall running time is $\mathcal{O}(d_T + 1)$.

Recalling and combining the various lemmas of this subsection, we obtain Lemma 1.6, which we restate here for convenience.

Lemma 1.6. *The median edit distance of each of the families $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, and $\mathcal{S}_\varphi$ is bounded by $\mathcal{O}(d)$. Further, each of the multisets $\mathsf{Special}(P)$, $\mathsf{Special}(T)$, and $\mathsf{Special}_{\beta\varphi}(T)$ is of size $\mathcal{O}(d)$ and can be computed in $\mathcal{O}(d)$ time in the `PILLAR` model.*

4.2 Solving NewPeriodicMatches via DynamicPuzzleMatching: A Warm-up Algorithm

In this (sub-)section we obtain a first reduction from NewPeriodicMatches to DynamicPuzzleMatching. While overly naive in nature, it serves as an overview of the general structure of the more involved variants that follow in later (sub-)sections. In particular, in this (sub-)section, we discuss the following (easy) result.

Lemma 4.20. *Given an instance NewPeriodicMatches$(P,T,k,d,Q,\mathcal{A}_P,\mathcal{A}_T)$, after a preprocessing step that takes $\mathcal{O}(d + m/\max\{q,d\})$ time in the `PILLAR` model, we can compute $\mathrm{Occ}_k^E(P,T)$ using a DynamicPuzzleMatching$(k,\Delta,\mathcal{S}_\beta,\mathcal{S}_\mu,\mathcal{S}_\varphi)$ data structure, with $\Delta + \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$, that maintains a DPM-sequence of length $z = \mathcal{O}(m/\max\{q,d\})$ under*
- *$\mathcal{O}(d \cdot m/\max\{q,d\})$ calls to `DPM-Substitute` and*
- *$\mathcal{O}(m/\max\{q,d\})$ calls to `DPM-Query`.*

Consider a given instance NewPeriodicMatches $(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$ and write $\kappa := k + d_P + d_T$ and $\tau := q\lceil\kappa/2q\rceil$. Further, as before, let $P = \bigodot_{i=1}^{\beta_P} P[p_{i-1}..p_i)$ denote the τ-tile partition of P with respect to $\mathcal{A}_P$; let $T = \bigodot_{i=1}^{\beta_T} T[t_{i-1}..t_i)$ denote the τ-tile partition of T with respect to $\mathcal{A}_T$; and set $\Delta := 6\kappa$ and $z := \beta_P - 17$. Finally, we define the families of puzzle pieces according to Definition 4.8.

Recalling Remark 4.2 and Lemma 1.6, we readily confirm that our choices for Δ, $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, and $\mathcal{S}_\varphi$ indeed satisfy $z = \mathcal{O}(m/\max\{q,d\})$, $\Delta = \mathcal{O}(d)$, and $\delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$.

Intuitively, we proceed as follows: we initialize a DynamicPuzzleMatching$(k,\Delta,\mathcal{S}_\beta,\mathcal{S}_\mu,\mathcal{S}_\varphi)$ data structure with a DPM-sequence $\mathcal{I}$ of length z that represents P and $R_{\min J}$, and call `DPM-Query` to obtain $\mathrm{Occ}_k^E(\mathcal{I}_{\min J}) = \mathrm{Occ}_k^E(P, R_{\min J})$, that is, the k-error occurrences of P in $R_{\min J}$. Then, iterating over J, for a $j \in (\min J .. \max J]$, we use `DPM-Substitute` operations to transform $\mathcal{I}$ into a DPM-sequence that represents P and R_j; to then call `DPM-Query` to $\mathrm{Occ}_k^E(\mathcal{I}_j) = \mathrm{Occ}_k^E(P, R_j)$, that is, obtain the k-error occurrences of P in R_j.

Observe that by Corollary 3.11, we indeed obtain all k-error occurrences of P in T. As we have $|J| = \mathcal{O}(\beta_P) = \mathcal{O}(m/\max\{q,d\})$ (see Lemma 4.5), we can already see that the number of calls to `DPM-Query` is just as promised in Lemma 4.20. For a bound on the number of calls to `DPM-Substitute`, we need to be more precise on how we transform the DPM-sequence $\mathcal{I}$. We start with formally defining the DPM-sequence(s) $\mathcal{I}$.

Definition 4.21. *For each $j \in J$, we set $\mathcal{I}_j := (P_1, T_{j,1})(P_2, T_{j,2})\cdots(P_z, T_{j,z})$.*
Recall that by Lemmas 4.6 and 4.7 we have $\mathsf{val}_\Delta(P_1,\ldots,P_z) = P$ and $\mathsf{val}_\Delta(T_{j,1},\ldots,T_{j,z}) = R_j$.

Let us write $\mathcal{I} = (U_1, V_1)(U_2, V_2)\ldots(U_z, V_z)$ for the DPM-sequence maintained in the DynamicPuzzleMatching data structure. We initialize $\mathcal{I}$ with the DPM-sequence $\mathcal{I} \leftarrow \mathcal{I}_{\min J}$. Then, we iterate over J (starting from $\min J$) as follows.

- If $j > \min J$, for each $i \in [\,1\,.\,.\,z\,]$ with $T_{j-1,i} \neq T_{j,i}$, we substitute the i-th pair of strings with $(P_i, T_{j,i})$. We perform these updates in non-decreasing order with respect to $\big||T_{j,i}| - |P_i|\big| - \big||T_{j-1,i}| - |P_i|\big|$.
- We call `DPM-Query` to obtain $\mathrm{Occ}_k^E(\mathcal{I}_j) = \mathrm{Occ}_k^E(P, R_j)$; and we set $\mathrm{Occ}_k^E(P, T) \coloneqq \mathrm{Occ}_k^E(P, T) \cup (r_j + \mathrm{Occ}_k^E(P, R_j))$.

Verifying the Conditions of DynamicPuzzleMatching

Before we discuss the number of required calls to `DPM-Substitute`, we have to convince ourselves that $\mathcal{I}$ fulfills the constraints of DYNAMICPUZZLEMATCHING on the maintained DPM-sequence. In particular, we need to show that the following conditions are satisfied at all times.

(a) $U_1, V_1 \in \mathcal{S}_\beta$ and $U_z, V_z \in \mathcal{S}_\varphi$, and, for all $i \in (\,1\,.\,.\,z\,)$, $U_i, V_j \in \mathcal{S}_\mu$

(b) $\mathrm{tor}(\mathcal{I}) = \sum_{i=1}^z \big||U_i| - |V_i|\big| \leq \Delta/2 - k$.

Observe that before any `DPM-Query`, the DPM-sequence $\mathcal{I}$ is equal to the DPM-sequence $\mathcal{I}_j$. Now, while Condition (a) is satisfied by construction, we need to work slightly harder for Condition (b). First, if we have $\mathcal{I} = \mathcal{I}_j$ for some $j \in J$, then Lemma 4.10 yields

$$\mathrm{tor}(\mathcal{I}) = \mathrm{tor}(\mathcal{I}_j) = \sum_{i=1}^z \big||P_i| - |T_{j,i}|\big| \leq 3\kappa - k = \Delta/2 - k.$$

Next, consider the sequence of updates $\mathcal{U}$ between two sequences $\mathcal{I}_j$ and $\mathcal{I}_{j+1}$. Observe that by construction, we can split $\mathcal{U}$ into two disjoint parts: first the updates $\mathcal{U}_\leq$ that do not increase the torsion $\mathrm{tor}(\mathcal{I}_j)$, followed by the updates $\mathcal{U}_>$ that do. In particular, we can apply the following useful fact, which directly yields the desired properties.

▞ **Fact 4.22.** *Let C denote a non-negative integer and let $\mathcal{I}$ and $\mathcal{I}'$ denote sequences of length z that satisfy $\mathrm{tor}(\mathcal{I}) \leq C$ and $\mathrm{tor}(\mathcal{I}') \leq C$. Consider a sequence $\mathcal{U}$ of updates that*

- *transforms $\mathcal{I}$ into $\mathcal{I}'$, and*
- *is ordered such that all updates that do not increase the torsion of the DPM-sequence precede all updates that increase the torsion of the DPM-sequence.*

Then, any DPM-sequence $\mathcal{I}''$ obtained from $\mathcal{I}$ by applying a prefix of $\mathcal{U}$ satisfies $\mathrm{tor}(\mathcal{I}'') \leq C$. ◾

Bounding the Number of Calls to `DPM-Substitute`

We return to bounding the number of calls to `DPM-*` operations. In particular, our next goal is to upper-bound the number of pairs that our algorithm substitutes to transform $\mathcal{I}_{j-1}$ to $\mathcal{I}_j$ for some $j \in (\,\min J\,.\,.\,\max J\,]$.

To that end, observe that we might need to substitute the head or the tail of the DPM-sequence. Further, for each $i \in (\,1\,.\,.\,z\,)$, the i-th pair is substituted only if at least one of $T_{j-1,i}$ and $T_{j,i}$ differs from $Q^\infty[\,0\,.\,.\,\tau + \Delta\,)$; that is, if either $T_{j-1,i}$ or $T_{j,i}$ is special. Now, recall from Lemma 1.6 that the set $\mathsf{Special}(T) \cup \mathsf{Special}(P)$ is of size $\mathcal{O}(d)$—thus, for each $j \in J$, we perform $\mathcal{O}(d)$ substitutions.

In total, we hence perform $\mathcal{O}(\kappa|J|) = \mathcal{O}(\kappa \cdot m/\tau) = \mathcal{O}(d \cdot m/\max\{q, d\})$ calls to `DPM-Substitute`; recall that earlier we already bounded the number of calls to `DPM-Update` by $\mathcal{O}(|J|) = \mathcal{O}(m/\max\{q, d\})$.

Analyzing the Preprocessing Time

Now for bounding the preprocessing time, we need to argue about the time required for constructing the initial DPM-sequence $\mathcal{I}_{\min J}$, as well as the time required to compute the substitutions that are to be performed.

For computing $\mathcal{I}_{\min J}$ we turn to Lemmas 4.4 and 3.12, which yield the sequences $(r_j)_{j \in J}$, $(r'_j)_{j \in J}$, $(p_i)_{i=0}^{\beta_P}$, and $(t_i)_{i=0}^{\beta_T}$, from which we can easily obtain $\mathcal{I}_{\min J}$ in time $\mathcal{O}(z + \kappa) = \mathcal{O}(m/\max\{q, d\} + d)$.

For computing which pieces to update, Lemma 1.6 yields the special pieces of P and T in a suitable representation in time $\mathcal{O}(d)$ (in the `PILLAR` model).

To summarize our result, let us recall Lemma 4.20, which we have just proved.

Lemma 4.20. *Given an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$, after a preprocessing step that takes $\mathcal{O}(d + m/\max\{q, d\})$ time in the* `PILLAR` *model, we can compute* $\mathrm{Occ}_k^E(P, T)$ *using a DYNAMICPUZZLEMATCHING$(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ data structure, with $\Delta + \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$, that maintains a DPM-sequence of length $z = \mathcal{O}(m/\max\{q, d\})$ under*

- $\mathcal{O}(d \cdot m/\max\{q, d\})$ *calls to* `DPM-Substitute` *and*
- $\mathcal{O}(m/\max\{q, d\})$ *calls to* `DPM-Query`.

Observe that combining Theorem 1.5 and Lemma 4.20, we obtain an algorithm for NEWPERIODIC-MATCHES that requires $\widetilde{\mathcal{O}}(d^3 + \Delta^2 d) = \widetilde{\mathcal{O}}(d^3)$ time for preprocessing, $\widetilde{\mathcal{O}}(\Delta z) = \widetilde{\mathcal{O}}(d \cdot m/\max\{q, d\})$ time for initialization, and $\widetilde{\mathcal{O}}(\Delta \kappa \cdot m/\tau) = \mathcal{O}(d^2 \cdot m/\max\{q, d\})$ time for processing updates and queries. This yields an overall running time of $\widetilde{\mathcal{O}}(d^3 + d^2 \cdot m/\max\{q, d\})$, which is by far too much. In particular, in the remainder of this section (and, by extension, this part), we aim to replace the factor $m/\max\{q, d\}$ with a *small* power of d.

4.3 Solving NewPeriodicMatches via DynamicPuzzleMatching, Improvement 0: Replacing Pair Substitutions with Pair Insertions and Pair Deletions

In our quest to replace the factor $\mathcal{O}(m/\tau) = \mathcal{O}(m/\max\{q, d\})$ in Lemma 4.20 by a small power of d, we take a small detour to give an alternative algorithm for obtaining $\mathcal{I}_j$ from $\mathcal{I}_{j-1}$. In particular, observe that up until now, we have not used the `DPM-Delete` and `DPM-Insert` operations, as we could make do with just calling `DPM-Substitute`. While we do not decrease the overall number of calls to `DPM-*`, in this (sub-)section, we decrease the number of such calls that involve special pieces to $\mathcal{O}(\kappa^2) = \mathcal{O}(d^2)$. Formally, we call an internal pair of pieces in $\mathcal{I}$ *canonical* if it does not contain a special piece.

Definition 4.23. *For an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$, write $\Delta := 6(k + d_P + d_T)$ and $\tau := q\lceil \kappa/2q \rceil$. We call a pair of pieces* canonical *if it equals*

$$\mathcal{Q} := (Q^\infty[0 \mathinner{.\,.} \tau + \Delta), Q^\infty[0 \mathinner{.\,.} \tau + \Delta)).$$

Now, we obtain the following variant of Lemma 4.20.

Lemma 4.24. *Given an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$, after a preprocessing step that takes $\mathcal{O}(d + m/\max\{q, d\})$ time in the* `PILLAR` *model, we can compute* $\mathrm{Occ}_k^E(P, T)$ *using a DYNAMICPUZZLEMATCHING$(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ data structure, with $\Delta + \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$, that maintains a DPM-sequence of length $z = \mathcal{O}(m/\max\{q, d\})$ under*

- $\mathcal{O}(d \cdot m/\max\{q, d\})$ *calls to* `DPM-Delete` *and* `DPM-Insert` *that delete and insert, respectively, a canonical pair,*
- $\mathcal{O}(d^2)$ *calls to* `DPM-Substitute`, *and*
- $\mathcal{O}(m/\max\{q, d\})$ *calls to* `DPM-Query`.

Proof. In general, we follow the algorithm from Lemma 4.20; however, we use a different method for transforming $\mathcal{I}_{j-1}$ to $\mathcal{I}_j$ (for $j \in (\min J \mathinner{.\,.} \max J]$).

First, observe that, by Lemmas 4.16 and 4.17, we need to perform $\mathcal{O}(\kappa) = \mathcal{O}(d)$ calls to `DPM-Substitute` for the head and the tail of the maintained DPM-sequence $\mathcal{I}$. We can thus focus on updates involving pairs of internal pieces.

Now, as an illustrative example, consider an internal pair $\mathcal{P}_i^{j-1} \coloneqq (P_i, T_{j-1,i})$, and suppose that only $T_{j-1,i}$ is special (that is, different from $Q^\infty[0\mathinner{.\,.}\Delta + \tau)$). Further, suppose that the pairs $\mathcal{P}_{i-1}^{j-1} \coloneqq (P_{i-1}, T_{j-1,i-1})$ and $\mathcal{P}_{i+1}^{j-1} \coloneqq (P_{i+1}, T_{j-1,i+1})$ are both canonical. Now, observe that in $\mathcal{I}_j$, all pieces of P get aligned to a piece of T that is one piece to the right compared to $\mathcal{I}_{j-1}$, that is, we need to construct the pairs

$$\mathcal{P}_{i-1}^j \coloneqq (P_{i-1}, T_{j,i-1}) = (P_{i-1}, T_{j-1,i}) = \mathcal{P}_i^{j-1} \quad \text{and} \quad \mathcal{P}_i^j \coloneqq (P_i, T_{j,i}) = (P_i, T_{j-1,i+1}) = \mathcal{Q}.$$

Now, instead of calling `DPM-Substitute` to directly create $\mathcal{P}_{i-1}^j$ and $\mathcal{P}_i^j$ (as we did in Lemma 4.20), we `DPM-Delete` the canonical pair $\mathcal{P}_{i-1}^{j-1}$ and `DPM-Insert` it after (the unchanged) $\mathcal{P}_{i-1}^j = \mathcal{P}_i^{j-1}$. Observe that by doing so, we replaced the previous calls to `DPM-Substitute` with an equal number of calls to `DPM-Delete` and `DPM-Insert` involving copies of $\mathcal{Q}$.

Generalizing the above example, unless both pieces P_i and $T_{j-1,i}$ (or both P_i and $T_{j,i}$) are special (and if $i \in (2\mathinner{.\,.}z-1)$), we can replace the calls to `DPM-Substitute` with an equal number of calls to `DPM-Delete` and `DPM-Insert` that delete and insert a copy of $\mathcal{Q}$, respectively. Now, by Lemma 1.6, there are at most $\mathcal{O}(d)$ special pieces in P and T each—hence, over all $\mathcal{I}_j, j \in J$ in total, only at most $\mathcal{O}(d^2)$ pairs contain two special pieces. Replacing all other calls to `DPM-Substitute`, we obtain the claimed result. ◢

Remark 4.25. Observe that in the worst case, no algorithm can iterate over all $\mathcal{I}_j$ for $j \in J$ with $o(\kappa \cdot m/\tau)$ updates. Consider the example, where for some constant c, every $\lfloor cz/\kappa \rfloor$-th piece within each of the sequences $P_2, \ldots, P_{z-1}$ and $T_{\min J+2}, \ldots, T_{\max J+z-1}$ is special and no two special pieces are equal. We can easily convince ourselves that for each pair $j_1, j_2 \in J$, we have $\delta_E(\mathcal{I}_{j_1}, \mathcal{I}_{j_2}) = \Omega(\kappa)$ and hence $\Omega(\kappa \cdot m/\tau)$ updates are required for exactly computing each $\mathcal{I}_j$. ◢

Remark 4.26. Observe that for each $j \in J$, the DPM-sequence $\mathcal{I}_j$ without its head and tail consists of $\mathcal{O}(\kappa)$ pairs that contain at least one special piece and $\mathcal{O}(\kappa)$ runs of canonical pairs. Then, roughly speaking the following claims hold (which we prove—in a more general setting—in Section 4.4): each run of canonical pairs in $\mathcal{I}_{j-1}$ remains intact in $\mathcal{I}_j$ up to the potential insertion or deletion of a canonical pair. In addition, we can decide whether we need any of these two potential updates by inspecting the pairs adjacent to the run in scope.

Thus, in the transformation of $\mathcal{I}_{j-1}$ to $\mathcal{I}_j$ we have to insert or delete $\mathcal{O}(\kappa)$ canonical pairs. The crucial observation in Section 4.4 now is that for very long runs of $\mathcal{Q}$, inserting or deleting a single copy of $\mathcal{Q}$ *does not matter*, that is, the resulting k-error occurrences are essentially the same. This allows us to skip many of the calls to `DPM-Insert` and `DPM-Delete`, decreasing their number to a power of $\mathcal{O}(d)$. Observe that this allows us to circumvent Remark 4.25, as we compute the sequences $\mathcal{I}_j$ only approximately. ◢

4.4 Solving NewPeriodicMatches via DynamicPuzzleMatching, Improvement 1: Trimming Long Perfectly Periodic Segments

In this (sub-)section, we extend and formalize the idea that Remark 4.26 hinted at. In particular, we exploit that insertions and deletions of canonical pairs $\mathcal{Q}$ that extend or shrink a "long" run of canonical pairs do not alter the answer to a DynamicPuzzleMatching query. Based on this, we iterate over "compressed" versions of the sequences in scope, while still being able to compute all sought approximate occurrences. To simplify our exposition, let us define a Trim operator for DPM-sequences.

Definition 4.27. *Consider a DPM-sequence $\mathcal{I}$, where the elements of some subset $\mathsf{Plain}(\mathcal{I})$ of the internal canonical DPM-pairs are labeled as* plain. *(We call all other pieces* non-plain.*)*

For a positive integer α, we write $\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)$ for the DPM-sequence obtained from $\mathcal{I}$ by removing exactly one plain DPM-pair from any one contiguous subsequence of plain DPM-pairs of length at least $\alpha+1$; if $\mathcal{I}$ does not contain any such subsequence, we set $\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha) := \mathcal{I}$. (That is, $\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)$ is a DPM-sequence of length at least $z-1$.)

Further, we write $\mathsf{Trim}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha) := \mathsf{Tm}^{\star}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)$ for an iterated application of Tm until the DPM-sequence remains unchanged. (That is, in $\mathsf{Trim}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)$ every contiguous subsequence of plain DPM-pairs is of length at most α.) ◾

Remark 4.28. For technical reasons, we may not want to include all canonical pairs $\mathcal{Q}$ in the set $\mathsf{Plain}(\mathcal{I})$. In particular, write $\mathsf{Red}(P) \supseteq \mathsf{Special}(P)$ for a set of *red* pieces of P and write $\mathsf{Red}(T) \supseteq \mathsf{Special}(T)$ for a set of red pieces of T. Now, for a $j \in J$, we set

$$\mathsf{Plain}(\mathcal{I}_j) := \{(P_i, T_{j,i}) : i \in (1\,..\,z) \text{ and } P_i \notin \mathsf{Red}(P) \text{ and } T_{j,i} \notin \mathsf{Red}(T).$$

Abusing notation, we write $\mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), \alpha) := \mathsf{Trim}(\mathcal{I}_j, \mathsf{Plain}(\mathcal{I}_j), \alpha)$. ◾

Remark 4.29. It is easy to see that for every $j \in J$, we have $|\mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), \alpha)| = \mathcal{O}(\alpha|\mathsf{Red}(P)|\,|\mathsf{Red}T|)$: in the DPM-sequence $\mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), \alpha)$, the set $\mathsf{Plain}(\mathcal{I}_j)$ contains at most $\mathcal{O}(|\mathsf{Red}(P)|\,|\mathsf{Red}(T)|)$ sequences of contiguous pairs, each of length of at most α. In particular, we have $|\mathsf{Trim}(\mathcal{I}_j, \mathsf{Special}(P), \mathsf{Special}(T), \alpha)| = \mathcal{O}(d^2\alpha)$ by Lemma 1.6. ◾

As a key result of this (sub-)section, we proceed to show that we can trim the DPM-sequences $\mathcal{I}_j$ with $\alpha = k+1$ and still recover all k-error occurrences of P in R_j.

Lemma 4.30. *For any $j \in J$, we have*

$$\mathrm{Occ}_k^E(P, R_j) = \mathrm{Occ}_k^E(\mathcal{I}_j) = \mathrm{Occ}_k^E(\mathsf{Trim}(\mathcal{I}_j, \mathsf{Special}(P), \mathsf{Special}(T), k+1)).$$ ◾

Now, we can concisely state the main algorithmic result of this (sub-)section.

Lemma 4.31. *Given an instance $\textsc{NewPeriodicMatches}(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$, after a preprocessing step that takes $\mathcal{O}(d^3\alpha \log n)$ time in the $\mathtt{PILLAR}$ model model, we can compute*

$$\mathcal{G} := \bigcup_{j \in J} (r_j + \mathrm{Occ}_k^E(\mathsf{Trim}(\mathcal{I}_j, \mathsf{Special}(P), \mathsf{Special}(T), k+1)))$$

as $\mathcal{O}(d^4)$ arithmetic progressions with difference $\tau = \mathcal{O}(\max\{q, d\})$, using a $\textsc{DynamicPuzzleMatching}$ $(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ data structure, with $\Delta + \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$, that maintains a DPM-sequence of length $z = \mathcal{O}(d\alpha)$ under
- $\mathcal{O}(d^3)$ *calls to* $\mathtt{DPM\text{-}Delete}$ *and* $\mathtt{DPM\text{-}Insert}$ *that delete and insert, respectively, a canonical pair,*
- $\mathcal{O}(d^2)$ *calls to* $\mathtt{DPM\text{-}Substitute}$*, and*
- $\mathcal{O}(d^2)$ *calls to* $\mathtt{DPM\text{-}Query}$*.*

◾

We proceed with a proof of Lemma 4.30; formally, we prove the following, slightly stronger statement.

Lemma 4.32. *Fix a string $\hat{Q}$ and integers $k \geq 0$ and $\Delta > 0$. Further, consider a DPM-sequence $\mathcal{I}$ whose pieces form Δ-puzzles and that has a torsion of $\mathrm{tor}(\mathcal{I}) \leq \Delta/2 - k$, and a set $\mathsf{Plain}(\mathcal{I})$ of DPM-pairs labeled as plain, where $\mathsf{Plain}(\mathcal{I}) \subseteq \{(U_i, V_i) \in \mathcal{I} : i \in (1\,..\,z) \text{ and } U_i = V_i = \hat{Q}\}$.*

For any $\alpha \geq 1$, we have $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)) \supseteq \mathrm{Occ}_k^E(\mathcal{I})$ and $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), k+\alpha)) = \mathrm{Occ}_k^E(\mathcal{I})$.

Proof. Let us assume that the Tm-operation indeed removes a plain piece of $\mathcal{I}$, otherwise there is nothing to prove. Further, let us fix an $\alpha \geq 1$. Now, first, we convince ourselves that the Tm operation produces valid Δ-puzzles.

☞ **Claim 4.33.** Given a Δ-puzzle $\mathcal{X} := X_1, \ldots, X_i, X_{i+1}, \ldots, X_z$ with $X_i = X_{i+1}$ and value $X := \mathsf{val}_\Delta(\mathcal{X})$, the sequence

$$\mathcal{X}' := X_1, \ldots, X_{i-1}, X_{i+1}, \ldots, X_z = X_1, \ldots, X_i, X_{i+2}, \ldots, X_z$$

forms a Δ-puzzle with value

$$\mathsf{val}_\Delta(\mathcal{X}') = X[\, 0 \mathbin{..} \xi_i + \lfloor \Delta/2 \rfloor \,) X[\, \xi_i + |X_i| - \lceil \Delta/2 \rceil \mathbin{..} |X| \,) = X[\, 0 \mathbin{..} \xi_i + p \,) X[\, \xi_{i+1} + p \mathbin{..} |X| \,),$$

where $\xi_i := \sum_{t=1}^{i-1} |X_t| - \Delta$ and $p \in [\, 0 \mathbin{..} |X_i| - \Delta \,]$.

Proof. By definition, we have $X_i[\, |X_i| - \Delta \mathbin{..} |X_i| \,) = X_{i+1}[\, 0 \mathbin{..} \Delta \,)$ and $X_{i+1}[\, |X_{i+1}| - \Delta \mathbin{..} |X_{i+1}| \,) = X_{i+2}[\, 0 \mathbin{..} \Delta \,)$. As $X_i = X_{i+1}$, we also have $X_i[\, |X_i| - \Delta \mathbin{..} |X_i| \,) = X_{i+2}[\, 0 \mathbin{..} \Delta \,)$, yielding the claim. ⌟

Now, let us write $\mathcal{I} = (U_1, V_1)(U_2, V_2) \cdots (U_z, V_z)$ and $U := \mathsf{val}_\Delta(U_1, \ldots, U_z)$ and $V := \mathsf{val}_\Delta(V_1, \ldots, V_z)$. Further, for each $j \in (\, 1 \mathbin{..} z \,)$, write $U[\, u_j \mathbin{..} u'_j \,) = U_j$ and $V[\, v_j \mathbin{..} v'_j \,) = V_j$; that is $u_j = \sum_{t=1}^{j-1}(|U_t| - \Delta)$ and $v_j = \sum_{t=1}^{j-1}(|V_t| - \Delta)$. With Claim 4.33 in mind, we also write $\hat{U}_j := U[\, \hat{u}_j \mathbin{..} \hat{u}'_j \,)$, where

$$\hat{u}_j := u_j + \lfloor \Delta/2 \rfloor = \sum_{t=1}^{j-1}(|U_t| - \Delta) + \lfloor \Delta/2 \rfloor \qquad \text{and}$$

$$\hat{u}'_j := \hat{u}_j + |U_j| - \Delta = \sum_{t=1}^{j-1}(|U_t| - \Delta) + |U_j| - \lceil \Delta/2 \rceil = u'_j - \lceil \Delta/2 \rceil.$$

For convenience, we define $\hat{u}'_1$ and $\hat{u}_z$ analogously. Observe that we have $\hat{u}'_j = \hat{u}_{j+1}$, that is, we obtain a partition of $U = U[\, 0 \mathbin{..} \hat{u}'_1 \,) \hat{U}_2 \cdots \hat{U}_{z-1} U[\, \hat{u}_z \mathbin{..} |U| \,)$.

We proceed to show $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)) \supseteq \mathrm{Occ}_k^E(\mathcal{I})$. To that end, suppose that the Tm-operation removes the DPM-pair (U_i, V_i) from $\mathcal{I}$. We may assume that $U_i = U_{i+1} = V_i = V_{i+1}$, as the Tm-operation may not fully remove a subsequence of plain DPM-pairs and we can relabel the remaining plain pair, if necessary. In particular, we can interpret $\hat{U}_i$ as the fragment cut out of U due to the Tm-operation.

Now, write

$$U' := \mathsf{val}_\Delta(U_1, \ldots, U_{i-1}, U_{i+1}, \ldots, U_z) \quad \text{and} \quad V' := \mathsf{val}_\Delta(V_1, \ldots, V_{i-1}, V_{i+1}, \ldots, V_z),$$

and consider an optimal alignment $\mathcal{A} : U \rightsquigarrow V[\, a \mathbin{..} b \,)$ of cost at most k. As $\mathcal{I}$ has a bounded torsion, we can see that $\mathcal{A}$ indeed aligns $\hat{U}_i$ within V_i—in fact, it does so for every fragment $\hat{U}_i$.

☞ **Claim 4.34.** For any $j \in (\, 1 \mathbin{..} z \,)$, the fragment $\hat{V}_j := V[\, \hat{v}_j \mathbin{..} \hat{v}'_j \,) := \mathcal{A}(\hat{U}_j)$ is contained in V_j.

Proof. Observe that we have by construction

$$v_j = u_j + \sum_{t=1}^{j-1}(|V_t| - |U_t|) \leq u_j + \sum_{t=1}^{z} \bigl||U_t| - |V_t|\bigr| = u_j + \mathrm{tor}(\mathcal{I}) \leq u_j + \lfloor \Delta/2 \rfloor - k \leq a + \hat{u}_j - k \leq \hat{v}_j.$$

Symmetrically, we obtain

$$\hat{v}'_j \leq b + \hat{u}'_j - |U| + k \leq |V| - |U| + u'_j - \lceil \Delta/2 \rceil + k$$

$$\leq |V| - |U| + u'_j - \mathrm{tor}(\mathcal{I}) = |V| - |U| + u'_j - \sum_{t=1}^{z}\bigl||U_t| - |V_t|\bigr|)$$

$$\leq |V| - |U| + u'_j - \sum_{t=j+1}^{z}(|V_t| - |U_t|) = v'_j,$$

thus completing the proof. ⌟

Next, we use Corollary 2.7 to cut out the pair (U_i, V_i) from $\mathcal{A}$; we have

$$\delta_E(U, V[a \mathinner{..} b)) \geq \delta_E(U[0 \mathinner{..} \hat{u}_i), V[a \mathinner{..} \hat{v}_i)) + \delta_E(U[\hat{u}_i + |U_i| - \Delta \mathinner{..} |U|), V[\hat{v}_i + |U_i| + \Delta \mathinner{..} b))$$

Now, Claim 4.34 ensures that $\hat{V}_i$ is indeed contained in V_i. Adding that $V_i[0 \mathinner{..} |U_i - \Delta)V_{i+1}$ has a period $|U_i| - \Delta$, we see that $V[\hat{v}_i \mathinner{..} \hat{v}_i + |U_i| - \Delta)$ is a rotation of $V_i[0 \mathinner{..} |U_i| - \Delta)$. Thus, we indeed obtain V':

$$
\begin{aligned}
&\delta_E(U[0 \mathinner{..} \hat{u}_i), V[a \mathinner{..} \hat{v}_i)) + \delta_E(U[\hat{u}_i + |U_i| - \Delta \mathinner{..} |U|), V[\hat{v}_i + |U_i| + \Delta \mathinner{..} b)) \\
&\quad = \delta_E(U'[0 \mathinner{..} \hat{u}_i), V'[a \mathinner{..} \hat{v}_i)) + \delta_E(U'[\hat{u}_i \mathinner{..} |U'|), V'[\hat{v}_i \mathinner{..} b - |U_i| + \Delta)) \\
&\quad \geq \delta_E(U', V'[a \mathinner{..} b - |U_i| + \Delta)).
\end{aligned}
$$

In total, we constructed an alignment $\mathcal{A}' : U' \rightsquigarrow V'[a \mathinner{..} b - |U_i| + \Delta)$ of cost at most k—and hence completed the proof of $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), \alpha)) \supseteq \mathrm{Occ}_k^E(\mathcal{I})$.

Now, to prove $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), k+\alpha)) = \mathrm{Occ}_k^E(\mathcal{I})$, and more specifically $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), k+\alpha)) \subseteq \mathrm{Occ}_k^E(\mathcal{I})$, we first observe that in $\mathcal{I}$, the DPM-pair (U_i, V_i) is part of a contiguous subsequence $\mathcal{X} := (U_{i-k'}, V_{i-k'}) \cdots (U_i, V_i) \cdots (U_{i+k''}, V_{i+k''})$ of length (at least) $k'+k''+1 \geq k+2$, where all DPM-pairs that are equal to (U_i, V_i). Again, write

$$U' := \mathsf{val}_\Delta(U_1, \ldots, U_{i-1}, U_{i+1}, \ldots, U_z) \quad \text{and} \quad V' := \mathsf{val}_\Delta(V_1, \ldots, V_{i-1}, V_{i+1}, \ldots, V_z),$$

and consider an optimal alignment $\mathcal{A} : U' \rightsquigarrow V'[a \mathinner{..} b)$ of cost at most k. As the alignment $\mathcal{A}$ can make at most k edits, at least one DPM-pair of the contiguous subsequence

$$\mathcal{X}' := (U_{i-k'}, V_{i-k'}) \cdots (U_{i-1}, V_{i-1})(U_{i+1}, V_{i+1}) \cdots (U_{i+k''}, V_{i+k''})$$

gets aligned without any edits; as all pairs of $\mathcal{X}$ are the same, we may thus assume without loss of generality that $\hat{V}_{i-1} = \hat{U}_{i-1}$. In particular, we can insert (U_i, V_i) into $\mathcal{X}'$ after (U_{i-1}, V_{i-1}) such that $\hat{U}_i = \hat{V}_i$; that is (using Fact 2.2), we have

$$
\begin{aligned}
\delta_E(U', V'[a \mathinner{..} b)) &= \delta_E(U'[0 \mathinner{..} \hat{u}_i), V'[a \mathinner{..} \hat{v}_i)) + 0 + \delta_E(U'[\hat{u}_i' \mathinner{..} |U|), V'[\hat{v}_i' \mathinner{..} b)) \\
&= \delta_E(U[0 \mathinner{..} \hat{u}_i), V[a \mathinner{..} \hat{v}_i)) + 0 \\
&\quad + \delta_E(U[\hat{u}_i' + |U_i| - \Delta \mathinner{..} |U|), V[\hat{v}_i' + |U_i| - \Delta \mathinner{..} b + |U_i| - \Delta)) \\
&\geq \delta_E(U, V[a \mathinner{..} b + |U_i| - \Delta)).
\end{aligned}
$$

In total, we constructed an alignment $\mathcal{A}' : U \rightsquigarrow V[a \mathinner{..} b + |U_i| - \Delta)$ of cost at most k—and hence completed the proof of $\mathrm{Occ}_k^E(\mathsf{Tm}(\mathcal{I}, \mathsf{Plain}(\mathcal{I}), k+\alpha)) \subseteq \mathrm{Occ}_k^E(\mathcal{I})$. ⌟

Combining Lemmas 4.10 and 4.32, we obtain Lemma 4.30, which we restate here for convenience.

▰ **Lemma 4.30**. *For any $j \in J$, we have*

$$\mathrm{Occ}_k^E(P, R_j) = \mathrm{Occ}_k^E(\mathcal{I}_j) = \mathrm{Occ}_k^E(\mathit{Trim}(\mathcal{I}_j, \mathit{Special}(P), \mathit{Special}(T), k+1)).$$
⌟

▰ **Remark 4.35**. Observe that we could trim sequences of plain pairs to be even shorter than $k+1$. By Lemma 4.32, processing such a "over-trimmed" sequence naively might result in false-positive reports of k-error occurrences, though. In essence, the remainder of this part shows when and how we are able to avoid such false-positive reports when trimming with $\alpha = \widetilde{\Theta}(\sqrt{d})$. ⌟

Observe that trimming the DPM-sequences alone is not enough to obtain Lemma 4.31: in the interesting case where $\kappa^3 = o(m/\tau)$, we cannot afford to consider each of the $\mathcal{O}(m/\tau)$ DPM-sequences $\mathcal{I}'_j := \mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), k+1)$ separately. However, for most values $j \in J$, we have $\mathcal{I}'_{j-1} = \mathcal{I}'_j$. Further, if $r_j - r_{j-1} = \tau$, then we have $r_j + \mathrm{Occ}_k^E(P, R_j) = r_{j-1} + \mathrm{Occ}_k^E(P, R_{j-1}) + \tau$. Consequently, we can return all occurrences compactly as arithmetic progressions. (Also recall from Lemma 3.12 that the sequence $(r_j)_{j \in J}$ is the concatenation of $\mathcal{O}(d_T + 1)$ arithmetic progressions with difference τ.) While these observations alone *still* are not enough to obtain a faster algorithm for NEWPERIODICMATCHES, they do yield an alternative algorithm with (roughly) the same running time as Fact 2.12.

With Remark 4.35 in mind, we give a slightly more general algorithm that allows for an arbitrary threshold α up to which we may trim the runs of plain pairs in the sequences $\mathcal{I}_j$. In particular, we prove the following more general variant of Lemma 4.31.

⚑ **Lemma 4.36** ($\mathtt{TrimmedPM}(T, P, k, Q, \mathcal{A}_P, \mathcal{A}_T, \mathsf{Red}(P), \mathsf{Red}(T), \alpha)$). *Suppose we are given an instance NEWPERIODICMATCHES$(P, T, k, d, Q, \mathcal{A}_P, \mathcal{A}_T)$, a positive integer α, as well as sets of puzzle pieces $\mathsf{Red}(P) \supseteq \mathsf{Special}(P)$ and $\mathsf{Red}(T) \supseteq \mathsf{Special}(T)$ of size $\mathcal{O}(d)$ each.*

After a preprocessing step that takes $\mathcal{O}(d^2 \alpha \log n)$ time in the $\mathtt{PILLAR}$ model, we can compute

$$\mathcal{G} := \bigcup_{j \in J}(r_j + \mathrm{Occ}_k^E(\mathit{Trim}(\mathcal{I}_j, \mathit{Red}(P), \mathit{Red}(T), \alpha)))$$

as $\mathcal{O}(d^3 \alpha)$ arithmetic progressions with difference $\tau = \mathcal{O}(\max\{q, d\})$, using a DYNAMICPUZZLEMATCHING $(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ data structure, with $\Delta + \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi) = \mathcal{O}(d)$, that maintains a DPM-sequence of length $z = \mathcal{O}(d\alpha)$ under

- *$\mathcal{O}(d^2 \alpha)$ calls to $\mathtt{DPM-Delete}$ and $\mathtt{DPM-Insert}$ that delete and insert, respectively, a canonical pair,*
- *$\mathcal{O}(d^2)$ calls to $\mathtt{DPM-Substitute}$, and*
- *$\mathcal{O}(d^2 \alpha)$ calls to $\mathtt{DPM-Query}$.*

◾

Proof. As in Sections 4.2 and 4.3, we set $\mathcal{S}_\beta$, $\mathcal{S}_\mu$, $\mathcal{S}_\varphi$ according to Definition 4.8 and choose $\Delta := 6\kappa = 6(k + d_P + d_T)$. Further, for each $j \in J$, set $\mathcal{I}'_j := \mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), \alpha)$.

In contrast to the warm-up algorithms, we initialize the DYNAMICPUZZLEMATCHING data structure with $\mathcal{I} \leftarrow \mathcal{I}'_{\min J}$ (instead of $\mathcal{I}_{\min J}$). Now, in general we follow the approach of the algorithms from Sections 4.2 and 4.3: we transform $\mathcal{I}$ to be equal to $\mathcal{I}'_j$ and then call $\mathtt{DPM-Query}$ to obtain the contained k-error occurrences. However, we skip over an $\mathcal{I}'_j$ if it is equal to $\mathcal{I}'_{j+1}$ and instead extend the arithmetic progressions obtained from the last call to $\mathtt{DPM-Query}$ by τ.

▱ **Claim 4.37.** In $\mathcal{O}(\kappa^2 \alpha \log n)$ time in total, we can compute, for all $j \in J \setminus \{\min J\}$ with $\mathcal{I}'_{j-1} \neq \mathcal{I}'_j$, a sequence update_j of DYNAMICPUZZLEMATCHING update operations that transforms $\mathcal{I}'_{j-1}$ to $\mathcal{I}'_j$, such that if we perform these updates one by one, all intermediate sequences satisfy conditions (a) and (b) of DYNAMICPUZZLEMATCHING. Further, these sequences of updates are returned in increasing order with respect to j and contain in total

- $\mathcal{O}(d^2 \alpha)$ calls to $\mathtt{DPM-Delete}$ and $\mathtt{DPM-Insert}$ that delete and insert, respectively, a canonical pair and
- $\mathcal{O}(d^2)$ calls to $\mathtt{DPM-Substitute}$.

Proof. We store $\mathcal{I}$ as a doubly-linked list I. We compute sequences that consist of the following types of updates:

- list-sub(pointer, $(P_i, T_{j,i})$): given a handle pointer to an element of I, substitute $(P_i, T_{j,i})$ for this element.
- list-ins(pointer, $(P_i, T_{j,i})$): given a handle pointer to an element of I, insert $(P_i, T_{j,i})$ before this element.

- list-del(pointer): given a pointer pointer to an element of I, delete its successor in I.

While this interface is different from that of DYNAMICPUZZLEMATCHING, each update can be mapped to a DYNAMICPUZZLEMATCHING update as long as we are able to perform the following operation: given a handle to an element of I, return its rank, that is, the number of elements that precede it in I. This operation can be implemented in $\mathcal{O}(\log n)$ time by maintaining a balanced binary search tree over the elements of I at the cost of an $\mathcal{O}(\log n)$-time additive overhead for each update operation.

In $\mathcal{O}(\kappa \log \kappa)$ time, we store the elements of $\mathsf{Red}(P) \cup \{P_1, P_z\}$ in a doubly linked list L_P and we store the elements of $\mathsf{Red}(T)$ in a doubly linked list L_X, each sorted with respect to their starting positions. At all times, we store a bidirectional pointer between each element of each of L_P and L_T and the DPM-pair that contains it in I. (We do not explicitly mention when such pointers need to be updated.) For both P and T, we write $\mathsf{pointer}(P_i)$ (or $\mathsf{pointer}(T_i)$) for the pointer from the element of L_P (or L_T) corresponding to P_i (or T_i) to the element of I that contains P_i (or T_i).

We insert all updates in a global min-heap with keys in $(\min J \mathinner{\ldotp\ldotp} \max J] \times \mathbb{Z} \cup \{-\infty\}$. When we are done generating updates, we pop them from the heap, inserting an update with key (j, p) in the min-heap update_j with priority p. These priorities ensure that the conditions on the maintained sequence are satisfied at all times. We can easily convince ourselves that this process indeed returns sequences update_j in increasing order with respect to j.

Similar to the algorithm that was sketched in Section 4.3, we consider a few types of updates to $\mathcal{I}$:

- $\mathcal{O}(\kappa)$ substitutions of the head/tail of the sequence,
- $\mathcal{O}(\kappa^2)$ updates that involve a pair that contains a red fragment, and
- $\mathcal{O}(\kappa^2 \alpha)$ insertions/deletions of a plain pair, effectively shrinking/expanding a run of plain pairs.

In order to develop some more intuition, let us examine how the length of the run of plain pairs that succeeds the head of $\mathcal{I}$ changes in the course of the algorithm. This length may decrease as a DPM pair A containing some $T_t \in \mathsf{Red}(T)$ approaches the head of $\mathcal{I}$ while we slide P on T. When such a DPM-pair disappears in the process of transforming $\mathcal{I}_{t-1}$ to $\mathcal{I}_t$, the length of the considered run might increase depending on where the leftmost pair with a red piece in $\mathcal{I}_t$ is; roughly speaking, the run of plain pairs after A in $\mathcal{I}_{t-1}$ becomes the run of plain pairs that succeeds the head of $\mathcal{I}_t$.

We say that a run of plain pairs is *enclosed* by the two DPM-pairs that are adjacent to it. Suppose that in each of $\mathcal{I}_{j-1}$ and $\mathcal{I}_j$ there is a run of plain pairs that is enclosed by a pair that contains a fragment $P_i \in \mathsf{Red}(P) \cup \{P_1, P_z\}$ and a pair that contains a fragment $T_t \in \mathsf{Red}(T)$. The lengths of these two runs differ by at most one; details are provided below. Observe that any run of plain pairs in $\mathcal{I}_{j-1}$ that is not enclosed by a pair that contains a fragment $P_i \in \mathsf{Red}(P) \cup \{P_1, P_z\}$ and a pair that contains a fragment $T_t \in \mathsf{Red}(T)$ either remains intact in $\mathcal{I}_j$ or is shifted to the left by one position. In particular, each change to the length of a run as we transform $\mathcal{I}'_{j-1}$ to $\mathcal{I}'_j$ can be attributed to a single pair of fragments $P_i \in \mathsf{Red}(P) \cup \{P_1, P_z\}$ and $T_t \in \mathsf{Red}(T)$ getting closer to (or farther from) each other.

We first issue updates for the head and the tail of the sequence.[16] To this end, we first compute the union $\mathsf{HeadTail}$ of the sets

$$\{j \in J : T_{j,1} \neq Q^\infty[-\kappa \mathinner{\ldotp\ldotp} 2\tau + \Delta\,)\} \quad \text{and} \quad \{j \in J : T_{j,z} \neq Q^\infty[z\tau \mathinner{\ldotp\ldotp} y_P + \kappa\,)\},$$

with its elements sorted in increasing order, in $\mathcal{O}(\kappa)$ time using Lemmas 4.16 and 4.17. For each $j \in \mathsf{HeadTail} \setminus \min J$, for each $e \in \{1, z\}$ we issue an update $\mathsf{list\text{-}sub}(\mathsf{pointer}(P_e), (P_e, T_{j,e}))$ with key $(j, \bigl||T_{j,e}| - |P_e|\bigr| - \bigl||T_{j-1,e}| - |P_e|\bigr|)$. Further, for each $j \in \mathsf{HeadTail} \setminus \max J$, for each $e \in \{1, z\}$, issue update $\mathsf{list\text{-}sub}(\mathsf{pointer}(P_e), (P_e, T_{j+1,e}))$ with key $(j + 1, \bigl||T_{j+1,e}| - |P_e|\bigr| - \bigl||T_{j,e}| - |P_e|\bigr|)$.

---

[16] Some of the issued updates might be redundant or duplicate, but this is not a problem.

Let us now issue all updates that involve some red fragment. For each pair of fragments $P_i \in \mathsf{Red}(P)$ and $T_t \in \mathsf{Red}(T)$ we do the following.

- If $t - i \in (\min J \mathinner{\ldotp\ldotp} \max J]$, we issue the following updates:
  - list-sub$(\mathsf{pointer}(P_i), (P_i, T_t))$ with key $(t - i, ||T_t| - |P_i|| - ||T_{t-1}| - |P_i||)$, and
  - if $i + 1 < z$ and $P_{i+1} \notin \mathsf{Red}(P)$, list-del$(\mathsf{pointer}(P_i))$ with key $(t - i, -\infty)$.
- If $t - i + 1 \in (\min J \mathinner{\ldotp\ldotp} \max J]$, we issue the following updates:
  - if $i > 2$ and $P_{i-1} \notin \mathsf{Red}(P)$, list-ins$(\mathsf{pointer}(P_i), (P_{i-1}, T_t))$ with key $(t - i + 1, ||T_t| - |P_{i-1}||)$, and
  - if $T_{t+1} \notin \mathsf{Red}(T)$, list-sub$(\mathsf{pointer}(P_i), (P_i, T_{t+1}))$ with key $(t - i + 1, ||T_{t+1}| - |P_i|| - ||T_t| - |P_i||)$.

Next, we show how to compute updates that allow the algorithm to maintain the trimmed length of the run of plain pairs that is enclosed by a pair containing some $P_i \in \mathsf{Red}(P) \cup \{P_1, P_z\}$ and a pair containing some $T_t \in \mathsf{Red}(T)$, when such a run exists.

First, we consider the case where the pair containing P_i precedes the pair containing $T_t = T_{j-1, t-(j-1)}$ in some $\mathcal{I}'_{j-1}$ and these two pairs enclose a non-empty run of plain pairs. In this case, the run would either shrink in $\mathcal{I}'_j$ or retain its length α. Let the successor of P_i in L_P be P_v. Further, if T_t is not the first element in L_T let its predecessor in L_T be $T_u = T_{j-1, u-(j-1)}$; otherwise, set $u := -\infty$. The following conditions must be satisfied:

(i) $t - (j - 1) \leq v$, which is equivalent to $j > t - v$, so that P_v is not in a pair strictly between the two pairs and hence none of the elements of $\mathsf{Red}(P)$ is,

(ii) $u - (j - 1) \leq i$, which is equivalent to $j > u - i$, so that there is no non-plain pair containing some element of $\mathsf{Red}(T)$ between the two pairs, and

(iii) $t - (j - 1) - i > 1$, which is equivalent to $j \leq t - i - 1$, so that the precedence condition is satisfied and the two pairs are not adjacent.

Further, the run should shrink in $\mathcal{I}'_j$ only if the two pairs are close enough, that is, if $t - (j - 1) - i - 1 \leq \alpha$, which is equivalent to $j > t - i - \alpha - 1$. Hence, for each

$$j \in (\max\{\min J, t - v, u - i, t - i - \alpha - 1\} \mathinner{\ldotp\ldotp} \min\{\max J, t - i - 1\}],$$

we issue an update list-del$(\mathsf{pointer}(P_i))$ with key $(j, -\infty)$. Observe that the total number of such issued updates is at most $(t - i - 1) - (t - i - \alpha - 1) = \alpha$.

We now consider the complementary case where in each of $\mathcal{I}'_{j-1}$ and $\mathcal{I}'_j$, the pair containing P_i succeeds the pair containing T_t and these two pairs either enclose a non-empty run of plain pairs or are adjacent.[17] In this case, the run would either be expanded by one plain pair in $\mathcal{I}'_j$ or retain its length α. Let the predecessor of P_i in L_P be P_y. Further, if T_t is not the last element in L_T let its successor in L_T be $T_w = T_{j-1, w-(j-1)}$; otherwise, set $w := \infty$. The following conditions must be satisfied:

(i) $t - (j - 1) > y$, which is equivalent to $j \leq t - y$, so that P_w is strictly to the left of both pairs in $\mathcal{I}'_{j-1}$, as if this is not the case then the condition on $\mathcal{I}'_j$ would not be satisfied,

(ii) $w - (j - 1) > i$, which is equivalent to $j \leq w - i$, so that T_y (if it exists) is strictly to the right of both pairs in $\mathcal{I}'_{j-1}$, as if this is not the case then the condition on $\mathcal{I}'_j$ would not be satisfied,

(iii) $t - (j - 1) < i$, which is equivalent to $j > t - i + 1$, so that the precedence condition is satisfied.

Further, the run should be expanded in $\mathcal{I}'_j$ only if the two pairs are close enough, that is, if $i - (t - (j - 1)) - 1 < \alpha$, which is equivalent to $j \leq \alpha + t - i + 1$. Hence, for each

$$j \in (\max\{\min J, t - i + 1\} \mathinner{\ldotp\ldotp} \min\{\max J, t - y, w - i, \alpha + t - i + 1\}],$$

---

[17] Observe that the pairs can only be adjacent in $\mathcal{I}'_{j-1}$.

we issue an update list-ins(pointer(P_i), ($Q^\infty[0\mathinner{.\,.}\tau+\Delta)$, $Q^\infty[0\mathinner{.\,.}\tau+\Delta)$)) plain pair in the (potentially empty) run preceding the pair that contains P_i with key $(j,-\infty)$. Observe that the total number of such issued updates is at most $(\alpha+t-i+1)-(t-i+1)=\alpha$.

Clearly, we can compute all updates for a given pair (in the intermediate interface) in $\mathcal{O}(\alpha)$ time.

The conditions of DYNAMICPUZZLEMATCHING are clearly satisfied at initialization. Condition (a) is satisfied at all times as all the updates involving the head or tail of only involve pairs in $\mathcal{S}_\beta\times\mathcal{S}_\beta$ and $\mathcal{S}_\varphi\times\mathcal{S}_\varphi$, respectively, while all remaining updates involve pairs in $\mathcal{S}_\mu\times\mathcal{S}_\mu$. Condition (b) is satisfied at all times; it is satisfied for each constructed $\mathcal{I}'_j$ by Lemma 4.10 as $\mathcal{I}'_j$ is a subsequence of $\mathcal{I}_j$, and for all other constructed sequences by a direct application of Fact 4.22, which is applicable because of how the updates in each update_j are sorted. ⌟

We are now ready to complete the reduction to DYNAMICPUZZLEMATCHING; see Algorithm 2 for a pseudocode implementation.

Using Claim 4.37, we initialize the set

$$\mathcal{U}=\{\mathsf{update}_j:j\in(\min J\mathinner{.\,.}\max J]\text{ and }\mathcal{I}'_{j-1}\neq\mathcal{I}'_j\}.$$

To simplify the construction of certain arithmetic progressions later in the proof, we insert to $\mathcal{U}$ the following single-element (void) sequences of updates. For each $j\in(\min J\mathinner{.\,.}\max J]$ such that $r_j-r_{j-1}\neq\tau$ and $\mathcal{I}'_{j-1}=\mathcal{I}'_j$, we insert to $\mathcal{U}$ the sequence update_j that consists of a single element: $\mathtt{DPM\text{-}Substitute}((P_1,T_{j,1}),1)$, that is, the substitution of the first pair of the maintained sequence with $(P_1,T_{j,1})$. By Lemma 3.12, we only generate $\mathcal{O}(d_T+1)$ new updates.

Now, we initialize the sequence $\mathcal{I}$ with $\mathcal{I}'_{\min J}$. We consider the sequences of updates update_j in $\mathcal{U}$ in increasing order with respect to j and apply them to the maintained sequence $\mathcal{I}$. Prior to performing the sequence of updates update_j we do the following. Suppose that the maintained sequence $\mathcal{I}$ corresponds to $\mathcal{I}'_t$, that is, either update_j is the first element of $\mathcal{U}$ and $t=\min J$ or the previously applied sequence of updates was update_t. First, we compute $\mathrm{Occ}_k^E(\mathcal{I}'_t)$ using a $\mathtt{DPM\text{-}Query}$. Observe that we have $\mathcal{I}'_t=\cdots=\mathcal{I}'_{j-1}$ and hence $\mathrm{Occ}_k^E(\mathcal{I}'_t)=\cdots=\mathrm{Occ}_k^E(\mathcal{I}'_{j-1})$. Further, for each $i\in(t\mathinner{.\,.}j)$, we have $r_i-r_{i-1}=\tau$ and hence $r_i+\mathrm{Occ}_k^E(\mathcal{I}'_i)=r_t+\mathrm{Occ}_k^E(\mathcal{I}'_t)+(i-t)\tau$. We can thus efficiently return $\bigcup_{i=t}^{j-1}(r_i+\mathrm{Occ}_k^E(\mathcal{I}'_t))$ as the union of $\mathcal{O}(|\mathrm{Occ}_k^E(\mathcal{I}'_i)|)$ arithmetic progressions

$$\bigcup_{a\in\mathrm{Occ}_k^E(\mathcal{I}'_t)}\{r_t+a+i\cdot\tau:i\in[0\mathinner{.\,.}j-t)\}.$$

(Observe that in the case where $t=j-1$, each of the constructed arithmetic progressions consists of a single element.) Finally, we apply the sequence of updates update_j. The case where there are no more updates to be processed is treated analogously (with j set to $\max J+1$). The upper bound on the number of returned arithmetic progressions is embedded in the analysis of the time complexity of the algorithm.

We now proceed to analyze the time required by Algorithm 2.

First, recall that the elements of $\mathcal{U}$ as returned by Claim 4.37 in $\mathcal{O}(\kappa^2\alpha\log n)$ time are sorted in increasing order with respect to j and a representation of $(r_j)_{j\in J}$ as the concatenation of $\mathcal{O}(d_T+1)$ arithmetic progressions with difference τ can be computed in $\mathcal{O}(d_T+1)$ time in the $\mathtt{PILLAR}$ model due to Lemma 3.12. Thus, Lines 4–6 can be implemented in $\mathcal{O}(\kappa^2\alpha)$ time in a parallel left-to-right scan of $\mathcal{U}$ and the aforementioned representation of $(r_j)_{j\in J}$, maintaining the sortedness of $\mathcal{U}$.

Next, observe that $\mathcal{I}'_{\min J}$ is of length $\mathcal{O}(\kappa\alpha)$ as it consists of a head, a tail, $\mathcal{O}(\kappa)$ non-plain pairs and $\mathcal{O}(\kappa)$ runs of plain pairs, each of length at most α. This sequence can be computed in $\mathcal{O}(\kappa\alpha+\kappa\log\kappa)$ time: the head and the tail are computed in $\mathcal{O}(\log(\kappa+1))$ time due to Corollary 4.15 and Lemmas 4.16 and 4.17, while the pairs consisting of internal pieces can be computed by processing the elements of

Algorithm 2 Reduction of computing $\mathcal{G}$ to an instance of DYNAMICPUZZLEMATCHING.

---

1 TrimmedPM(T, P, k, Q, $\mathcal{A}_P$, $\mathcal{A}_T$, $Red(P)$, $Red(T)$, α)

2 $\mathcal{G} \leftarrow \varnothing$;

3   Construct set $\mathcal{U}$ of sequences update_j specified in Claim 4.37;   // $\mathcal{O}(\kappa^2\alpha\log n)$ time

4   **foreach** $j \in (\min J \mathbin{..} \max J]$ *such that* $r_j - r_{j-1} \neq \tau$ *and* $\mathcal{I}'_{j-1} = \mathcal{I}'_j$ **do**

5 $\mathsf{update}_j \leftarrow \{\text{DPM-Substitute}((P_1, T_{j,1}),1)\}$;

6 $\mathcal{U} \leftarrow \mathcal{U} \cup \{\mathsf{update}_j\}$;

7 $\mathcal{I} \leftarrow \mathcal{I}'_{\min J}$;   // $\mathcal{O}(\kappa\alpha + \kappa\log\kappa)$ time

8 $t \leftarrow \min J$;

9   **foreach** $j \in \{i \in (\min J \mathbin{..} \max J + 1] : \mathsf{update}_i \neq \varnothing \text{ or } i = \max J + 1\}$ *in incr. order* **do**

10 $\mathrm{Occ}_k^E(\mathcal{I}'_j) \leftarrow \text{DPM-Query}$;

11   **foreach** $a \in \mathrm{Occ}_k^E(\mathcal{I}'_j)$ **do**

12 $\mathcal{G} \leftarrow \mathcal{G} \cup \{r_t + a + i \cdot \tau : i \in [0 \mathbin{..} j - t)\}$;

13   **if** $j \leq \max J$ **then**

14   Apply the sequence of updates update_j on $\mathcal{I}$;

15 $t \leftarrow j$;

16   **return** $\mathcal{G}$;

$Red(P)$, $Red(T)$ in a left-to-right manner, after sorting them with respect to their starting positions in $\mathcal{O}(\kappa\log\kappa)$ time.

The for-loop of Lines 9–15 is executed $\mathcal{O}(\kappa^2\alpha)$ times as this is an upper bound on the number of sequences of updates that are processed. For each $j \in J$, due to Lemma 4.10, we have $\sum_{i=1}^{z_j} \big||F_{j,i}|-|G_{j,i}|\big| \leq \sum_{i=1}^{z} \big||T_{j,i}|-|P_i|\big| \leq 3\kappa - k$ and hence $|\mathrm{Occ}_k^E(F_j, G_j)| = \mathcal{O}(\kappa)$. Thus, each query in Line 10 returns a set of size $\mathcal{O}(\kappa)$. Hence, apart from the time required for DYNAMICPUZZLEMATCHING updates and queries, each iteration of the for-loop takes $\mathcal{O}(\kappa)$ time, for a total of $\mathcal{O}(\kappa^3\alpha)$ time. It readily follows that the output of the algorithm consists in $\mathcal{O}(\kappa^3\alpha)$ arithmetic progressions.

All in all, in $\mathcal{O}(\kappa^3\alpha + \kappa^2\alpha\log n) = \mathcal{O}(d^3\alpha + d^2\alpha\log n)$ time, the problem in scope reduces to an instance of DYNAMICPUZZLEMATCHING with $\mathcal{S}_\beta + \mathcal{S}_\mu + \mathcal{S}_\varphi = \mathcal{O}(d)$ (due to Lemma 1.6), $\Delta = 6\kappa = \mathcal{O}(d)$, $\mathcal{I}$ initialized as a sequence of length $\mathcal{O}(d\alpha)$, and $\mathcal{O}(d^2\alpha)$ updates and queries. ◾

As a direct consequence, we obtain an alternative algorithm for the NEWPERIODICMATCHES problem with a competitive running time, that is, we obtain an algorithm for NEWPERIODICMATCHES, which is slower than the algorithm implied by Fact 2.12 only by a poly-logarithmic factor.

⬛ **Corollary 4.38**. *We can solve the NEWPERIODICMATCHES problem in $\mathcal{O}(d^4\log n\log d)$ time in the* PILLAR *model.*

Proof. By Lemma 4.30, in order to solve an instance of NEWPERIODICMATCHES, it suffices to construct sets $\mathsf{Special}(P)$ and $\mathsf{Special}(T)$ in $\mathcal{O}(d)$ time in the PILLAR model using Lemma 1.6 and call the algorithm underlying Lemma 4.36 with $Red(X) \coloneqq \mathsf{Special}(X)$ for each $X \in \{P, T\}$ and $\alpha \coloneqq k + 1$.

Using Theorem 1.5 for the DYNAMICPUZZLEMATCHING problem, we need $\mathcal{O}(d^3\log^2 d)$ time for preprocessing, $\mathcal{O}(d^4\log d)$ time for initialization, and $\mathcal{O}(d^3 \cdot d\log n\log d)$ time for processing updates and queries. Overall, the required time is thus $\mathcal{O}(d^4\log n\log d)$. ◾

In the following sections, our goal is to reduce the NEWPERIODICMATCHES problem to an instance TrimmedPM($T, P, k, Q, \mathcal{A}_T, \mathcal{A}_P, Red(P), Red(T), O(\sqrt{d})$) with $|Red(P)| + |Red(T)| = \mathcal{O}(d)$.

5 Faster NewPeriodicMatches: Additional Combinatorial Insights

While the improved algorithms in [CKW20] crucially relied on analyzing and understanding the structure of the *pattern*, we obtain the improvements in this work by additionally analyzing and understanding the structure of the *text*. Throughout this section, we fix a text T, a pattern P, a primitive string Q, and parameters d_P, d_T, and d stemming from an instance of NewPeriodicMatches. Our goal is to classify each position of the text where a k-error occurrence of P can start (cf. Fact 2.8 in this regard) as either *heavy* or *light*, using the notion of *locked fragments*, depending on some parameter η. The heavy positions are few and can be covered by a few *heavy ranges* of small total length. The set of light positions may be large, but, as explained and exploited in Section 6, for each light position $v \in \mathrm{Occ}_k^E(P, T)$, any alignment $\mathcal{A} : P \leadsto T[v \mathinner{\ldotp\ldotp} w)$ of cost $\min_t \delta_E(P, T[v \mathinner{\ldotp\ldotp} t))$ is quite restricted: it does not makes edit operations outside the vicinity (interpreted as $\mathcal{O}(\eta\tau)$ positions) of locked fragments of the text and the pattern.

5.1 Locked Fragments and their Properties

We intend to understand and analyze the structure of the text by using a more elaborate version of the marking schemes used in [CKW20]. In particular, we heavily rely on the notion of *locked fragments* in text and pattern from [CKW20].

Definition 5.1 ([CKW20, 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 $L = S$.*

Let us also recall an intuitive example from [CKW20].

Example 5.2 ([CKW20]). Fix a primitive string Q and consider a string $U = Q^{k+1}SQ^{k+1}$ with $\delta_E(U, {}^*Q^*) \leq k$. In any optimal alignment of U with a substring of Q^∞ with at most k edits, at least one of the leading $k + 1$ occurrences of Q in U is matched exactly and at least one of the trailing $k + 1$ occurrences of Q in U is matched exactly. Hence, all occurrences preceding (or succeeding) said exactly matched occurrence of Q are also matched exactly. Thus, U is locked with respect to Q.

As in [CKW20], we also need the slightly stronger notion of an *h-locked prefix* of a string.

Definition 5.3 ([CKW20, Definition 5.10]). *Let S denote a string, let Q denote a primitive string, and let $h \geq 0$ denote an integer. We say that a prefix L of S is* h-locked *(with respect to Q) if at least one of the following holds:*
- *For every $p \in [0 \mathinner{\ldotp\ldotp} |Q|)$, if $\delta_E(L, \mathrm{rot}^p(Q)^*) \leq h$, then $\delta_E(L, \mathrm{rot}^p(Q)^*) = \delta_E(L, Q^\infty[-p \mathinner{\ldotp\ldotp} j|Q|))$ for some $j \in \mathbb{Z}$.*
- *We have $L = S$.*

Given a string S (that is either the pattern or the text), we intend to construct locked fragments covering all errors of S with respect to ${}^*Q^*$ such that the total length of these locked fragments is roughly proportional to the product of $|Q|$ and $\delta_E(S, {}^*Q^*)$.

Lemma 5.4 (Locked(S, Q, d, h), [CKW20, Lemma 6.9]). *Let S denote a string, let Q denote a primitive string, let d_S denote a positive integer such that $\delta_E(S, {}^*Q^*) \leq d_S$ and $|S| \geq (2d_S + 1)|Q|$, and let $h \in \mathbb{Z}_{\geq 0}$.*
   Then, there is an algorithm that computes disjoint locked fragments $L_1, \ldots, L_\ell \preceq S$ such that

- $S = L_1 \cdot \bigodot_{i=1}^{\ell-1}(Q^{\alpha_i} L_{i+1})$ *for positive integers* $\alpha_1, \ldots, \alpha_{\ell-1}$;[18]

- L_1 *is an* h-*locked prefix of* S *and* L_ℓ *is a suffix of* S;

- $\delta_E(S, {}^*Q^*) = \sum_{i=1}^{\ell} \delta_E(L_i, {}^*Q^*)$ *and* $\delta_E(L_i, {}^*Q^*) > 0$ *for* $i \in (1 \mathinner{.\,.} \ell)$; *and*

- $\displaystyle\sum_{i=1}^{\ell} |L_i| \le (5|Q|+1)\delta_E(S, {}^*Q^*) + 2(h+1)|Q|$.

The algorithm takes $\mathcal{O}(d_S^2 + h)$ *time in the* `PILLAR` *model.* ◪

5.2 Analyzing the Text Using Locked Fragments

Our marking scheme is similar to the marking scheme used in the proof of [CKW20, Theorem 5.2]. As both T and P are close to being periodic, we can compute locked fragments with respect to a common string Q (which is given as a parameter in the call of NewPeriodicMatches).

We start with some intuition for how different locked fragments of T and P influence the edit distance between T and P. To that end, let $\mathcal{L}^P$ denote the set of locked fragments computed for P, let $\mathcal{L}^T$ denote the set of locked fragments computed for T, and consider an alignment $\mathcal{A} : P \rightsquigarrow T[t \mathinner{.\,.} t')$. If $\mathcal{A}$ aligns a locked fragment $L^P \in \mathcal{L}^P$ to a fragment of Q^∞, that is, $\mathcal{A}(L^P)$ does not overlap any locked fragment of $T[t \mathinner{.\,.} t')$, we obtain $\delta_E^{\mathcal{A}}(L^P, \mathcal{A}(L^P)) \ge \delta_E(L^P, {}^*Q^*)$. Symmetrically, any locked fragment $L^T \in \mathcal{L}^T$ of $T[t \mathinner{.\,.} t')$, for which $\mathcal{A}^{-1}(L^T)$ does not overlap any locked fragment of P satisfies $\delta_E^{\mathcal{A}}(\mathcal{A}^{-1}(L^T), L^T) \ge \delta_E(L^T, {}^*Q^*)$.

However, if $\mathcal{A}$ aligns a substring U of a locked fragment $L^P \in \mathcal{L}^P$ to a substring V of a locked fragment $L^T \in \mathcal{L}^T$, the situation is not as straightforward. (One can consider the cleaner case where U is L^P and V is L^T.) Suppose for simplicity that there exist integers x and y such that $\delta_E(U, {}^*Q^*) = \delta_E(U, Q^\infty[x \mathinner{.\,.} y))$ and $\delta_E(V, {}^*Q^*) = \delta_E(V, Q^\infty[x \mathinner{.\,.} y))$. In this case, by the triangle inequality, we have that

$$\delta_E(U, {}^*Q^*) + \delta_E(V, {}^*Q^*) - \min\{\delta_E(U, {}^*Q^*), \delta_E(V, {}^*Q^*)\} \le \delta_E(U, V) \le \delta_E(U, {}^*Q^*) + \delta_E(V, {}^*Q^*).$$

In other words, by aligning these substrings of locked fragments, we can hope to "save" at most $\min\{\delta_E(U, {}^*Q^*), \delta_E(V, {}^*Q^*)\}$ edits compared to the $\delta_E(U, {}^*Q^*) + \delta_E(V, {}^*Q^*)$ upper bound (when aligning the same locked fragments to substrings of Q^∞ instead).

To quantify said potential "savings", we give marks to each position t in the text corresponding to the total number of marks potentially saved in alignment $\mathcal{A} : P \rightsquigarrow T[t \mathinner{.\,.} t')$. Roughly speaking, for each pair of fragments L^T and L^P, we place $\min\{\delta_E(L^T, {}^*Q^*), \delta_E(L^P, {}^*Q^*)\}$ marks at position t, if they may overlap in any alignment $\mathcal{A} : P \rightsquigarrow T[t \mathinner{.\,.} t')$ with at most $\hat{\kappa}$ insertions and deletions; $\hat{\kappa}$ can be thought to be $\mathcal{O}(d)$. In what follows, when we check whether two fragments overlap, we thus allow for a (small) slack $\hat{\kappa}$. Formally, the marking scheme is captured in Definition 5.5.

⬚ **Definition 5.5.** *For a text* T, *a pattern* P, *a primitive string* Q *with* $\delta_E(P, {}^*Q^*) = d_P$ *and* $\delta_E(T, {}^*Q^*) = d_T$, *and corresponding sets of locked fragments* $\mathcal{L}^P := \mathtt{Locked}(P, Q, d_P, \partial_P)$ *and* $\mathcal{L}^T := \mathtt{Locked}(T, Q, d_T, 0)$, *write* mk *for the function that maps an integer* v *to a (weighted) number of locked fragments in* $\mathcal{L}^P$ *that (almost) overlap locked fragments in* $\mathcal{L}^T$ *when aligning* P *to position* v.

---

[18] This item is not stated in [CKW20, Lemma 6.9]. However, it readily follows from the construction algorithm underlying that lemma and it is already used in [CKW20] (for instance, in the proof of [CKW20, Claim 5.17]). We believe that adding a proof of this item here would not be instructive.

Formally, we first define a function $\mathrm{mk} : \mathbb{Z} \times \mathbb{Z}_{\geq 0} \times \mathcal{L}^P \times \mathcal{L}^T \to \mathbb{Z}_{\geq 0}$ *by*

$$\mathrm{mk}(v, \hat{\kappa}, L^P = P[\ell_P \mathbin{..} r_P), L^T = T[\ell_T \mathbin{..} r_T))$$

$$:= \begin{cases} \delta_E(L^T, {}^*Q^*), & \text{if } \ell_P = 0 \text{ and } v \in (\ell_T - \hat{\kappa} - r_P \mathbin{..} r_T + \hat{\kappa} - \ell_P); \\ \min\{\delta_E(L^T, {}^*Q^*), \delta_E(L^P, {}^*Q^*)\}, & \text{if } \ell_P \neq 0 \text{ and } v \in (\ell_T - \hat{\kappa} - r_P \mathbin{..} r_T + \hat{\kappa} - \ell_P); \\ 0, & \text{otherwise.} \end{cases}$$

Now, set

$$\mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T) := \sum_{L^P \in \mathcal{L}^P} \sum_{L^T \in \mathcal{L}^T} \mathrm{mk}(v, \hat{\kappa}, L^P, L^T).$$

When $\hat{\kappa}$, $\mathcal{L}^P$, and $\mathcal{L}^T$ are clear from context we may just write $\mathrm{mk}(v)$ for $\mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T)$.

We continue with a set of useful observations about our marking scheme. Fix sets of locked fragments $\mathcal{L}^P := \mathtt{Locked}(P, Q, d_P, \partial_P)$ and $\mathcal{L}^T := \mathtt{Locked}(T, Q, d_T, 0)$, with $\partial_P = \mathcal{O}(d)$.

Lemma 5.6. *For every* $\hat{\kappa} \in \mathbb{Z}_{\geq 0}$, *we have*

$$\sum_{v \in \mathbb{Z}} \mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T) \leq 2d_T(d_P + 2)\hat{\kappa} + 2d_T(6d_P + 3d_T + \partial_P + 2)|Q| + 2d_P|Q| = \mathcal{O}(d^2(\hat{\kappa} + |Q|)).$$

Proof. Fix a locked fragment $L^P = P[\ell_P \mathbin{..} r_P) \in \mathcal{L}^P$ and a locked fragment $L^T = P[\ell_T \mathbin{..} r_T) \in \mathcal{L}^T$. If $\ell_P \neq 0$, then the definition of mk yields

$$\sum_{v \in \mathbb{Z}} \mathrm{mk}(v, \hat{\kappa}, L^P, L^T) \leq \sum_{v \in (\ell_T - \hat{\kappa} - r_P \mathbin{..} r_T + \hat{\kappa} - \ell_P)} \min\{\delta_E(L^T, {}^*Q^*), \delta_E(L^P, {}^*Q^*)\}$$

$$\leq \min\{\delta_E(L^T, {}^*Q^*), \delta_E(L^P, {}^*Q^*)\}(|L^T| + |L^P| + 2\hat{\kappa})$$

$$\leq \delta_E(L^T, {}^*Q^*)(|L^P| + 2\hat{\kappa}) + \delta_E(L^P, {}^*Q^*)|L^T|.$$

Similarly, if $\ell_P = 0$, then the definition of mk yields

$$\sum_{v \in \mathbb{Z}} \mathrm{mk}(v, \hat{\kappa}, L^P, L^T) \leq \sum_{v \in (\ell_T - \hat{\kappa} - r_P \mathbin{..} r_T + \hat{\kappa} - \ell_P)} \delta_E(L^T, {}^*Q^*)$$

$$\leq \delta_E(L^T, {}^*Q^*)(|L^T| + |L^P| + 2\hat{\kappa})$$

$$\leq \delta_E(L^T, {}^*Q^*)(|L^P| + 2\hat{\kappa}) + d_T|L^T|.$$

Overall, we have

$$\sum_{a \in \mathbb{Z}} \mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T) \leq \sum_{L^P \in \mathcal{L}^P} \sum_{L^T \in \mathcal{L}^T} (\delta_E(L^T, {}^*Q^*)(|L^P| + 2\hat{\kappa}) + \delta_E(L^P, {}^*Q^*)|L^T|) + d_T \sum_{L^T \in \mathcal{L}^T} |L^T|$$

$$\leq d_T \sum_{L^P \in \mathcal{L}^P} (|L^P| + 2\hat{\kappa}) + (d_P + d_T) \sum_{L^T \in \mathcal{L}^T} |L^T|$$

$$\leq d_T(d_P + 2)2\hat{\kappa} + d_T \sum_{L^P \in \mathcal{L}^P} |L^P| + (d_P + d_T) \sum_{L^T \in \mathcal{L}^T} |L^T|$$

$$\leq 2d_T(d_P + 2)\hat{\kappa} + d_T(5|Q| + 1)d_P + 2d_T(\partial_P + 1)|Q| + (d_P + d_T)(5|Q| + 1)d_T$$
$$+ (d_P + d_T)2|Q|$$

$$\leq 2d_T(d_P + 2)\hat{\kappa} + 2d_T(6d_P + 3d_T + \partial_P + 2)|Q| + 2d_P|Q|.$$

Definition 5.7. *For a set* $\mathcal{U}$ *of fragments of a string* S *and an interval* $I \subseteq \mathbb{Z}$, *we write* $\mathcal{U}_I = \{S[x \mathbin{..} y) \in \mathcal{U} : [x \mathbin{..} y) \cap I \neq \varnothing\}$.

Recall that, as stated in Fact 2.8, $\mathrm{Occ}_k^E(P,T) \cap (n-m+k \mathinner{\ldotp\ldotp} n) = \varnothing$. We partition the remaining positions of the text into two groups. Intuitively, a position $t \in [0 \mathinner{\ldotp\ldotp} n-m+k]$ is *light* if it has a few marks, and there cannot be an alignment $\mathcal{A}: P \rightsquigarrow T[t \mathinner{\ldotp\ldotp} t')$ with at most $\hat{\kappa}$ insertions and deletions that aligns the first position of P, the last position of L_1^P, or the last position of P against a portion of a locked fragment of T.

⌐ **Definition 5.8**. *For any fixed thresholds $\hat{\kappa} \in \mathbb{Z}_{\geq 0}$ and $\eta \in \mathbb{Z}_{>0}$, we say that a position $v \in [0 \mathinner{\ldotp\ldotp} n-m+k]$ is* light *if the following conditions are simultaneously satisfied:*

- $\mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T) < \eta$,
- $\mathcal{L}_{[v-\hat{\kappa} \mathinner{\ldotp\ldotp} v+\hat{\kappa})}^T = \varnothing$,
- $\mathcal{L}_{[v+m-\hat{\kappa} \mathinner{\ldotp\ldotp} v+m+\hat{\kappa})}^T = \varnothing$, and
- $\mathcal{L}_{[v+|L_1^P|-\hat{\kappa} \mathinner{\ldotp\ldotp} v+|L_1^P|+\hat{\kappa})}^T = \varnothing$.

Otherwise, the position $v \in [0 \mathinner{\ldotp\ldotp} n-m+k]$ is called heavy. *We denote the sets of heavy and light positions by $\mathbb{H}$ and $\mathbb{L}$, respectively.* ⌐

⌐ **Lemma 5.9** ($\mathrm{Heavy}(P,T,k,d,Q,\mathcal{L}^P,\mathcal{L}^T,\hat{\kappa},\eta)$). *Consider an instance of the NEWPERIODICMATCHES problem, families $\mathcal{L}^P = \mathrm{Locked}(P,Q,d_P,\partial_P)$ and $\mathcal{L}^T = \mathrm{Locked}(T,Q,d_T,0)$, as well as thresholds $\kappa \in \mathbb{Z}_{\geq 0}$ and $\eta \in \mathbb{Z}_{>0}$,*

The set $\mathbb{H}$ of heavy positions, represented as the union of $\mathcal{O}(d^2)$ disjoint integer ranges, can be computed in $\mathcal{O}(d^2 \log \log d)$ time.

Proof. We implement the marking process according to Definition 5.5, assigning η extra marks to all positions made heavy due to the last three items in Definition 5.8. Formally, we produce the following $\mathcal{O}(d^2)$ weighted intervals:

- $(\ell_T - \hat{\kappa} - r_P \mathinner{\ldotp\ldotp} r_T + \hat{\kappa} - \ell_P)$ of weight $\min\{\delta_E(L^T, {}^*Q^*), \delta_E(L^P, {}^*Q^*)\}$ for each locked fragment $L^P = P[\ell_P \mathinner{\ldotp\ldotp} r_P) \in \mathcal{L}^P$ with $\ell_P \neq 0$ and $L_T = T[\ell_T \mathinner{\ldotp\ldotp} r_T) \in \mathcal{L}^T$,
- $(\ell_T - \hat{\kappa} - r_P \mathinner{\ldotp\ldotp} r_T + \hat{\kappa} - \ell_P)$ of weight $\delta_E(L^T, {}^*Q^*)$ for each locked fragment $L^P = P[\ell_P \mathinner{\ldotp\ldotp} r_P) \in \mathcal{L}^P$ with $\ell_P = 0$ and $L_T = T[\ell_T \mathinner{\ldotp\ldotp} r_T) \in \mathcal{L}^T$,
- $(\ell_T - \hat{\kappa} \mathinner{\ldotp\ldotp} r_T + \hat{\kappa})$ of weight η for each locked fragment $L_T = T[\ell_T \mathinner{\ldotp\ldotp} r_T) \in \mathcal{L}^T$,
- $(\ell_T - \hat{\kappa} - m \mathinner{\ldotp\ldotp} r_T + \hat{\kappa} - m)$ of weight η for each locked fragment $L_T = T[\ell_T \mathinner{\ldotp\ldotp} r_T) \in \mathcal{L}^T$,
- $(\ell_T - \hat{\kappa} - |L_1^P| \mathinner{\ldotp\ldotp} r_T + \hat{\kappa} - |L_1^P|)$ of weight η for each locked fragment $L_T = T[\ell_T \mathinner{\ldotp\ldotp} r_T) \in \mathcal{L}^T$.

The heavy positions are exactly those positions in $[0 \mathinner{\ldotp\ldotp} n-m+k]$ that are contained in intervals of total weight at least η; they can be computed using a sweep-line procedure with events corresponding to interval endpoints. The number of events is $\mathcal{O}(d^2)$, so the output consists of $\mathcal{O}(d^2)$ disjoint ranges. In terms of the running time, the bottleneck is sorting the events using the algorithm of [AHNR98]. ⌐

⌐ **Lemma 5.10**. *The total number of heavy positions does not exceed*

$$|\mathbb{H}| \leq \frac{2d_T(d_P + 2 + 3\eta)\hat{\kappa} + 2d_T(6d_P + 3d_T + \partial_P + 2 + 9\eta)|Q| + 2d_P|Q| + 6|Q|\eta + 2\hat{\kappa}\eta}{\eta}$$

$$= \mathcal{O}((d^2/\eta + d)(\hat{\kappa} + |Q|)).$$

Moreover, for any integer $b \in \mathbb{Z}_{>0}$,

$$\left| \bigcup_{v \in \mathbb{H}} [v - b \mathinner{\ldotp\ldotp} v + b] \right|$$

$$\leq \frac{2d_T(d_P + 2 + 3\eta)(\hat{\kappa} + b) + 2d_T(6d_P + 3d_T + \partial_P + 2 + 9\eta)|Q| + 2d_P|Q| + 6|Q|\eta + 2(\hat{\kappa} + 2b)\eta}{\eta}$$

$$= \mathcal{O}((d^2/\eta + d)(\hat{\kappa} + b + |Q|)).$$

Proof. By Lemma 5.6, the number of positions violating the first condition of Definition 5.8 does not exceed $\left(2d_T(d_P+2)\hat{\kappa}+2d_T(6d_P+3d_T+\partial_P+2)|Q|+2d_P|Q|\right)/\eta=\mathcal{O}(d^2(\hat{\kappa}+|Q|)/\eta)$. As for the remaining conditions, each locked fragment $L^T\in\mathcal{L}^T$ may yield at most $3(|L^T|+2\hat{\kappa})$ heavy positions, for a total of:

$$\sum_{L^T\in\mathcal{L}^T}3(|L^T|+2\hat{\kappa})\le 3(5|Q|+1)d_T+3\cdot 2|Q|+6(d_T+2)\hat{\kappa}\le 18d_T|Q|+6|Q|+6d_T\hat{\kappa}+2\hat{\kappa}=\mathcal{O}(d(\hat{\kappa}+|Q|)).$$

As for the second claim, observe that if v is heavy, then all positions in $[v-b\mathinner{\ldotp\ldotp}v+b]\cap[0\mathinner{\ldotp\ldotp}n-m+k]$ would be heavy if we increased the threshold $\hat{\kappa}$ to $\hat{\kappa}+b$. This is because the left endpoint of all intervals considered in the proof of Lemma 5.9 contains a $-\hat{\kappa}$ term whereas the right endpoint contains a $+\hat{\kappa}$ term (and there is no other dependency on $\hat{\kappa}$). Further, $\left|\bigcup_{v\in\mathbb{H}}[v-b\mathinner{\ldotp\ldotp}v+b]\setminus[0\mathinner{\ldotp\ldotp}n-m+k]\right|\le 2b$, since $\mathbb{H}\subseteq[0\mathinner{\ldotp\ldotp}n-m+k]$. ◾

6 A Faster Algorithm for NewPeriodicMatches

In this section, we present the core of our improvements: a faster algorithm for NEWPERIODICMATCHES.

⬛ **Lemma 1.3** (NewPeriodicMatches(P, T, k, d, Q, $\mathcal{A}_P$, $\mathcal{A}_T$)). *We can solve the NEWPERIODIC-MATCHES problem in $\mathcal{O}(d^{3.5}\sqrt{\log n\log d})$ time in the* PILLAR *model.* ◾

Let us fix an instance of the NEWPERIODICMATCHES problem, sets of locked fragments $\mathcal{L}^P\coloneqq\texttt{Locked}(P,Q,d_P,2d)$ and $\mathcal{L}^T\coloneqq\texttt{Locked}(T,Q,d_T,0)$, that is, we have $\partial_P\coloneqq 2d$, an integer $\hat{\kappa}\coloneqq 7d\ge\partial_P+2k+d_P+d_T>\kappa$ and an integer threshold $\eta\coloneqq\max\{1,\lfloor\sqrt{d}/\sqrt{\log(n+1)\log(d+1)}\rfloor\}\le d$. Further, consider a partition of the positions of T in $[0\mathinner{\ldotp\ldotp}n-m+k]$ into heavy and light using Lemma 5.9.

First, we show how to compute k-error occurrences that start at heavy positions; this is a straightforward application of Lemmas 3.12 and 3.13. Then, we reduce the problem of computing k-error occurrences that start at light positions to an instance $\texttt{TrimmedPM}(T,P,k,Q,\mathcal{A}_T,\mathcal{A}_P,\mathsf{Red}(P),\mathsf{Red}(T),\alpha))$, where $\mathsf{Red}(P)$ and $\mathsf{Red}(T)$ are both of size $\mathcal{O}(d)$ and contain all the internal pieces that (almost) overlap locked fragments, and $\alpha=o(\sqrt{d})$.

6.1 Computing Occurrences Starting at Heavy Positions

⬛ **Lemma 6.1** (HeavyMatches(P, T, k, d, Q, $\mathcal{A}_P$, $\mathcal{A}_T$, $\mathbb{H}$)). *Given an instance of the NEW-PERIODICMATCHES problem and the set $\mathbb{H}$ of heavy positions constructed using Lemma 5.9, $\mathrm{Occ}_k^E(P,T)\cap\mathbb{H}$ can be computed in $\mathcal{O}(d^3+d^4/\eta)$ time in the* PILLAR *model.*

Proof. Recall that the set $\mathbb{H}$ is represented as the union of $\mathcal{O}(d^2)$ disjoint heavy ranges (listed in the left-to-right order).

Our algorithm starts with an application Lemma 3.12 to construct the sequence $(r_j)_{j\in J}$, represented as a concatenation of $\mathcal{O}(d)$ arithmetic progressions. Our first goal is to enumerate elements of the set $J'\coloneqq\{j\in J:[r_j\mathinner{\ldotp\ldotp}r_{j+1})\cap\mathbb{H}\ne\varnothing\}$. For this, we simultaneously traverse the sequence $(r_j)_{j\in J}$ along with the heavy ranges constituting $\mathbb{H}$. For each heavy range $H\subseteq\mathbb{H}$, we list all $j\in J$ such that $[r_j\mathinner{\ldotp\ldotp}r_{j+1})\cap H\ne\varnothing$ (only the smallest such j might have already been listed for an earlier heavy range). In the second phase, we compute $O\coloneqq\bigcup_{j\in J'}(r_j+\mathrm{Occ}_k^E(P,R_j))\cap[r_j\mathinner{\ldotp\ldotp}r_{j+1})$ using Lemma 3.13, making sure that the positions are listed in the left-to-right order. Finally, we simultaneously scan O and the heavy ranges constituting $\mathbb{H}$, reporting all positions $v\in O\cap\mathbb{H}$.

As for correctness, observe that $R_j=T[r_j\mathinner{\ldotp\ldotp}r_j')$, so $O\subseteq\mathrm{Occ}_k^E(P,T)$ and, due to the final filtering step, we output a subset of $\mathrm{Occ}_k^E(P,T)\cap\mathbb{H}$. To prove the converse inclusion, consider a position $v\in\mathrm{Occ}_k^E(P,T)\cap\mathbb{H}$. By Corollary 3.11, there exists $j\in J$ such that $v\in(r_j+\mathrm{Occ}_k^E(P,R_j))\cap[r_j\mathinner{\ldotp\ldotp}r_{j+1})$ and, by the definition of J', we also have $j\in J'$. Consequently, v is indeed reported.

As for the complexity analysis, let us first compute the running time in terms of $|J'|$. Constructing the sequence $(r_j)_{j\in J}$ costs $\mathcal{O}(d)$ time. The set J' can be computed in $\mathcal{O}(|J'| + d + d^2) = \mathcal{O}(|J'| + d^2)$ time. The applications of Lemma 3.13 cost $\mathcal{O}(d^2)$ time each (in the PILLAR model), for a total of $\mathcal{O}(|J'| \cdot d^2)$ time in the PILLAR model. This is also a (crude) upper bound on the output size, so the final filtering step works in $\mathcal{O}((1 + |J'|)d^2)$ time. Overall, the running time in the PILLAR model is $\mathcal{O}((1 + |J'|)d^2)$.

It remains to bound $|J'|$. Observe that each interval $[r_j \mathbin{..} r_{j+1})$ (possibly except for the last one with $j = \max J$) has length at most $\tau + d_T$. Moreover, the total length of any s intervals $[r_j \mathbin{..} r_{j+1})$ with $0 < r_j < r_{j+1} < n$ is at least $s\tau - d_T$. On top of that, there might be one non-empty interval with $r_j = 0$ and one non-empty interval with $r_{j+1} = n$. We conclude that

$$|J'| \leq 2 + \frac{\sum_{j\in J'\setminus\{\max J\}} |[r_j \mathbin{..} r_{j+1})| + d_T}{\tau} \leq 2 + \frac{\left|\bigcup_{v\in\mathbb{H}}[v - \tau - d_T \mathbin{..} v + \tau + d_T]\right| + d_T}{\tau}.$$

Since $\tau = \Theta(\max(|Q|, d))$ and $d_T, d_P, k, \hat{\kappa} = \mathcal{O}(d)$, the bound of Lemma 5.10 yields that $|J'| = \mathcal{O}(d + d^2/\eta)$. Hence, the total running time is $\mathcal{O}(d^3 + d^4/\eta)$ just as claimed. ◾

6.2 Computing Occurrences Starting at Light Positions

6.2.1 Combinatorial Insights

The following fact follows directly from Definition 5.8.

Fact 6.2. *For any light position v in T and for any $w \in (v + m - \hat{\kappa} \mathbin{..} v + m + \hat{\kappa})$, we have $\mathcal{L}^T_{[v \mathbin{..} w)} = \{T[\ell \mathbin{..} r) \in \mathcal{L}^T : [\ell \mathbin{..} r) \subseteq [v \mathbin{..} v + m)\}$.* ◾

For each position v of T, set $\rho(v) \in [x_T \mathbin{..} y_T]$ such that $\mathcal{A}_T(T[v \mathbin{..} n)) = Q^\infty[\rho(v) \mathbin{..} y_T)$. Observe that

$$\delta_E(T, {}^*Q^*) = \delta_E(T[0 \mathbin{..} v), Q^\infty[x_T \mathbin{..} \rho(v))) + \delta_E(T[v \mathbin{..} n), Q^\infty[\rho(v) \mathbin{..} y_T))$$
$$= \delta_E(T[0 \mathbin{..} v), {}^*\mathrm{rot}^{-\rho(v)}(Q)) + \delta_E(T[v \mathbin{..} n), \mathrm{rot}^{-\rho(v)}(Q)^*).$$

Lemma 6.3. *For any two positions $v < w$ of T such that $\mathcal{L}^T_{[v]} = \mathcal{L}^T_{[w]} = \varnothing$, we have*

$$\delta_E(T[v \mathbin{..} w), {}^*Q^*) = \delta_E(T[v \mathbin{..} w), Q^\infty[\rho(v) \mathbin{..} \rho(w))) = \sum_{L\in\mathcal{L}^T_{[v \mathbin{..} w)}} \delta_E(L, {}^*Q^*).$$

Proof. Set $\mathcal{L}^T_{[v \mathbin{..} w)} = \{L^T_j \in \mathcal{L}^T : j \in [j_1 \mathbin{..} j_2]\}$, and observe that all elements of this set are fragments of $T[v \mathbin{..} w)$. Now, we have

$$\delta_E(T[v \mathbin{..} w), Q^\infty[\rho(v) \mathbin{..} \rho(w)))$$
$$= \delta_E(T, Q^\infty[x_T \mathbin{..} y_T)) - \delta_E(T[0 \mathbin{..} v), Q^\infty[x_T \mathbin{..} \rho(v))) - \delta_E(T[w \mathbin{..} n), Q^\infty[\rho(w) \mathbin{..} y_T)))$$
$$\leq d_T - \sum_{j=1}^{j_1-1} \delta_E(L^T_j, {}^*Q^*) - \sum_{j=j_2+1}^{\ell^T} \delta_E(L^T_j, {}^*Q^*)$$
$$= \sum_{j=j_1}^{j_2} \delta_E(L^T_j, {}^*Q^*)$$
$$\leq \delta_E(T[v \mathbin{..} w), {}^*Q^*)$$
$$\leq \delta_E(T[v \mathbin{..} w), Q^\infty[\rho(v) \mathbin{..} \rho(w))).$$
◾

Lemma 6.4. *Consider a light position v of T. If $\delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) < \partial_P$, then there is a* $w \in [v \mathinner{\ldotp\ldotp} n]$ *such that*

$$\delta_E(P, T[v \mathinner{\ldotp\ldotp} w)) \leq \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) + \sum_{L \in \mathcal{L}_{[v \mathinner{\ldotp\ldotp} v+m)}^T} \delta_E(L, {}^*Q^*).$$

Proof. We first prove two auxiliary claims.

Claim 6.5. We have $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) = \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) < \partial_P + d_P$.

Proof. The inequality $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) \geq \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*)$ holds trivially.

By Lemma 5.4, we have $P = L_1^P Q^{\alpha_1} L_2^P Q^{\alpha_2} \cdots Q^{\alpha_{\ell^P}-1} L_{\ell^P}^P$ for some non-negative integers α_i. Since L_1^P is ∂_P-locked, we have $\delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) = \delta_E(L_1^P, Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} jq))$ for some non-negative integer j. Further, for each $i \in (1 \mathinner{\ldotp\ldotp} \ell^P)$, we have $\delta_E(L_i^P, {}^*Q^*) = \delta_E(L_i^P, Q^{\beta_i})$ for some non-negative integer β_i. Finally, we have $\delta_E(L_{\ell^P}^P, {}^*Q^*) = \delta_E(L_{\ell^P}^P, Q^\infty[0 \mathinner{\ldotp\ldotp} x))$ for some non-negative integer x.

Set $\gamma = \alpha_1 + \sum_{i=2}^{\ell^P-1}(\alpha_i + \beta_i)$. The above discussion implies that there is an alignment of P with the prefix $Q[\rho(v) \mathinner{\ldotp\ldotp} jq) Q^\gamma Q^\infty[0 \mathinner{\ldotp\ldotp} x)$ of $Q[\rho(v) \mathinner{\ldotp\ldotp} jq) Q^\infty$ that costs $\delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*)$, thus proving that $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) \leq \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*)$. This concludes the proof of the claimed equality.

The claimed inequality holds since $\delta_E(L_1^P, \mathrm{rot}^{-\rho(a)}(Q)^*) < \partial_P$ and $\sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) \leq d_P$. ⬭

Claim 6.6. There is a $w \in (v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa}) \cap [0 \mathinner{\ldotp\ldotp} n]$ such that $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) = \delta_E(P, Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} \rho(w)))$.

Proof. Recall that $\mathcal{L}_{[v+m-\hat{\kappa} \mathinner{\ldotp\ldotp} v+m+\hat{\kappa})}^T = \varnothing$ holds because v is a light position of T. Since $\mathcal{L}^T$ contains a suffix of T, we conclude that either $n \leq v + m - \hat{\kappa}$ or $v + m + \kappa \leq n$. The former case contradicts $v \leq n - m + k < n - m + \hat{\kappa}$, so $T[v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa})$ must be a fragment of T disjoint with all locked fragments in $\mathcal{L}^T$. This means that $\mathcal{A}_T$ matches $T[v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa})$ without any edits to $\mathcal{A}_T(T[v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa}))$. Since $(v, \rho(v)) \in \mathcal{A}_T$ and the cost of $\mathcal{A}_T$ is d_T, the fragment $\mathcal{A}_T(T[v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa}))$ must contain $Q^\infty[\rho(v) + m - \hat{\kappa} + d_T \mathinner{\ldotp\ldotp} \rho(v) + m + \hat{\kappa} - d_T)$. Consequently, for every $y \in (\rho(v) + m - \hat{\kappa} + d_T \mathinner{\ldotp\ldotp} \rho(v) + m + \hat{\kappa} - d_T)$, there exists $w \in (v + m - \hat{\kappa} \mathinner{\ldotp\ldotp} v + m + \hat{\kappa}) \cap [0 \mathinner{\ldotp\ldotp} n]$ such that $y = \rho(w)$. Since $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) \leq \partial_P + d_P < \hat{\kappa} - d_T$ holds by Claim 6.5, such w exist in particular for y chosen so that $\delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) = \delta_E(P, Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} y))$. ⬭

We get the desired bound by combining Claims 6.5 and 6.6, Fact 6.2, , and Lemma 6.3 via the triangle inequality; observe that Lemma 6.3 is applicable because $\mathcal{L}_{[w]}^T \subseteq \mathcal{L}_{[v+m-\hat{\kappa} \mathinner{\ldotp\ldotp} v+m+\hat{\kappa})}^T = \varnothing$.

$$
\begin{aligned}
&\delta_E(P, T[v \mathinner{\ldotp\ldotp} w)) \\
&\quad \leq \delta_E(P, Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} \rho(w))) + \delta_E(T[v \mathinner{\ldotp\ldotp} w), Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} \rho(w))) && \text{(triangle inequality)} \\
&\quad = \delta_E(P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \delta_E(T[v \mathinner{\ldotp\ldotp} w), Q^\infty[\rho(v) \mathinner{\ldotp\ldotp} \rho(w))) && \text{(Claim 6.6)} \\
&\quad = \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) + \sum_{L \in \mathcal{L}_{[v \mathinner{\ldotp\ldotp} w)}^T} \delta_E(L, {}^*Q^*) && \text{(Claim 6.5 and Lemma 6.3)} \\
&\quad = \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) + \sum_{L \in \mathcal{L}_{[v \mathinner{\ldotp\ldotp} v+m)}^T} \delta_E(L, {}^*Q^*). && \text{(Fact 6.2)}
\end{aligned}
$$

This completes the proof of the lemma. ⬛

For a position v and a locked fragment $L \in \mathcal{L}^P \cup \mathcal{L}^T$, we write $\mathrm{mk}(v, L)$ for the number of marks placed in v due to pairs of locked fragments that contain L; formally:

$$\mathrm{mk}(v, L) = \begin{cases} \sum_{L^T \in \mathcal{L}^T} \mathrm{mk}(v, \hat{\kappa}, L, L^T) & \text{if } L \in \mathcal{L}^P; \\ \sum_{L^P \in \mathcal{L}^P} \mathrm{mk}(v, \hat{\kappa}, L^P, L) & \text{if } L \in \mathcal{L}^T. \end{cases}$$

Definition 6.7. *For a light position v of T, set*

$$\mathcal{D}(v) := \{L_1^P\} \cup \{L \in \mathcal{L}^P \cup \mathcal{L}^T_{[v\,..\,v+m)} : \mathrm{mk}(v, L) < \delta_E(L, {}^*Q^*)\}.$$

Let us now provide some intuition on what follows. Consider an alignment $\mathcal{A} : P \rightsquigarrow T[v\,..\,w)$, where v is a light position of T and the cost of $\mathcal{A}$ is not larger than k. In the next lemma, we essentially lower bound the cost of the restriction of such an alignment $\mathcal{A}$ to each locked fragment. For instance, we lower bound $\delta_E(L, \mathcal{A}(L))$ for each $L \in \mathcal{L}^P$. Our lower bound is positive only for elements of $\mathcal{D}(v)$. Consider some locked fragment $L \in \mathcal{D}(v) \cap \mathcal{L}^P$ other than L_1^P. Roughly speaking, at most $\mathrm{mk}(v, L)$ errors of L with a fragment Q' of Q^∞ cancel out with errors between $\mathcal{A}(L)$ and Q', yielding a lower bound $\delta_E(L, \mathcal{A}(L)) \geq \delta_E(L, {}^*Q^*) - \mathrm{mk}(v, L)$. Then, the definition of $\mathcal{D}(v)$ guarantees that, for any $L \in \mathcal{D}(v) \cap \mathcal{L}^P$, $\mathcal{A}(L)$ is disjoint from all $L' \in \mathcal{D}(v) \cap \mathcal{L}^T$. One can exploit this property to obtain a lower bound for the cost of $\mathcal{A}$ by showing that we can sum over the lower bounds for individual locked fragments in $\mathcal{D}(v)$. In fact, we use this reasoning to lower bound the cost of an alignment between two other strings, obtained from P and T, respectively, via the deletion of some fragments; this happens in the proof of Lemma 6.20 in Section 6.2.2.

Lemma 6.8. *Consider a light position v of T and a locked fragment $L \in \mathcal{L}^P \cup \mathcal{L}^T_{[v\,..\,v+m)}$.*
1. *If $L = T[v + \ell\,..\,v + r) \in \mathcal{L}^T$, then every $U \in \{P[i\,..\,j) : i \geq \ell - k \text{ and } j \leq r + k\}$ satisfies*

$$\delta_E(L, U) \geq \delta_E(L, {}^*Q^*) - \mathrm{mk}(v, L).$$

2. *If $L = P[\ell\,..\,r) \in \mathcal{L}^P \setminus \{L_1^P\}$, then every $U \in \{T[i\,..\,j) : i \geq v + \ell - k \text{ and } j \leq v + r + k\}$ satisfies*

$$\delta_E(L, U) \geq \delta_E(L, {}^*Q^*) - \mathrm{mk}(v, L).$$

3. *If L is L_1^P, then for every $U \in \{T[v\,..\,v + j) : j \in [\,|L| - k\,..\,|L| + k\,]\}$ satisfies*

$$\delta_E(L, U) \geq \delta_E(L, \mathrm{rot}^{-\rho(v)}(Q)^*) - \mathrm{mk}(v, L).$$

Proof. Consider some $L \in \mathcal{L}^P \cup \mathcal{L}^T_{[v\,..\,v+m)}$. If $L \in \mathcal{L}^T$, let $Y = P$; otherwise, let $Y = T$. Further, let $U := Y[u\,..\,w)$ denote any of the fragments specified in the statement of the lemma, and let

$$\zeta := \begin{cases} \delta_E(L_1^P, \mathrm{rot}^{-\rho(v)}(Q)^*) & \text{if } L = L_1^P, \\ \delta_E(L, {}^*Q^*) & \text{otherwise.} \end{cases}$$

Consult Figure 7 for an illustration of the setting.

Let x and y denote integers that satisfy $\delta_E(U, Q^\infty[x\,..\,y)) = \delta_E(U, {}^*Q^*)$. In the case where $L = L_1^P$, we choose $x = \rho(v)$ and $y = \rho(w)$ so that $Q^\infty[x\,..\,y)$ is a prefix of $\mathrm{rot}^{-\rho(v)}(Q)^\infty$. This is allowed by Lemma 6.3 because $\mathcal{L}^T_{[v-\hat{\kappa}\,..\,v+\hat{\kappa})} = \mathcal{L}^T_{[v+|L_1^P|-\hat{\kappa}\,..\,v+|L_1^P|+\hat{\kappa})} = \varnothing$, and it ensures that the following inequality holds in all three cases (the inequality holds trivially if $L \neq L_1^P$):

$$\delta_E(L, Q^\infty[x\,..\,y)) \geq \zeta. \tag{6}$$
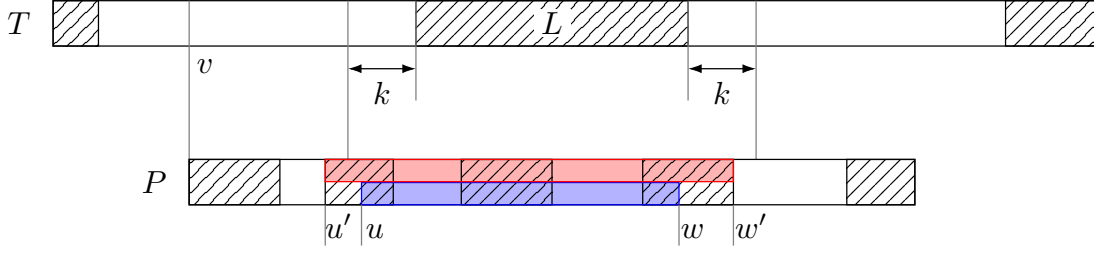
■ **Figure 7** An illustration of the setting in the proof of Lemma 6.8. Locked fragments are distinguished by diagonal stripes. In this example, L is the sole element of $\mathcal{L}^T_{[\,v\,..\,v+m\,)}$ and $Y = P$. The fragment $U = P[\,u\,..\,w\,)$ is shaded in blue, the fragment $P[\,u'\,..\,w'\,)$ is shaded in red.

Further, our marking scheme implies

$$\mathrm{mk}(v, L) \geq \sum_{K \in \mathcal{L}^Y_{[\,u\,..\,w\,)}} \delta_E(K, {}^*Q^*). \tag{7}$$

Let $Y[\,u'\,..\,w'\,)$ denote the fragment of Y that is covered by $Y[\,u\,..\,w\,)$ and the elements of $\mathcal{L}^Y_{[\,u\,..\,w\,)}$. We have

$$\delta_E(U, Q^\infty[\,x\,..\,y\,)) = \delta_E(U, {}^*Q^*) \leq \delta_E(Y[\,u'\,..\,w'\,), {}^*Q^*) = \sum_{K \in \mathcal{L}^Y_{[\,u\,..\,w\,)}} \delta_E(K, {}^*Q^*), \tag{8}$$

where the last equality follows from the properties of locked fragments as computed by Lemma 5.4.

We are now ready to prove the claimed inequality.

$$
\begin{aligned}
\delta_E(L, U) &\geq |\delta_E(L, Q^\infty[\,x\,..\,y\,)) - \delta_E(U, Q^\infty[\,x\,..\,y\,))| && \text{(triangle inequality)} \\
&\geq \delta_E(L, Q^\infty[\,x\,..\,y\,)) - \delta_E(U, Q^\infty[\,x\,..\,y\,)) \\
&\geq \zeta - \delta_E(U, Q^\infty[\,x\,..\,y\,)) && \text{(due to (6))} \\
&\geq \zeta - \sum_{K \in \mathcal{L}^Y_{[\,i\,..\,j\,)}} \delta_E(K, {}^*Q^*) && \text{(due to (8))} \\
&\geq \zeta - \mathrm{mk}(v, L). && \text{(due to (7))}
\end{aligned}
$$

This completes the proof of the lemma. ◢

6.2.2 Shrinking Runs of Plain Pairs to Length $o(\sqrt{d})$

We set $\mathsf{Red}(P)$ to be equal to

$$\mathsf{Special}(P) \cup \{P_i : i \in (\,1\,..\,z\,) \text{ and } \exists_{P[\,\ell\,..\,r\,) \in \mathcal{L}^P}[\,\ell - 13\hat\kappa\,..\,r + 13\hat\kappa\,) \cap [\,p_i\,..\,p_{i+1} + \Delta\,) \neq \varnothing\}.$$

Similarly, we set $\mathsf{Red}(T)$ to be equal to

$$\mathsf{Special}(T) \cup \{T_i : i \in (\,\min J + 1\,..\,\max J + z\,) \text{ and } \exists_{T[\,\ell\,..\,r\,) \in \mathcal{L}^T}[\,\ell - 13\hat\kappa\,..\,r + 13\hat\kappa\,) \cap [\,t_i\,..\,t_{i+1} + \Delta\,) \neq \varnothing\}.$$

Next, we upper-bound the sizes of these two sets and show how to construct them efficiently.

▗ **Lemma 6.9.** *Given P, T, $\mathcal{A}_P$, $\mathcal{A}_T$, $\mathcal{L}^P$, and $\mathcal{L}^T$, the sets $\mathsf{Red}(P)$ and $\mathsf{Red}(T)$ are of size $\mathcal{O}(d)$ and can be constructed in time $\mathcal{O}(d \log \log d)$ in the* PILLAR *model.*

Proof. We start with upper-bounding the size of each of the sets $\mathsf{Red}(P)$ and $\mathsf{Red}(T)$. $\mathsf{Special}(P)$ and $\mathsf{Special}(T)$ are of size $\mathcal{O}(d)$ due to Lemma 1.6. The task is therefore to bound, for each of P and T, the number of internal pieces that are within $13\hat{\kappa}$ positions of a locked fragment.

Let (S, β_S) denote either of (P, β_P) or (T, β_T) and $\mathcal{L}^S = \{L[\ell_j \mathinner{.\,.} h_j) : j \in [1 \mathinner{.\,.} \ell^S]\}$. It suffices to upper bound the number of tiles $S[s_{i-1} \mathinner{.\,.} s_i)$ with $i \in (1 \mathinner{.\,.} \beta_S)$ in the τ-tile partition of S (with respect to $\mathcal{A}_S$) such that $[s_{i-1} \mathinner{.\,.} s_i)$ overlaps $[\ell_j - \Delta - 13\hat{\kappa} \mathinner{.\,.} h_j + 13\hat{\kappa})$ for some $j \in [1 \mathinner{.\,.} \ell^S]$; that is,

$$|\{i \in (1 \mathinner{.\,.} \beta_S) : \exists_{j \in [1 \mathinner{.\,.} \ell^S]}[\ell_j - \Delta - 13\hat{\kappa} \mathinner{.\,.} h_j + 13\hat{\kappa}) \cap [s_{i-1} \mathinner{.\,.} s_i) \neq \varnothing\}|.$$

Intuitively, we extend each locked fragment in $\mathcal{L}^S$ by $\Delta + 13\hat{\kappa}$ characters to the left and by $13\hat{\kappa}$ characters to the right, thus covering at most $\|\mathcal{L}^S\| + \ell^S \cdot (2 \cdot 13 + 6)\hat{\kappa} = \|\mathcal{L}^S\| + 32\ell^S\hat{\kappa}$ positions of S, since $\Delta = 6\kappa \leq 6\hat{\kappa}$. Let us first upper bound the size μ of the set M of tiles (other than the first and last ones) that are fully covered by these "extended" locked fragments, that is, $M = \{i \in (0 \mathinner{.\,.} \beta_S) : [s_{i-1} \mathinner{.\,.} s_i) \subseteq \bigcup_{j \in [1 \mathinner{.\,.} \ell^S]}[\ell_j - \Delta - 13\hat{\kappa} \mathinner{.\,.} h_j + 13\hat{\kappa})\}$. We have

$$\sum_{i \in M}(\tau - |S[s_{i-1} \mathinner{.\,.} s_i)|) \leq \sum_{i \in M}\delta_E(S[s_{i-1} \mathinner{.\,.} s_i), Q[(i-1)\tau \mathinner{.\,.} i\tau)) \leq \delta_E(S, {}^*Q^*)$$

and hence

$$\mu \cdot \tau - \delta_E(S, {}^*Q^*) \leq \sum_{i \in M}|S[s_{i-1} \mathinner{.\,.} s_i)| \leq \|\mathcal{L}^S\| + 32\ell^S\hat{\kappa},$$

which is equivalent to

$$\mu \leq \frac{\|\mathcal{L}^S\| + 32\ell^S\hat{\kappa} + \delta_E(S, {}^*Q^*)}{\tau}.$$

Finally, we have to account for the at most $2\ell^S$ tiles that overlap "extended" locked fragments but are not fully contained in them; we have at most two such tiles for each $L \in \mathcal{L}^S$.

Since $\ell^S = \mathcal{O}(\delta_E(S, {}^*Q^*))$ and $\|\mathcal{L}^S\| = \mathcal{O}(\delta_E(S, {}^*Q^*) \cdot q)$ by Lemma 5.4, $\delta_E(S, {}^*Q^*) = \mathcal{O}(\kappa)$, and $\tau = \Theta(\max\{\kappa, q\})$, we have

$$2\ell^S + \frac{\|\mathcal{L}^S\| + 32\ell^S\hat{\kappa} + \delta_E(S, {}^*Q^*)}{\tau} = \mathcal{O}\left(\kappa + \frac{\kappa q + \kappa\hat{\kappa} + \kappa}{\tau}\right) = \mathcal{O}(\hat{\kappa}) = \mathcal{O}(d).$$

Let us now show how to efficiently construct the sets in scope. First, recall that $\mathsf{Special}(P)$ and $\mathsf{Special}(T)$ can be constructed in $\mathcal{O}(d)$ time in the `PILLAR` model due to Lemma 1.6. The remaining elements of $\mathsf{Red}(P)$ (resp. $\mathsf{Red}(T)$) can be computed in a simultaneous left-to-right scan of:

- the representation of starting positions of internal pieces p_i (resp. t_i) as $\mathcal{O}(d)$ arithmetic progressions, which can be computed in $\mathcal{O}(d)$ time due to Lemma 4.4;
- the $\mathcal{O}(d)$ locked fragments of P (resp. T) sorted with respect to their starting positions.

The $\mathcal{O}(d \log \log d)$ time required for sorting the starting positions of the locked fragments, using the algorithm of [AHNR98], is the bottleneck of the algorithm in the `PILLAR` model. ∎

Definition 6.10. *For $j \in J$, set $\mathcal{I}'_j := (F_{j,1}, G_{j,1}) \cdots (F_{j,z_j}, G_{j,z_j}) := \mathsf{Trim}(\mathcal{I}_j, \mathsf{Red}(P), \mathsf{Red}(T), 2\eta + 30)$, and $F_j := \mathsf{val}_\Delta(F_{j,1}, \ldots, F_{j,z_j})$ and $G_j := \mathsf{val}_\Delta(G_{j,1}, \ldots, G_{j,z_j})$.*

Further, let $F_{j,i}$ and $G_{j,i}$ correspond to the fragments $F_j[f_{j,i} \mathinner{.\,.} f_{j,i+1} + \Delta)$ and $G_j[g_{j,i} \mathinner{.\,.} g_{j,i+1} + \Delta)$, respectively, that is, we set $f_{j,1} = g_{j,1} = 0$ and, for $i \in [2 \mathinner{.\,.} z_j + 1]$, $f_{j,i} = (\sum_{x<i}|F_{j,x}|) - \Delta$ and $g_{j,i} = (\sum_{x<i}|G_{j,x}|) - \Delta$. ∎

The focus of the remainder of Section 6.2.2 is to prove the following lemma.

▟ **Lemma 6.11**. $\operatorname{Occ}_k^E(P,T) \cap \mathbb{L} = \left(\bigcup_{j \in J}(r_j + \operatorname{Occ}_k^E(\mathcal{I}_j'))\right) \cap \mathbb{L} = \left(\bigcup_{j \in J}(r_j + \operatorname{Occ}_k^E(F_j, G_j))\right) \cap \mathbb{L}.$ ◾

The combination of Corollary 3.11 and Lemma 4.32 directly yields the following.

▟ **Corollary 6.12**. $\left(\bigcup_{j \in J}(r_j + \operatorname{Occ}_k^E(F_j, G_j))\right) \supseteq \left(\bigcup_{j \in J}(r_j + \operatorname{Occ}_k^E(P, R_j))\right) = \operatorname{Occ}_k^E(P,T).$ ◾

The following—skippable—example illustrates that, for some $j \in J$ and $p \in \operatorname{Occ}_k^E(F_j, G_j)$, we might have $r_j + p \notin \operatorname{Occ}_k^E(P,T)$ if $r_j + p \in \mathbb{H}$.

▟ **Example 6.13**. Let $k > 6$ denote an integer and set $\Sigma = \{\mathtt{a}, \mathtt{b}\}$. Further, set $U = \mathtt{a}^i\mathtt{b}$ and $V = \mathtt{a}^{i+1}\mathtt{b}$ for $i = \lfloor 2k/3 \rfloor$. Set $Q = UVU$ and $P = Q^y(UVV)^{k+1}Q^y$ and $T = Q^w(VVU)^{k+1}Q^w$ for some integers y and w that satisfy $w > y > k$, ensuring that $d := 2k \le |P|/8|Q|$. We have a valid instance of the NEWPERIODICMATCHES problem with $d_P = d_T = k + 1$ and $\tau = \Theta(q)$.

Consider a set of locked fragments (with respect to Q) for each of P and T such that the only locked fragment of P (resp. T) that is not its prefix or suffix is $L_2^P := P[y|Q| \mathinner{.\,.} y|Q| + (k+1)(3i+5)) = (UVV)^{k+1}$ (or $L_2^T := [w|Q| \mathinner{.\,.} w|Q| + (k+1)(3i+5)) = (VVU)^{k+1}$).[19] Let $\eta = \mathcal{O}(\sqrt{k})$ denote the threshold used in the marking and assume that k is large enough so that $\eta < k$. Consider some $j \in J$ and a position p of R_j such that $Q^y(VVU)^{k+1}Q^y$ is a prefix of $R_j[p \mathinner{.\,.} r_j')$, noting that $r_j + p$ has at least $\min\{\delta_E(L_2^P, {}^*Q^*), \delta_E(L_2^T, {}^*Q^*)\} = k + 1 > \eta$ marks, and is thus heavy.

We next argue that $\min_t(\delta_E(P, R_j[p \mathinner{.\,.} t))) > k$. Toward a contradiction, suppose that there exists an integer t' and an alignment $\mathcal{A} : P \rightsquigarrow R_j[p \mathinner{.\,.} t')$ of cost at most k. Then, at least one of the first/last y copies of Q in P is matched exactly by $\mathcal{A}$ since $y > k$. In addition, at least one of the copies of UVV is matched exactly by $\mathcal{A}$. Thus, $\mathcal{A}$ makes at least $|U|$ edits in order to "synchronize" a copy of VVU in P to a copy of VVU in T, and then at least $|U|$ more edits to "synchronize" copies of Q. Hence, we have $\min_t(\delta_E(P, R_j[p \mathinner{.\,.} t))) \ge 2|U| = 2(\lfloor 2k/3 \rfloor + 1) \ge 4k/3 > k$.

On the other hand, for some integers $y_1 \le y$ and $y_2 \le c\eta$, where c is a constant independent from k, we have

$$F_j = Q^{y_1}(UVV)^{k+1}Q^{y_2} = Q^{y_1}U(VVU)^{k+1}VUQ^{y_2-1} \text{ and}$$
$$G_j[p \mathinner{.\,.} |G_j|) = Q^{y_1}(VVU)^{k+1}Q^{y_2}W, \text{ for some string } W.$$

It is easy to observe that

$$\min_t \delta_E(F_j, G_j[p \mathinner{.\,.} t))$$
$$\le \delta_E(Q^{y_1}, Q^{y_1}) + \delta_E(U, \varepsilon) + \delta_E((VVU)^{k+1}, (VVU)^{k+1}) + \delta_E(VUQ^{y_2-1}, Q^{y_2-1}UV)$$
$$\le |U| + \delta_E(VUQ^{y_2-1}, U^{3y_2-1}) + \delta_E(U^{3y_2-1}, Q^{y_2-1}UV)$$
$$= i + 1 + 2y_2.$$

We can assume that k is large enough so that $k > 12c\eta$, in which case $i + 1 + 2y_2 < 2k/3 + k/6 + k/6 = k$. The point is that the implied alignment pays $|U|$ to "synchronize" copies of VVU, but it can then afford to pay for the (fewer than k) misaligned copies of Q without needing to "synchronize" again. ◾

It remains to show that, for any light position $p \in \left(\bigcup_{j \in J}(r_j + \operatorname{Occ}_k^E(F_j, G_j))\right)$, we have $p \in \operatorname{Occ}_k^E(P,T)$. The following lemma demonstrates that, for each $j \in J$, we can restrict our attention to a subset of the positions of R_j.

---

[19] These are not precisely the locked fragments that would be computed by the algorithm underlying Lemma 5.4, but they are consistent with the properties that need to be satisfied and are easier to work with for the sake of this example.

Lemma 6.14. *Consider some $j \in J$ and a position p of R_j such that $r_j + p$ is a light position of T. If $\delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j+p)}(Q)^*) \geq \partial_P$, then $p \notin \mathrm{Occ}_k^E(F_j, G_j)$.*

Proof. By the definition of $\mathsf{Red}(T)$, it follows that $|F_j| \geq |L_1^P|$ and any prefix of F_j of length at most $|L_1^P| + 13\hat{\kappa}$ is also a prefix of P. Combined with the upper bound on the sum of length-differences from Lemma 4.10, this also implies that any prefix of G_j of length at most $|L_1^P| + 13\hat{\kappa} - 3\kappa$ is also a prefix of R_j. Further, recall that we have $|R_j| \leq |P| + 3\kappa - k$ due to Lemma 4.10. Consequently, since $|R_j| - |F_j| = |P| - |G_j|$, $\mathrm{Occ}_k^E(F_j, G_j) \subseteq [0 \mathbin{.\,.} |G_j| - |F_j| + k] \subseteq [0 \mathbin{.\,.} 3\kappa]$.

It thus suffices to consider the case where $p \in [0 \mathbin{.\,.} 3\kappa]$. Let W denote a prefix of $G_j[p \mathbin{.\,.} |G_j|)$ that satisfies $\delta_E(F_j[0 \mathbin{.\,.} |L_1^P|), W) = \min_t \delta_E(F_j[0 \mathbin{.\,.} |L_1^P|), G_j[p \mathbin{.\,.} t))$. We distinguish between two cases:

- If $|W| \in [\,|L_1^P| - k \mathbin{.\,.} |L_1^P| + k\,]$, then W is a prefix of $T[r_j + p \mathbin{.\,.} n)$ since $p + |W| \leq 3\kappa + |L_1^P| + k \leq |L_1^P| + 13\hat{\kappa} - 3\kappa$. Thus, by Lemma 6.8, we have

$$\delta_E(L_1^P, W) \geq \delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j+p)}(Q)^*) - \mathrm{mk}(r_j + p, L_1^P) \geq \partial_P - \mathrm{mk}(r_j + p, L_1^P) \geq 2d - \eta \geq d > k.$$

- Otherwise, we have $\delta_E(F_j[0 \mathbin{.\,.} |L_1^P|), W) \geq \big||W| - |L_1^P|\big| > k$.

To conclude the proof, it suffices to observe that $\min_t \delta_E(F_j, G_j[p \mathbin{.\,.} t)) \geq \min_t \delta_E(L_1^P, G_j[p \mathbin{.\,.} t)) > k$, and hence $p \notin \mathrm{Occ}_k^E(F_j, G_j)$. ∎

In what follows, for convenience, we assume that, for each run of plain pairs of $\mathcal{I}_j$ that has been trimmed, the deleted pairs correspond to a suffix of this run. This yields a natural mapping from pairs of $\mathcal{I}_j'$ to pairs of $\mathcal{I}_j$.

Definition 6.15. *For each $j \in J$ and each $i \in [1 \mathbin{.\,.} z_j]$, let $\mathsf{orig}(j, i)$ denote the number of pairs to the left of pair $(F_{j,i}, G_{j,i})$ that were deleted in the process of obtaining $\mathcal{I}_j'$ from $\mathcal{I}_j$. We say that pair $(F_{j,i}, G_{j,i})$ originates from pair $(P_{i+\mathsf{orig}(j,i)}, T_{j,i+\mathsf{orig}(j,i)})$.* ∎

Let each internal pair of pieces $(F_{j,i}, G_{j,i})$ inherit the color of $(P_{i+\mathsf{orig}(j,i)}, T_{j,i+\mathsf{orig}(j,i)})$. In addition, mark the first and the last pairs as not plain.

Definition 6.16. *For $j \in J$ and $i \in [1 \mathbin{.\,.} z_j]$, we say that $F_j[x_1 \mathbin{.\,.} x_2)$ (or $G_j[x_1 \mathbin{.\,.} x_2)$) has an overlap with a pair $(F_{j,i}, G_{j,i})$ if and only if $[x_1 \mathbin{.\,.} x_2) \cap [f_{j,i} \mathbin{.\,.} f_{j,i+1} + \Delta) \neq \varnothing$ (or $[x_1 \mathbin{.\,.} x_2) \cap [g_{j,i} \mathbin{.\,.} g_{j,i+1} + \Delta) \neq \varnothing$).* ∎

Definition 6.17. *Consider some $j \in J$ and a contiguous sequence $\mathcal{M} = (F_{j,i_1}, G_{j,i_1}) \cdots (F_{j,i_2}, G_{j,i_2})$ of pairs in $\mathcal{I}_j'$ that are either all plain or all not plain (that is, $\mathcal{M}$ is monochromatic). Fix an $X \in \{F_j, G_j\}$. For $i \in [1 \mathbin{.\,.} z_j]$, if $X = F$, we set $x_{j,i} = f_{j,i}$; otherwise, we set $x_{j,i} = g_{j,i}$.*
For a non-negative integer γ, we say that a fragment $X[x_1 \mathbin{.\,.} x_2)$ of $X \in \{F_j, G_j\}$ is γ-contained by $\mathcal{M}$ when the following two conditions are satisfied: (a) $x_1 \geq x_{j,i_1} + \gamma$ or $i_1 = 1$ and (b) $x_2 \leq x_{j,i_2+1} + \Delta - \gamma$ or $i_2 = z_j$. ∎

Fact 6.18. *If a fragment U of F or G is Δ-contained by a monochromatic sequence $\mathcal{M}$ of contiguous pairs in $\mathcal{I}_j'$, then U only overlaps pairs in $\mathcal{M}$.* ∎

Lemma 6.19. *For $j \in J$, consider a monochromatic sequence $\mathcal{M} = (F_{j,i_1}, G_{j,i_1}) \cdots (F_{j,i_2}, G_{j,i_2})$ of contiguous pairs in $\mathcal{I}_j'$. Let $\{X, Y\} = \{F_j, G_j\}$ and $X[x_1 \mathbin{.\,.} x_2)$ be a fragment of X that is γ-contained by $\mathcal{M}$ for $\gamma \geq 13\kappa$. All the pairs of $\mathcal{I}_j'$ that overlap $Y[y_1 \mathbin{.\,.} y_2) := Y[\max\{0, x_1 - 4\kappa\} \mathbin{.\,.} \min\{|Y|, x_2 + 4\kappa\})$ belong to $\mathcal{M}$.*

Proof. For $i \in [1 .. z_j]$, if $X = F_j$ and $Y = G_j$, let $x_{j,i} = f_{j,i}$ and $y_{j,i} = g_{j,i}$; otherwise, let $x_{j,i} = g_{j,i}$ and $y_{j,i} = f_{j,i}$.

Recall that the sum of length-differences of pairs in $\mathcal{I}'_j$ is at most $3\kappa - k$, that is, $|x_{j,i} - y_{j,i}| \leq 3\kappa - k$ for all i. Let us first show that $Y[y_1 .. y_2)$ does not overlap any of the first $i_1 - 1$ pairs. The case where $i_1 = 1$ is trivial. In the remaining case,

$$y_1 \geq x_1 - 4\kappa \geq x_{j,i_1} + \gamma - 4\kappa \geq y_{j,i_1} + 13\kappa - 4\kappa - (3\kappa - k) \geq y_{j,i_1} + (13 - 7)\kappa = y_{j,i_1} + \Delta.$$

We next show that $Y[y_1 .. y_2)$ does not overlap any of the last $z_j - i_2$ pairs, thus concluding the proof. The case where $i_2 = z_j$ is trivial. In the remaining case,

$$y_2 \leq x_2 + 4\kappa \leq x_{j,i_2+1} + \Delta - \gamma + 4\kappa \leq y_{j,i_2+1} + \Delta - 13\kappa + 4\kappa + (3\kappa - k) \leq y_{j,i_2+1} + \Delta - (13-7)\kappa = y_{j,i_2+1}.$$

We are done by a direct application of Fact 6.18. ∎

⬛ **Lemma 6.20.** *Consider some $j \in J$ and a position $p \in \mathrm{Occ}^E_k(F_j, G_j)$ such that $r_j + p$ is a light position of T. Then, $r_j + p \in \mathrm{Occ}^E_k(P, T)$.*

Proof. First, observe that we may assume that $\mathcal{I}'_j \neq \mathcal{I}_j$; otherwise, the statement follows trivially. To avoid clutter, we drop the subscript j when referring to F_j and G_j and simply call them F and G, respectively.

Let b denote a position of G and let $\mathcal{B} : F \rightsquigarrow G[p .. b)$ denote an alignment of cost $\min_t \delta_E(F, G[p .. t)) \leq k$. We intend to show that, in this case, there exist a position c of R_j and an alignment $\mathcal{C} : P \rightsquigarrow R_j[p .. c)$ of the same cost.

Set

$$\mathcal{L}^T_{[r_j+p .. r_j+p+m)} = \{L^T_i : i \in [i_1 .. i_2]\}.$$

Observe that there is a natural mapping of each locked fragment $L^P_i \in \mathcal{L}^P$ to a fragment of F, which we denote by L^F_i; we denote the set of fragments in the image of this mapping by $\mathcal{L}^F$. Similarly, there is a natural mapping of each locked fragment $L^T_i \in \{L^T_i : i \in [i_1 .. i_2]\}$ to a fragment of G, which we denote by L^G_i; we denote the set of fragments in the image of this mapping by $\mathcal{L}^G$. Let $\mathcal{D}'(p)$ consist of the images of the locked fragments in $\mathcal{D}(r_j + p)$ under these mappings. Observe that each fragment $L \in \mathcal{L}^F \cup \mathcal{L}^G$ only overlaps non-plain pairs. For a fragment $L^X_y \in \mathcal{L}^X$, where $X \in \{P, T, F, G\}$, let $L^X_y = X[\ell^X_y .. r^X_y)$. The following claim follows instantly.

▱ **Claim 6.21.** Each fragment $L \in \mathcal{L}^F \cup \mathcal{L}^G$ is $13\hat{\kappa}$-contained by a sequence of contiguous non-plain pairs in $\mathcal{I}'_j$. ▱

We next essentially show that, if we were to mark positions of G_j based on overlaps of pairs of fragments in $\mathcal{L}^F \times \mathcal{L}^G$, consistently with Definition 5.5, position p of G_j would get the same number of marks as position $r_j + p$ of T.

▱ **Claim 6.22.** For all $x \in [1 .. |\mathcal{L}^P|]$ and $y \in [i_1 .. i_2]$, we have

$$\left|[\ell^G_y - \hat{\kappa} .. r^G_y + \hat{\kappa}) \cap [p + \ell^F_x .. p + r^F_x)\right| = \left|[\ell^T_y - \hat{\kappa} .. r^T_y + \hat{\kappa}) \cap [r_j + p + \ell^P_x .. r_j + p + r^P_x)\right|.$$

Proof. Consider sequences $\mathcal{M}_x = (F_{j,w_1}, G_{j,w_1}) \cdots (F_{j,w_2}, G_{j,w_2})$ and $\mathcal{M}_y = (F_{j,w_3}, G_{j,w_3}) \cdots (F_{j,w_4}, G_{j,w_4})$ of contiguous non-plain pairs of $\mathcal{I}'_j$ that $13\hat{\kappa}$-contain L^F_x and L^G_y, respectively, and are maximal in the sense that they cannot be extended and remain monochromatic. (Such sequences exist by Claim 6.21.)

First, consider the case where $\mathcal{M}_x$ and $\mathcal{M}_y$ do not coincide. We treat the case where $\mathcal{M}_x$ lies to the left of $\mathcal{M}_y$; the other case can be handled analogously. In this case, $w_2 < z_j$ and $w_3 > 1$. Recall that

$p \in [0 \mathinner{\ldotp\ldotp} 3\kappa]$. By Lemma 6.19, we have that $G[p+\ell_x^F \mathinner{\ldotp\ldotp} p+r_x^F)$ only overlaps pairs in $\mathcal{M}_x$ and hence it is disjoint from $[\ell_y^G - \hat{\kappa} \mathinner{\ldotp\ldotp} r_y^G + \hat{\kappa})$ since $p + r_x^F \leq g_{j,w_3} \leq \ell_y^G - \hat{\kappa}$. Now, observe that $\operatorname{orig}(j, w_1) \leq \operatorname{orig}(j, w_3)$ and hence

$$r_j + p + r_x^P = r_j + \operatorname{orig}(j, w_1)\tau + p + r_x^F < r_j + \operatorname{orig}(j, w_3)\tau + \ell_y^G - \hat{\kappa} = \ell_y^T - \hat{\kappa}.$$

Thus, in this case, both considered intersections are empty.

Otherwise, $\mathcal{M}_x$ and $\mathcal{M}_y$ coincide. We then have

$$r_j + \operatorname{orig}(j, w_1)\tau + [p + \ell_x^F \mathinner{\ldotp\ldotp} p + r_x^F) = r_j + [p + \ell_x^P \mathinner{\ldotp\ldotp} p + r_x^P) = [r_j + p + \ell_x^P \mathinner{\ldotp\ldotp} r_j + p + r_x^P)$$

and

$$r_j + \operatorname{orig}(j, w_1)\tau + [\ell_y^G - \hat{\kappa} \mathinner{\ldotp\ldotp} r_y^G + \hat{\kappa}) = r_j + [\ell_y^T - r_j - \hat{\kappa} \mathinner{\ldotp\ldotp} r_y^T - r_j + \hat{\kappa}) = [\ell_y^T - \hat{\kappa} \mathinner{\ldotp\ldotp} r_y^T + \hat{\kappa}).$$

The statement readily follows in the considered case. ⌟

Let $\mathcal{B} = (f_v, g_v)_{v=0}^u$. For each $L_i^F \in \mathcal{D}'(p)$, let $[a_i^F \mathinner{\ldotp\ldotp} b_i^F) \subseteq [0 \mathinner{\ldotp\ldotp} u]$ so that $L_i^F = F[f_{a_i^F} \mathinner{\ldotp\ldotp} f_{b_i^F})$ and $\mathcal{B}(L_i^F) = G[g_{a_i^F} \mathinner{\ldotp\ldotp} g_{b_i^F})$. Symmetrically, for each $L_i^G \in \mathcal{D}'(p)$, let $[a_i^G \mathinner{\ldotp\ldotp} b_i^G) \subseteq [0 \mathinner{\ldotp\ldotp} u]$ so that $L_i^G = G[g_{a_i^G} \mathinner{\ldotp\ldotp} g_{b_i^G})$ and $\mathcal{B}^{-1}(L_i^G) = F[f_{a_i^G} \mathinner{\ldotp\ldotp} f_{b_i^G})$. Further, let $\mathcal{E}$ denote the multiset union of the multisets

$$\mathcal{E}^F = \{[a_i^F \mathinner{\ldotp\ldotp} b_i^F) : L_i^F \in \mathcal{L}^F \cap \mathcal{D}'(p)\} \quad \text{and} \quad \mathcal{E}^G = \{[a_i^G \mathinner{\ldotp\ldotp} b_i^G) : L_i^G \in \mathcal{L}^G \cap \mathcal{D}'(p)\}.$$

In fact, the following claim implies that the multiplicity of each element of $\mathcal{E}$ is one.

▷ **Claim 6.23.** The intervals in $\mathcal{E}$ are pairwise disjoint.

Proof. First, observe that the elements of each of $\mathcal{L}^F$ and $\mathcal{L}^G$ are pairwise disjoint and hence the elements of each of $\mathcal{E}^F$ and $\mathcal{E}^F$ are pairwise disjoint.

Now, consider any two fragments $L_x^F \in \mathcal{D}'(p)$ and $L_y^G \in \mathcal{D}'(p)$. Observe that since $L_x^F, L_y^T \in \mathcal{D}(r_j + p)$, we have $\operatorname{mk}(r_j + p, \hat{\kappa}, L_x^P, L_y^T) = 0$. Hence, $[\ell_y^T - \hat{\kappa} \mathinner{\ldotp\ldotp} r_y^T + \hat{\kappa}) \cap [r_j + p + \ell_x^P \mathinner{\ldotp\ldotp} r_j + p + r_x^P) = \varnothing$. By Claim 6.22, we also have $[g_{a_y^G} - \hat{\kappa} \mathinner{\ldotp\ldotp} g_{b_y^G} + \hat{\kappa}) \cap [p + f_{a_x^F} \mathinner{\ldotp\ldotp} p + f_{b_x^F}) = \varnothing$.

Since the cost of $\mathcal{B}$ is no more than k, $g_i \in [p + f_i - k \mathinner{\ldotp\ldotp} p + f_i + k]$ holds for all $i \in [0 \mathinner{\ldotp\ldotp} u]$. In particular, we have $[g_{a_y^G} - k - 1 \mathinner{\ldotp\ldotp} g_{b_y^G} + k] \supseteq [p + f_{a_y^G} - 1 \mathinner{\ldotp\ldotp} p + f_{b_y^G} + 1)$. Since $[g_{a_y^G} - \hat{\kappa} \mathinner{\ldotp\ldotp} g_{b_y^G} + \hat{\kappa}) \supseteq [g_{a_y^G} - k - 1 \mathinner{\ldotp\ldotp} g_{b_y^G} + k]$, we then have $[f_{a_y^G} - 1 \mathinner{\ldotp\ldotp} f_{b_y^G} + 1) \cap [f_{a_x^F} \mathinner{\ldotp\ldotp} f_{b_x^F}) = \varnothing$, that is, $[f_{a_y^G} \mathinner{\ldotp\ldotp} f_{b_y^G}] \cap [f_{a_x^F} \mathinner{\ldotp\ldotp} f_{b_x^F}] = \varnothing$. This implies $[a_y^G \mathinner{\ldotp\ldotp} b_y^G) \cap [a_x^F \mathinner{\ldotp\ldotp} b_x^F) = \varnothing$; consequently, the intervals in $\mathcal{E}$ are pairwise disjoint. ⌟

Set $\Lambda := \delta_E(L_1^P, \operatorname{rot}^{-\rho(r_j+p)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) + \sum_{i=i_1}^{i_2} \delta_E(L_i^T, {}^*Q^*).$

▷ **Claim 6.24.**

$$\sum_{L_i^F \in \mathcal{D}_j'(p)} \delta_E(L_i^F, \mathcal{B}(L_i^F)) + \sum_{L_i^G \in \mathcal{D}_j'(p)} \delta_E(L_i^G, \mathcal{B}^{-1}(L_i^G)) \geq \Lambda - 2\operatorname{mk}(r_j + p).$$

Proof. Using Lemma 6.8, we lower bound each individual term of the left-hand side of the proved inequality. We consider three cases.

1. Consider some $L_i^G = G[p + \ell \mathinner{\ldotp\ldotp} p + r) \in \mathcal{D}'_j(p)$ and let $L_i^T = [r_j + p + \ell' \mathinner{\ldotp\ldotp} r_j + p + r')$. Since $\delta_E(F, G[p \mathinner{\ldotp\ldotp} b)) \leq k$, we have $\mathcal{B}^{-1}(L_i^G) \preccurlyeq F[\max\{0, \ell - k\} \mathinner{\ldotp\ldotp} \min\{|F|, r + k\})$. By Claim 6.21, L_i^G is $13\hat{\kappa}$-contained by a sequence $\mathcal{M}$ of contiguous non-plain pairs in $\mathcal{I}'_j$. Then, by Lemma 6.19, $F[\max\{0, \ell - k\} \mathinner{\ldotp\ldotp} \min\{|F|, r + k\})$ only overlaps pairs of $\mathcal{I}'_j$ that are in $\mathcal{M}$ and hence it is a substring of $P[(\ell' - \ell) + \max\{0, \ell - k\} \mathinner{\ldotp\ldotp} (r' - r) + \min\{|F|, r + k\})$, which in turn is a substring of $P[(\max\{0, \ell' - k\} \mathinner{\ldotp\ldotp} \min\{|P|, r' + k\})$. Thus, $\delta_E(L_i^G, \mathcal{B}^{-1}(L_i^G)) \geq \delta_E(L_j^T, {}^*Q^*) - \mathrm{mk}(r_j + p, L_j^T)$ holds by Lemma 6.8.

2. Consider some $L_i^F = F[\ell \mathinner{\ldotp\ldotp} r) \in \mathcal{D}'_j(p) \setminus \{L_1^P\}$ and let $L_i^P = P[\ell' \mathinner{\ldotp\ldotp} r')$. Since $\delta_E(F, G[p \mathinner{\ldotp\ldotp} b)) \leq k$, we have $\mathcal{B}(L_i^F) \preccurlyeq G[\max\{p, p + \ell - k\} \mathinner{\ldotp\ldotp} \min\{|G|, p + r + k\})$. As before, by combining Claim 6.21 and Lemma 6.19 we get that $G[\max\{p, p + \ell - k\} \mathinner{\ldotp\ldotp} \min\{|G|, p + r + k\})$ is a substring of $T[r_j + (\ell' - \ell) + \max\{p, p + \ell - k\} \mathinner{\ldotp\ldotp} r_j + (r' - r) + \min\{|G|, p + r + k\})$, which in turn is a substring of $T[r_j + p + \max\{0, \ell' - k\} \mathinner{\ldotp\ldotp} r_j + p + \min\{|T|, r' + k\})$. Thus, $\delta_E(L_i^F, \mathcal{B}(L_i^F)) \geq \delta_E(L_i^P, {}^*Q^*) - \mathrm{mk}(r_j + p, L_i^P)$ holds by Lemma 6.8.

3. Lastly, since $\delta_E(F, G[p \mathinner{\ldotp\ldotp} b)) \leq k$, we have that $\mathcal{B}(L_1^F)$ is a prefix of $T[r_j + p \mathinner{\ldotp\ldotp} r_j + p + |L_1^P| + k)\}$, and hence $\delta_E(L_1^F, \mathcal{B}(L_1^F)) \geq \delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j + p)}(Q)^*) - \mathrm{mk}(r_j + p, L_1^P)$ holds by Lemma 6.8.

We now put everything together. In the following inequalities, we use the fact that, for each $L \in (\mathcal{L}^P \cup \mathcal{L}^T_{[r_j + p \mathinner{\ldotp\ldotp} r_j + p + m)}) \setminus \mathcal{D}(r_j + p, B)$, we have $\delta_E(L, {}^*Q^*) \leq \mathrm{mk}(r_j + p, L)$, and hence the sum of $\delta_E(L, {}^*Q^*) - \mathrm{mk}(r_j + p, L)$ over all such L is at most zero.

$$\sum_{L_i^F \in \mathcal{D}'_j(p)} \delta_E(L_i^F, \mathcal{B}(L_i^F)) + \sum_{L_i^G \in \mathcal{D}'_j(p)} \delta_E(\mathcal{B}^{-1}(L_i^G), L_i^G)$$

$$\geq \delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j + p)}(Q)^*) - \mathrm{mk}(r_j + p, L_1^P) + \sum_{L_i^P \in \mathcal{D}(r_j + p) \setminus \{L_1^P\}} (\delta_E(L_i^P, {}^*Q^*) - \mathrm{mk}(r_j + p, L_i^P))$$

$$+ \sum_{L_i^T \in \mathcal{D}(r_j + p)} (\delta_E(L_i^T, {}^*Q^*) - \mathrm{mk}(r_j + p, L_i^T))$$

$$\geq \delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j + p)}(Q)^*) + \sum_{i=2}^{\ell^P} \delta_E(L_i^P, {}^*Q^*) + \sum_{i=i_1}^{i_2} \delta_E(L_i^T, {}^*Q^*) - \sum_{i=1}^{\ell^P} \mathrm{mk}(r_j + p, L_i^P)$$

$$- \sum_{i=i_1}^{i_2} \mathrm{mk}(r_j + p, L_i^T)$$

$$= \Lambda - \sum_{i=1}^{\ell^P} \sum_{v=1}^{\ell^T} \mathrm{mk}(r_j + p, \hat{\kappa}, L_v^P, L_i^T) - \sum_{i=i_1}^{i_2} \sum_{v=1}^{\ell^P} \mathrm{mk}(r_j + p, \hat{\kappa}, L_v^P, L_i^T)$$

$$\geq \Lambda - 2\mathrm{mk}(r_j + p).$$

This concludes the proof of the claim. ⌟

Next, due to Lemma 6.14, we have $\delta_E(L_1^P, \mathrm{rot}^{-\rho(r_j + p)}(Q)^*) < \partial_P$. By a direct application of Lemma 6.4, we then have that $\min_t \delta_E(P, T[r_j + p \mathinner{\ldotp\ldotp} t)) \leq \Lambda$. If $\Lambda \leq k$, then we are done. For the remainder of the proof we thus consider the case where $\Lambda > k$, which, combined with the fact that $r_j + p$ is a light position of T, means that $\Lambda - 2\mathrm{mk}(r_j + p) > k - 2\mathrm{mk}(r_j + p) \geq k - 2\eta$.

⊐ **Claim 6.25.** For any run M of $2\eta + 30$ consecutive plain pairs $(F_{j,i_1}, G_{j,i_1}), \ldots, (F_{j,i_2}, G_{j,i_2})$ in $\mathcal{I}'_j$, there exists some $i \in (i_1 \mathinner{\ldotp\ldotp} i_2)$ for which $\delta_E(F[f_{j,i} \mathinner{\ldotp\ldotp} f_{j,i+1}), \mathcal{B}(F[f_{j,i} \mathinner{\ldotp\ldotp} f_{j,i+1}))) = 0$.

Proof. We have $|f_{j,i_1+14} - f_{j,i_1}| \geq 14\tau - d_T \geq 14\kappa/2 - d_T \geq 6\kappa = \Delta$. Similarly, $|f_{j,i_2-13} - f_{j,i_2+1}| \geq \Delta$. Thus, $U = F[f_{j,i_1+14} \mathinner{\ldotp\ldotp} f_{j,i_2-13})$ is Δ-contained by M and hence only overlaps pairs in M by Fact 6.18.

Now, each fragment $L_i^F \in \mathcal{L}^F$ is $13\hat{\kappa}$-contained by a sequence of contiguous non-plain pairs and hence only overlaps non-plain pairs. Further, for each fragment $L_i^G \in \mathcal{L}^G$, as shown in the proof of Claim 6.24, $\mathcal{B}^{-1}(L_i^G)$ only overlaps non-plain pairs. Observe that two fragments of F are necessarily disjoint if the sets of pairs that they overlap are disjoint, and hence $[\,f_{j,i_1+14} \mathinner{\ldotp\ldotp} f_{j,i_2-13}\,)$ is disjoint from all elements of $\mathcal{E}$. As the intervals in $\mathcal{E}$ are pairwise disjoint by Claim 6.23, using Claim 6.24, we obtain

$$
\begin{aligned}
k &\geq \delta_E(F, G[\,p \mathinner{\ldotp\ldotp} b\,)) \\
&\geq \delta_E(U, \mathcal{B}(U)) + \sum_{L_i^F \in \mathcal{D}_j'(p)} \delta_E(L_i^F, \mathcal{B}(L_i^F)) + \sum_{L_i^G \in \mathcal{D}_j'(p)} \delta_E(\mathcal{B}^{-1}(L_i^G), L_i^G) \\
&\geq \delta_E(U, \mathcal{B}(U)) + \Lambda - 2\mathrm{mk}(r_j + p) \\
&> \delta_E(U, \mathcal{B}(U)) + k - 2\eta.
\end{aligned}
$$

Hence, we have $\delta_E(U, \mathcal{B}(U)) < 2\eta$. This implies that $\sum_{i=i_1+14}^{i_2-14} \delta_E(F[\,f_{j,i} \mathinner{\ldotp\ldotp} f_{j,i+1}\,), \mathcal{B}(F[\,f_{j,i} \mathinner{\ldotp\ldotp} f_{j,i+1}\,))) < 2\eta$, and hence, since the number of summands in in the left-hand side of the inequality is $i_2 - 14 - (i_1 + 14) + 1 = 2\eta + 30 - 28 > \eta$ and each of these summands is a non-negative integer, the claim follows. $\quad\lrcorner$

We are now ready to conclude the proof of the lemma. Let $\mathcal{M}_1, \ldots, \mathcal{M}_w$ denote the trimmed runs of $2\eta + 30$ consecutive plain pairs in $\mathcal{I}_j'$ such that, for all i, $\mathcal{M}_i$ originated from a run with $\chi(i)$ more plain pairs. For $i \in [\,1 \mathinner{\ldotp\ldotp} w\,]$, let $\mathcal{M}_i = (F_{j,s_i}, G_{j,s_i}) \cdots (F_{j,e_i}, G_{j,e_i})$ let $\psi(i) \in (\,s_i \mathinner{\ldotp\ldotp} e_i\,)$ be such that and $\delta_E(F[\,f_{j,\psi(i)} \mathinner{\ldotp\ldotp} f_{j,\psi(i)+1}\,), \mathcal{B}(F[\,f_{j,\psi(i)} \mathinner{\ldotp\ldotp} f_{j,\psi(i)+1}\,))) = 0$; observe that $\psi(i)$ exists due to Claim 6.25. Let us now show how to construct $\mathcal{C}$ given $\mathcal{B} = (f_v, g_v)_{v=0}^u$. We intuitively achieve this by inserting $\chi(i)$ copies of $Q^\infty[\,0 \mathinner{\ldotp\ldotp} \tau\,)$ after each of $F[\,f_{j,\psi(i)} \mathinner{\ldotp\ldotp} f_{j,\psi(i)+1}\,)$ and $\mathcal{B}(F[\,f_{j,\psi(i)} \mathinner{\ldotp\ldotp} f_{j,\psi(i)+1}\,))$ and aligning them without errors, thus restoring the original length of each trimmed plain run, without changing the cost of the alignment. Initially, set $\mathcal{C} := \mathcal{B}$. Then, for each $\psi(i)$, in decreasing order

- replace each pair (f_v, g_v) of $\mathcal{C}$ that satisfies $f_v \geq f_{j,\psi(i)+1}$ with $(\chi(i) \cdot \tau + f_v, \chi(i) \cdot \tau + g_v)$, and
- insert $(f_{j,\psi(i)+1} + \mu, g_{j,\psi(i)+1} + \mu)_{\mu=0}^{\chi(i) \cdot \tau - 1}$ after $(f_{j,\psi(i)+1} - 1, g_{j,\psi(i)+1} - 1)$. $\quad\blacksquare$

We conclude with the proof of Lemma 6.11 as promised.

Lemma 6.11. $\mathrm{Occ}_k^E(P, T) \cap \mathbb{L} = \left(\bigcup_{j \in J} (r_j + \mathrm{Occ}_k^E(\mathcal{I}_j')) \right) \cap \mathbb{L} = \left(\bigcup_{j \in J} (r_j + \mathrm{Occ}_k^E(F_j, G_j)) \right) \cap \mathbb{L}$.

Proof. ($\subseteq$): This direction is an immediate consequence of Corollary 6.12.

($\supseteq$): This direction is an immediate consequence of Lemma 6.20. $\quad\blacksquare$

6.3 Combining the Partial Results: Faster NewPeriodicMatches

We are ready to prove our headline result—Lemma 1.3—which we restate here for convenience.

Lemma 1.3 (NewPeriodicMatches(P, T, k, d, Q, $\mathcal{A}_P$, $\mathcal{A}_T$)). *We can solve the NEWPERIODIC-MATCHES problem in* $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ *time in the* PILLAR *model.* $\quad\blacksquare$

Proof. Consistently with all previous sections, set

$$
\eta := \max\{1, \lfloor \sqrt{d}/\sqrt{\log(n+1)\log(d+1)} \rfloor\} \leq d \quad \text{and} \quad \kappa := k + d_P + d_T \quad \text{and} \quad \hat{\kappa} := 7d.
$$

First, observe that we may assume $\eta = \Theta(\sqrt{d}/\sqrt{(\log(n+1)\log(d+1))}$—otherwise, in the case where $\sqrt{d} \leq \sqrt{\log(n+1)\log(d+1)}$, PeriodicMatches($P$, T, k, d, Q) from Fact 2.12 already runs in time $\mathcal{O}(d^4) = \mathcal{O}(d^{3.5}\sqrt{\log(n+1)\log(d+1)})$.[20]

We proceed in roughly four steps.

---

[20] All running times in this proof are in the PILLAR model.

- First, in a preprocessing step, we identify the heavy positions $\mathbb{H}$ and the light positions $\mathbb{L}$ in T, as well as a filter $\mathcal{F}$ (according to Lemma 3.4) for where potential k-error occurrences may start.
- Next, we compute all k-error occurrences starting at a position in $\mathbb{H}$, using the `Verify`-based algorithm from Section 6.1. In particular, we obtain $\mathbb{H} \cap \mathrm{Occ}_k^E(P, T)$ as a set of positions.
- Next, we use `TrimmedPM` from Lemma 4.36 to obtain a candidate set $\mathcal{R}$ of potential starting positions of k-error occurrences (represented as $\mathcal{O}(d^3 \eta)$ disjoint arithmetic progressions with difference τ). From Section 6.2, we have $(\mathcal{R} \cap \mathcal{F}) \setminus (\mathbb{H} \cap \mathcal{F}) = \mathrm{Occ}_k^E(P, T) \cap \mathbb{L}$; hence we proceed to compute $\mathbb{H} \cap \mathcal{F}$ (represented as a set), and $\mathcal{R} \cap \mathcal{F}$ (represented as $\mathcal{O}(d^3 \eta)$ disjoint arithmetic progressions with difference τ), and thereafter compute their set difference to obtain $\mathrm{Occ}_k^E(P, T) \cap \mathbb{L}$ (represented as $\mathcal{O}(d^3 \eta)$ disjoint arithmetic progressions with difference τ).
- In a post-processing step, we union the two sets $\mathbb{H} \cap \mathrm{Occ}_k^E(P, T)$ and $\mathrm{Occ}_k^E(P, T) \cap \mathbb{L}$ and compute a representation as $\mathcal{O}(d^3)$ arithmetic progressions with difference q.

Preprocessing. We compute the sets of locked fragments $\mathcal{L}^P = \mathtt{Locked}(P, Q, d_P, \partial_P)$ and $\mathcal{L}^T = \mathtt{Locked}(T, Q, d_T, 0)$ in $\mathcal{O}(d^2)$ time using Lemma 5.4 and call $\mathtt{Heavy}(P, T, d, k, Q, \mathcal{L}^P, \mathcal{L}^T, \hat{\kappa}, \eta)$ to obtain the set $\mathbb{H}$ of heavy positions, represented as the union of $\mathcal{O}(d^2)$ disjoint integer ranges, in $\mathcal{O}(d^2 \log \log d)$ time (cf. Lemma 5.9).

Define J and each of r_j and $\mathcal{I}_j'$, for $j \in J$, as in Sections 3 and 4 and Definition 6.10, and set

$$\mathcal{R} := \bigcup_{j \in J} (r_j + \mathrm{Occ}_k^E(\mathcal{I}_j')).$$

Observe that by Lemma 6.11, $\mathrm{Occ}_k^E(P, T)$ is equal to the union

$$\mathrm{Occ}_k^E(P, T) = (\mathrm{Occ}_k^E(P, T) \cap \mathbb{H}) \cup (\mathrm{Occ}_k^E(P, T) \cap \mathbb{L}) = (\mathrm{Occ}_k^E(P, T) \cap \mathbb{H}) \cup (\mathcal{R} \cap \mathbb{L}).$$

Finally, for our filter, we recall that Lemma 3.4 yields

$$\mathcal{F} := \bigcup_{j \in \mathbb{Z}} [\, jq - x_T - \kappa - d_T \,..\, jq - x_T + \kappa + d_T \,] \supseteq \mathrm{Occ}_k^E(P, T).$$

Observe that we thus have

$$\mathrm{Occ}_k^E(P, T) \cap \mathbb{L} = (\mathcal{R} \cap \mathcal{F}) \setminus (\mathbb{H} \cap \mathcal{F}). \tag{9}$$

Heavy occurrences. By Lemma 6.1, we can compute $\mathrm{Occ}_k^E(P, T) \cap \mathbb{H}$ in time

$$\mathcal{O}(d^3 + d^4/\eta) = \mathcal{O}(d^{3.5}\sqrt{\log n \log d}).$$

Light occurrences. Let us now proceed to computing the remaining k-error occurrences, namely those that start at light positions. To this end, we first compute the $\mathcal{O}(d)$-size sets $\mathsf{Red}(P)$ and $\mathsf{Red}(T)$, defined in the beginning of Section 6.2.2, in $\mathcal{O}(d \log \log d)$ time (cf. Lemma 6.9). Then, we call $\mathtt{TrimmedPM}(T, P, k, Q, \mathcal{A}_P, \mathcal{A}_T, \mathsf{Red}(P), \mathsf{Red}(T), 2\eta + 30)$ from Lemma 4.36, which returns a representation of $\mathcal{R}$ as $\mathcal{O}(d^3 \cdot \eta)$ arithmetic progressions with difference $\tau := q \lceil \kappa/2q \rceil$. Plugging in Theorem 1.5 for the DYNAMICPUZZLEMATCHING data structure, $\mathtt{TrimmedPM}$ runs in time

$$\mathcal{O}(d^3 \cdot \eta \cdot \log n \log d) = \mathcal{O}(d^{3.5}\sqrt{\log n \log d}).$$

We move on to remove the surplus positions from the candidate set $\mathcal{R}$. We start by computing $\mathbb{H} \cap \mathcal{F}$.

☞ **Claim 6.26.** The set $\mathbb{H} \cap \mathcal{F}$ is of size $\mathcal{O}(d^3/\eta)$ and can be computed in time $\mathcal{O}(d^3/\eta)$.

Proof. The proof is similar to a part of the proof of Lemma 6.1.

Due to Lemma 5.9, our representation of $\mathbb{H}$ consists of $\mathcal{O}(d^2)$ disjoint integer ranges. In addition, due to Lemma 5.10, we have

$$|\mathbb{H}| = \mathcal{O}((d^2/\eta + d)(\hat{\kappa} + q)) = \mathcal{O}(d^2/\eta \cdot (d + q)).$$

Let us first upper-bound the size of $|\mathbb{H} \cap \mathcal{F}|$. We distinguish between two cases.

- First, if $q \leq d$, we have $|\mathbb{H} \cap \mathcal{F}| \leq |\mathbb{H}| = \mathcal{O}(d^3/\eta)$.
- As for the complementary case where $q > d$, observe that only $\mathcal{O}(d)$ out of any $\mathcal{O}(q)$ consecutive integers are in $\mathcal{F}$. Hence, $\mathbb{H} \cap \mathcal{F}$ is of size $\mathcal{O}(d^2 + |\mathbb{H}| \cdot d/q) = \mathcal{O}(d^3/\eta)$.

As for computing $\mathbb{H} \cap \mathcal{F}$, we first sort the $\mathcal{O}(d^2)$ integer ranges that comprise $\mathbb{H}$ with respect to their starting positions in $\mathcal{O}(d^2 \log d)$ time and then scan them from left to right, skipping positions in $\mathbb{Z} \setminus \mathcal{F}$. The running time of the scan is proportional to the total number of input ranges and output positions, and hence we are done. □

We move on to compute $\mathcal{R} \cap \mathcal{F}$. In what follows, for any integer x, let us say that the *residue modulo x* of a non-empty arithmetic progression whose difference is a multiple of x is the residue of any element of this arithmetic progression modulo x.

☞ **Claim 6.27.** We can compute a representation of $\mathcal{R} \cap \mathcal{F}$ as $\mathcal{O}(d^3\eta)$ disjoint arithmetic progressions with difference τ, sorted according to their starting positions, in $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ time.

Proof. In a linear scan of the $\mathcal{O}(d^3\eta)$ arithmetic progressions that comprise $\mathcal{R}$ as returned by the call to the algorithm `TrimmedPM`, we delete any arithmetic progression whose elements are not in $\mathcal{F}$. Then, we sort all arithmetic progressions according to their starting positions in time $\mathcal{O}(d^3\eta \log d) = \mathcal{O}(d^{3.5})$ and distribute them among $\mathcal{O}(d/q \cdot \tau) = \mathcal{O}(d^2)$ buckets according to their residues modulo τ. Finally, we scan linearly the arithmetic progressions in each bucket, greedily merging progressions that overlap (that is, we merge progressions if their union is also a valid arithmetic progression with difference τ). The number of the resulting arithmetic progressions is clearly upper-bounded by the size of the input, that is, $\mathcal{O}(d^3\eta)$; we sort them according to their starting positions in $\mathcal{O}(d^{3.5})$ time. □

Finally, we compute the set difference $(\mathcal{R} \cap \mathcal{F}) \setminus (\mathbb{H} \cap \mathcal{F}) = \mathrm{Occ}_k^E(P, T) \cap \mathbb{L}$.

☞ **Claim 6.28.** A representation of $\mathrm{Occ}_k^E(P, T) \cap \mathbb{L}$ as $\mathcal{O}(d^3\eta)$ disjoint arithmetic progressions with difference τ, sorted according to their positions, can be computed in time $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ time the `PILLAR` model.

Proof. We intend to use equation (9), relying on Claims 6.26 and 6.27 to compute $\mathbb{H} \cap \mathcal{F}$ and a representation of $\mathcal{R} \cap \mathcal{F}$ as $\mathcal{O}(d^3\eta)$ disjoint arithmetic progressions with difference τ in $\mathcal{O}(d^{3.5})$ time in total. Then, we process the elements of these two sets in $\mathcal{O}(d/q \cdot \tau) = \mathcal{O}(d^2)$ batches, where each batch contains all elements with a specific residue modulo τ. For such a fixed residue, we scan in parallel the arithmetic progressions in $\mathcal{R} \cap \mathcal{F}$ and the positions in $\mathbb{H} \cap \mathcal{F}$, both of which are sorted in increasing order. When some element of $\mathbb{H} \cap \mathcal{F}$ is contained in some arithmetic progression in $\mathcal{R} \cap \mathcal{F}$, this arithmetic progression is split into two parts, either of which may be empty. The number of the resulting arithmetic progressions is upper-bounded by the total size of the input, that is, $\mathcal{O}(d^3\eta)$; we sort them according to their starting positions in $\mathcal{O}(d^{3.5})$ time. □

Post-processing. Recall that the set $\mathrm{Occ}_k^E(P,T) \cap \mathbb{H}$ is computed explicitly using Lemma 6.1 in $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ time; its size is $\mathcal{O}(d^3/\eta)$ due to Lemma 3.4 and Claim 6.26. In addition, due to Claim 6.28, a representation of $\mathrm{Occ}_k^E(P,T) \cap \mathbb{L}$ as $\mathcal{O}(d^3\eta)$ disjoint arithmetic progressions with difference τ can be computed in $\mathcal{O}(d^{3.5}\sqrt{\log n \log d})$ time. Overall, by taking the union of the disjoint sets $\mathrm{Occ}_k^E(P,T) \cap \mathbb{H}$ and $\mathrm{Occ}_k^E(P,T) \cap \mathbb{L}$, we obtain a representation of $\mathrm{Occ}_k^E(P,T)$ as $\mathcal{O}(d^3\eta)$ disjoint arithmetic progressions with difference τ. In what follows, we show how to efficiently replace these arithmetic progressions with difference τ with arithmetic progressions with difference q. This allows us to benefit from the fact that $\mathrm{Occ}_k^E(P,T)$ can be decomposed to $\mathcal{O}(d^3)$ arithmetic progressions with difference q (see [CKW20, Main Theorem 7]) to decrease the size of the output.

First, we repeatedly merge any two arithmetic progressions that contain elements that are τ positions apart. Then, we process each of the $\mathcal{O}(d)$ relevant residues modulo q separately. Let us fix such a residue r. We maintain arrays B_r and S_r, each of size τ/q, during a left-to-right scan of the text. When some position μ of T is processed, for $i \in [0 \mathinner{.\,.} \tau/q)$,

- $B_r[i]$ stores a boolean variable indicating whether the successor of μ in $\mathbb{Z}$ with residue $iq+r$ modulo τ is in $\mathrm{Occ}_k^E(P,T)$.
- $S_r[i]$ stores the successor of μ in $\mathbb{Z}$ with residue $iq+r$ modulo τ that is *not* in $\mathrm{Occ}_k^E(P,T)$ if $B_r[i]=1$ and ∞ otherwise.

We maintain balanced binary trees over arrays B_r and S_r so that we can efficiently query for the leftmost non-zero element in any subarray of B_r and the minimum element of S_r, with an $\mathcal{O}(\log d)$-time additive overhead per update of B_r and S_r. All updates to B_r and S_r can be stored in a priority queue after a linear-time preprocessing of the arithmetic progressions in scope, prioritized by the value of μ that triggers them.

In our scan of T, we maintain the minimum position $\mu \in \mathrm{Occ}_k^E(P,T)$ of residue r modulo q that we have not reported so far. Given such a position μ, we query in $\mathcal{O}(\log d)$ time for the successor ν of μ in $\mathbb{Z}$ with residue $iq+r$ that is *not* in $\mathrm{Occ}_k^E(P,T)$:

- If there is any false entry in $B_r(\mu - r \mod q \mathinner{.\,.} \tau/q)$, then the sought integer can be retrieved in $\mathcal{O}(1)$ time given the leftmost such entry;
- else, if there is any false entry in $B_r[0 \mathinner{.\,.} \mu - r \mod q)$, then the sought integer can be retrieved in $\mathcal{O}(1)$ time given the leftmost such entry;
- else, all entries of B_r are set to true and hence the sought integer corresponds to the smallest integer stored in S_r.

Given ν, we report the arithmetic progression $\{\mu + iq : i \in [0 \mathinner{.\,.} (\nu - \mu)/q)]\}$. Then, we implicitly continue our scan of T by performing the necessary (precomputed) updates in B_r and S_r until we reach either the successor of ν in $\mathbb{Z}$ with residue r modulo q that is in $\mathrm{Occ}_k^E(P,T)$ or the end of the text. The former condition can be checked using array B_r in $\mathcal{O}(\log d)$ time prior to each update to either B_r or S_r. Then, we set said position as μ and repeat the above process.

Over all residues, the total time taken is $\mathcal{O}(d^2 + d^3\eta \log d) = \mathcal{O}(d^{3.5})$: we pay $\mathcal{O}(d\tau/q) = \mathcal{O}(d^2)$ time to initialize all arrays, while all other operations take total time proportional to the product of $\log d$ with the total size of the input and the output arithmetic progressions. We can benefit from the upper bound on the number of arithmetic progressions to which $\mathrm{Occ}_k^E(P,T)$ can be decomposed due to the greedy nature of the algorithm that computes them. This concludes the proof of this lemma. ∎

7 Faster Approximate Pattern Matching in Important Settings

In this section, we rely on known implementations (see [CKW20]) of the PILLAR model in the static, dynamic, and fully compressed settings, thereby lifting Main Theorem 2 to these settings.

7.1 An Algorithm for the Standard Setting

In the standard setting, we implement a handle to $S = X[\ell .. r)$ as a pointer to $X \in \mathcal{X}$ (which is stored explicitly) along with the indices ℓ and r. As was argued in detail in [CKW20], for the implementation of the PILLAR model in the standard setting it suffices to combine well-known results. Namely,

- Extract, Access, and Length admit trivial implementations;
- LCP queries can be efficiently implemented by constructing a generalized suffix tree for the strings in the collection [Far97] and preprocessing it for $\mathcal{O}(1)$-time lowest common ancestor queries [BF00], while for LCP^R we can use an analogous construction over the reverse strings of the strings in the collection;
- IPM queries can be answered in $\mathcal{O}(1)$ time by the linear-size data structure of Kociumaka et al. [KRRW15, Koc18].

The above discussion is summarized in the following statement.

Theorem 7.1 (see [CKW20, 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.*

Combining Theorem 7.1 and Main Theorem 2, we obtain an algorithm for pattern matching with edits that is faster than the algorithm of Cole and Hariharan [CH02], unless $k/\log(k+1) = \mathcal{O}(\log m)$ or both algorithms run in $\mathcal{O}(n)$ time.

Main Theorem 1. *Given a text T of length n, a pattern P of length m, and an integer threshold $k > 0$, we can compute the set $\text{Occ}_k^E(P, T)$ in $\mathcal{O}(n + n/m \cdot k^{3.5}\sqrt{\log m \log k})$ time.*

Observe that the algorithm encapsulated in Theorem 7.1 is faster than the classical $\mathcal{O}(nk)$-time algorithm of Landau and Vishkin [LV89] when $k = \omega(1)$ and $k^{2.5}\sqrt{\log k} = o(m/\sqrt{\log(m+1)})$.

Remark 7.2. Our algorithm also applies 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 $\text{Occ}_k^E(P, T)$ in time $\mathcal{O}(|T|/|P| \cdot k^{3.5}\sqrt{\log m \log k})$.

7.2 An Algorithm for the Dynamic Setting

Next, 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. [GKK+18] presented a data structure that efficiently maintains such a collection and allows for efficient longest common prefix queries. In [CKW20], it was (a) argued in detail that that the aforementioned data structure readily supports all PILLAR operations other than IPM queries, and (b) shown that it can be augmented to efficiently answer IPM queries; see also [Cha20, Section 4] for a more direct proof of the latter claim. The above discussion is formalized in the following statement.

Theorem 7.3 ([GKK+18, CKW20]). *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 $\mathcal{O}(\log^2 N)$.*[21]

---

[21] All running time bounds hold w.h.p.

A very recent work [KK22, Section 8] provides an alternative deterministic implementation of dynamic strings, supporting operations $\texttt{makestring}(U)$, $\texttt{concat}(U, V)$, $\texttt{split}(U, i)$ in $\mathcal{O}(|U| \log^{\mathcal{O}(1)} \log N)$, $\mathcal{O}(\log |UV| \log^{\mathcal{O}(1)} \log N)$, and $\mathcal{O}(\log |U| \log^{\mathcal{O}(1)} \log N)$, respectively, so that PILLAR operations can be performed in time $\mathcal{O}(\log N \log^{\mathcal{O}(1)} \log N)$ time.

Combining Theorem 7.3 and Main Theorem 2, we obtain the following result for approximate pattern matching under the edit distance for dynamic strings.

Main Theorem 3. *We can maintain a collection $\mathcal{X}$ of non-empty persistent strings of total length N subject to $\texttt{makestring}(U)$, $\texttt{concat}(U, V)$, and $\texttt{split}(U, i)$ operations that require $\mathcal{O}(\log N + |U|)$, $\mathcal{O}(\log N)$, and $\mathcal{O}(\log N)$ time, respectively, so that given two strings $P, T \in \mathcal{X}$, and an integer threshold $k > 0$, we can compute a representation of $\mathrm{Occ}_k^E(P, T)$ as $\mathcal{O}(|T|/|P| \cdot k^3)$ arithmetic progressions with the same difference in time $\mathcal{O}(|T|/|P| \cdot k^{3.5} \sqrt{\log |P| \log k} \log^2 N)$.[22]*

The result encapsulated in Main Theorem 3 should be compared to [CKW20, Main Theorem 3], which has the same complexity guarantees for processing updates, but answers approximate pattern matching queries under the edit distance in $\mathcal{O}(|T|/|P| \cdot k^4 \log^2 N)$ time w.h.p.

7.3 An Algorithm for 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.

We write $N_{\mathcal{G}}$ for the set of non-terminals of a context-free grammar $\mathcal{G}$ and call the elements of $\mathcal{S}_G \coloneqq N_{\mathcal{G}} \cup \Sigma$ *symbols*. Then, a *straight line program (SLP)* $\mathcal{G}$ is a context-free grammar that consists of a set $N_{\mathcal{G}} = \{A_1, \ldots, A_n\}$ of non-terminals, such that each $A_i \in N_{\mathcal{G}}$ is associated with a unique production rule $A_i \to f_{\mathcal{G}}(A_i) \in (\Sigma \cup \{A_j : j < i\})^*$. For SLPs given as input, we can assume without loss of generality that each production rule is of the form $A \to BC$ 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 $\mathsf{gen}(A) \in \Sigma^*$. The string $\mathsf{gen}(A)$ can be obtained from A by repeatedly replacing each non-terminal by its production. In addition, A is associated with its *parse tree* $\mathsf{PT}(A)$ 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 $A \to BC$, then $\mathsf{PT}(B)$ and $\mathsf{PT}(C)$ are attached (in this order).

Observe that if we traverse the leaves of $\mathsf{PT}(A)$ from left to right, spelling out the corresponding non-terminals, then we obtain $\mathsf{gen}(A)$. We say that $\mathcal{G}$ generates $\mathsf{gen}(\mathcal{G}) \coloneqq \mathsf{gen}(A_n)$. The parse tree $\mathsf{PT}_{\mathcal{G}}$ of $\mathcal{G}$ is then the parse tree of the starting symbol $A_n \in N_{\mathcal{G}}$.

First, as also observed in [CKW20], 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 $\mathcal{O}(n)$ $\texttt{makestring}(a)$ operations, for $a \in \Sigma$, and $\mathcal{O}(n)$ $\texttt{concat}$ 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 7.3, thus outperforming Theorem 7.4. Next, we outline a deterministic implementation of PILLAR operations in the fully compressed setting.

Following [CKW20], the handle of a fragment $S = X[\ell \mathinner{.\,.} r)$ consists of a pointer to the SLP $\mathcal{G} \in \mathcal{X}$ generating X along with the positions ℓ and r. This makes operation $\texttt{Extract}$ trivial. As argued in [CKW20], all remaining PILLAR operations admit efficient implementations in the considered setting.

---

[22] All running time bounds hold with high probability (that is, $1 - 1/N^{\Omega(1)}$). A deterministic version can be obtained at the cost of a $\mathrm{poly}(\log \log N)$-factor overhead.

- For operation `Length`, we precompute $|\text{gen}(A)|$ for each non-terminal A using dynamic programming.
- For operation `Access`, we use the data structure of Bille et al. [BLR$^+$15].
- For operations `LCP` and `LCP`^{R}, we use the data structure of I [I17] that is based on the recompression technique, which is due to [Jeż15, Jeż16].
- For operation `IPM`, we use a data structure presented in [CKW20, KK20] that is also based on the recompression technique.

The above discussion is summarized in the following statement.

⬚ **Theorem 7.4** (see [BLR$^+$15, I17, CKW20, KK20]). *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.* ⬛

We are now ready to present an efficient algorithm for approximate pattern matching under edit distance in the fully compressed setting. We choose to state our results using the deterministic implementation of the `PILLAR` model in this setting, that is, Theorem 7.4.

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 $P := \text{gen}(\mathcal{G}_T)$, and a threshold k and are required to compute the k-error occurrences of P in T.

Set $N := |T|$, $M := |P|$, 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 7.4. 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. We combine Theorem 7.4 with Main Theorem 2 to compute such occurrences. Finally, we combine the computed occurrences using dynamic programming.

Formally, for each non-terminal $A \in N_{\mathcal{G}_T}$, with production rule $A \to BC$, let

- $A_\ell := \text{gen}(B)[\max\{0, |\text{gen}(B)| - M - k + 1\} \mathbin{..} |\text{gen}(B)|)$,
- $A_r := \text{gen}(C)[0 \mathbin{..} \min\{M + k, |\text{gen}(C)|\})$, and
- $\text{cross}(A) := (\text{Occ}_k^E(P, A_\ell A_r) \cap [0 \mathbin{..} |A_\ell|)) \setminus \text{Occ}_k^E(P, A_\ell)$,

observing that $\text{Occ}_k^E(P, \text{gen}(A))$ can then be partitioned to

- $\text{Occ}_k^E(P, \text{gen}(B))$,
- $|\text{gen}(B)| - |\text{gen}(A_\ell)| + \text{cross}(A)$, and
- $|\text{gen}(B)| + \text{Occ}_k^E(P, \text{gen}(C))$.

Next, observe that, by combining Theorem 7.4 and Main Theorem 2, $\text{cross}(A)$ can be computed in time $\mathcal{O}(k^{3.5} \sqrt{\log M \log k} \log^2 N \log \log N)$ since $A_\ell A_r$ and A_ℓ are fragments of $\text{gen}(\mathcal{G}_T)$ of length at most $2(M + k - 1)$. Now, the size of $|\text{Occ}_k^E(P, T)|$ can be computed by a straightforward dynamic programming approach: for a non-terminal $A \in N_{\mathcal{G}_T}$, with production rule $A \to BC$, the number of k-errors occurrences of P in $\text{gen}(A)$ equals $|\text{Occ}_k^E(P, \text{gen}(B))| + |\text{cross}(A)| + |\text{Occ}_k^E(P, \text{gen}(C))|$. 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 k-error occurrences.[23]

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

⬚ **Main Theorem 4**. *Let $\mathcal{G}_T$ denote a straight-line program of size n generating a string T, let $\mathcal{G}_P$ denote a straight-line program of size m generating a string P, let $k > 0$ denote an integer threshold, and set $N := |T|$ and $M := |P|$. We can compute $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(m \log N + n \, k^{3.5} \log^2 N \sqrt{\log M \log k} \log \log N)$ and we can report the elements of $\text{Occ}_k^E(P, T)$ within $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time.* ⬛

The result encapsulated in Main Theorem 4 should be compared to [CKW20, Main Theorem 2], which summarizes an algorithm that computes $|\text{Occ}_k^E(P, T)|$ in time $\mathcal{O}(m \log N + n \, k^4 \log^2 N \log \log N)$ and can then report all elements of $\text{Occ}_k^E(P, T)$ within $\mathcal{O}(|\text{Occ}_k^E(P, T)|)$ extra time.

---

[23] Compare [CKW20, Main Theorem 2] for a similar algorithm.

Part II: Seaweeds

8 The Seaweed Monoid of Permutation Matrices

We start by introducing the terminology behind the definition of permutation matrices and their seaweed products. We mostly follow [KS21, Section 3.2] and [Tis07, Chapters 2 and 3], except that we use matrices whose rows and columns are indexed with non-empty integer intervals (finite or infinite). Consistently with previous work, we think of the plane reflected along the horizontal axis; that is, $(1,1)$ is below and to the right of $(0,0)$. For an interval $I \subseteq \mathbb{Z}$, denote

$$I^+ = \{i \in \mathbb{Z} : i \in I \text{ or } i - 1 \in I\}, \qquad {}^+I = \{i \in \mathbb{Z} : i \in I \text{ or } i + 1 \in I\},$$
$$I^- = \{i \in \mathbb{Z} : i \in I \text{ and } i + 1 \in I\}, \qquad {}^-I = \{i \in \mathbb{Z} : i \in I \text{ and } i - 1 \in I\}.$$

Moreover, for $S \subseteq \mathbb{Z}$, we define $\mathsf{span}(S) = \{i \in \mathbb{Z} : \exists_{s,s' \in S} s \leq i \leq s'\}$ as the smallest interval containing S.

The $(\min, +)$ product of matrices $A \in \mathbb{Z}_{\geq 0}^{I \times K}$ and $B \in \mathbb{Z}_{\geq 0}^{K \times J}$ is a matrix $A \odot B \in \mathbb{Z}_{\geq 0}^{I \times J}$ with entries defined as follows for $i \in I$ and $j \in J$:

$$(A \odot B)[i,j] = \min_{k \in K}\{A[i,k] + B[k,j]\}.$$

For $A \in \mathbb{Z}^{I \times J}$, the *density matrix* $A^\square \in \mathbb{Z}^{I^- \times J^-}$ has entries defined as follows for $i \in I^-$ and $j \in J^-$:

$$A^\square[i,j] = A[i+1,j] + A[i,j+1] - A[i,j] - A[i+1,j+1].$$

A matrix $A \in \mathbb{Z}^{I \times J}$ is a *Monge matrix* if all entries of the density matrix $A^\square$ are non-negative.

A matrix $A \in \{0,1\}^{I \times I}$ is a *permutation matrix* if each row and each column contains exactly one entry equal to 1. We define $\mathsf{span}(A) = {}^-\mathsf{span}(\{i \in I : A[i,i] = 0\})$, and we say that the matrix is *bounded* if $\mathsf{span}(A)$ is finite. Note that a permutation matrix A can be represented with a permutation $\sigma : {}^+\mathsf{span}(A) \to {}^+\mathsf{span}(A)$ such that $A[i,j] = 1$ if and only if $j = \sigma(i)$ (when $i \in {}^+\mathsf{span}(A)$) or $j = i$ (otherwise).

For a bounded permutation matrix $A \in \{0,1\}^{I \times I}$, the *distribution matrix* $A^\Sigma \in \mathbb{Z}_{\geq 0}^{I^+ \times I^+}$, has its entries defined as follows for $i, j \in I^+$ (the entries are finite because A is bounded):

$$A^\Sigma[i,j] = \sum_{i' \geq i} \sum_{j' < j} A[i',j'].$$

By constructing the data structure of Chan and Pătrașcu [CP10] for answering two-sided orthogonal range counting queries in 2D over the non-zero entries of A, one obtains the following.

Fact 8.1 ([Tis07, CKM20, KS21]). *After $\mathcal{O}(w\sqrt{\log w})$-time preprocessing of a permutation representing a permutation matrix A with $w = |\mathsf{span}(A)|$, any entry of A^Σ can be computed in $\mathcal{O}(\log w / \log \log w)$ time.*

The *seaweed product* of bounded permutation matrices $A, B \in \{0,1\}^{I \times I}$ is defined as $A \boxdot B := (A^\Sigma \odot B^\Sigma)^\square$.

Theorem 8.2 (Tiskin [Tis07, Tis15]). *For any two bounded permutation matrices A, B, the seaweed product $C := A \boxdot B$ is a permutation matrix with $\mathsf{span}(C) \subseteq \mathsf{span}(\mathsf{span}(A) \cup \mathsf{span}(B))$. Moreover, given the permutations representing A and B, the permutation representing C can be constructed in $\mathcal{O}(w \log w)$ time, where $w = |\mathsf{span}(\mathsf{span}(A) \cup \mathsf{span}(B))|$.*

For a matrix $A \in \mathbb{Z}^{I \times J}$ and an integer $s \in \mathbb{Z}$, we define the (diagonal) *shift* of A by s units as a matrix $A \searrow s \in \mathbb{Z}^{(I+s) \times (J+s)}$ such that $(A \searrow s)[i+s,j+s] = A[i,j]$ for $(i,j) \in I \times J$.

Fact 8.3. *The $\searrow$ operation distributes over the seaweed product, that is, $(A \searrow s) \boxdot (B \searrow s) = (A \boxdot B) \searrow s$ holds for all bounded permutation matrices A, B and integers $s \in \mathbb{Z}$.*

8.1 Alignment Graphs and Distance Matrices

■ **Definition 8.4.** *Given a set $M \subseteq \mathbb{Z}^2$, we define the* alignment graph $\mathsf{AG}(M)$ *with vertices $\mathbb{Z}^2$ and weighted edges:*

- $(x, y) \overset{1}{\longleftrightarrow} (x + 1, y)$ *for every $(x, y) \in \mathbb{Z}^2$ (horizontal edges),*
- $(x, y) \overset{1}{\longleftrightarrow} (x, y + 1)$ *for every $(x, y) \in \mathbb{Z}^2$ (vertical edges),*
- $(x, y) \overset{0}{\longleftrightarrow} (x + 1, y + 1)$ *for every $(x, y) \in \mathbb{Z}^2 \setminus M$ (diagonal edges).*

We denote the underlying distance function on $\mathbb{Z}^2$ by dist_M. ◾

We introduce an order $\prec$ on $\mathbb{Z}^2$ so that $(x, y) \prec (x', y')$ if and only if $x < x'$ and $y < y'$. For $S \subseteq \mathbb{Z}^2$, we denote by $\mathsf{LIS}(S)$ the maximum length of an increasing sequence (with respect to $\prec$) of points in S.

■ **Lemma 8.5.** *Consider a set $M \subseteq \mathbb{Z}^2$ and two points $p = (x, y), p' = (x', y')$ in $\mathbb{Z}^2$.*

- *If $x \leq x'$ and $y \leq y'$, then $\mathsf{dist}_M(p, p') = |x' - x| + |y' - y| - 2\mathsf{LIS}(([\,x \mathinner{.\,.} x'\,) \times [\,y \mathinner{.\,.} y'\,)) \setminus M)$.*
- *If $x \leq x'$ and $y \geq y'$, then $\mathsf{dist}_M(p, p') = |x' - x| + |y' - y|$.*
- *If $x \geq x'$ and $y \leq y'$, then $\mathsf{dist}_M(p, p') = |x' - x| + |y' - y|$.*
- *If $x \geq x'$ and $y \geq y'$, then $\mathsf{dist}_M(p, p') = |x' - x| + |y' - y| - 2\mathsf{LIS}(([\,x' \mathinner{.\,.} x\,) \times [\,y' \mathinner{.\,.} y\,)) \setminus M)$.*

Proof. Let us fix a shortest path between p and p' in $\mathsf{AG}(M)$ and a box $B := [\,i \mathinner{.\,.} i'\,] \times [\,j \mathinner{.\,.} j'\,]$ containing all vertices of this path. Observe that $\mathsf{dist}_M(p, p') = \mathsf{dist}_{M,B}(p, p')$, where $\mathsf{dist}_{M,B}$ denotes the distance function of the subgraph of $\mathsf{AG}(M)$ induced by B. By [KS21, Lemma 12], $\mathsf{dist}_{M,B}(p, p')$ satisfies the claimed formula in all four cases. ◾

■ **Definition 8.6.** *For a box $B = [\,i \mathinner{.\,.} i'\,] \times [\,j \mathinner{.\,.} j'\,]$, we define the* left-top *and* bottom-right *boundaries $(\mathsf{lt}_d^B)_{d \in [\,i - j' \mathinner{.\,.} i' - j\,]}, (\mathsf{br}_d^B)_{d \in [\,i - j' \mathinner{.\,.} i' - j\,]}$ so that*

$$\mathsf{lt}_d^B = \begin{cases} (i, i - d) & for\ d \in [\,i - j' \mathinner{.\,.} i - j\,], \\ (d + j, j) & for\ d \in [\,i - j \mathinner{.\,.} i' - j\,], \end{cases} \qquad \mathsf{br}_d^B = \begin{cases} (d + j', j') & for\ d \in [\,i - j' \mathinner{.\,.} i' - j'\,], \\ (i', i' - d) & for\ d \in [\,i' - j' \mathinner{.\,.} i' - j\,]. \end{cases}$$
◾

■ **Definition 8.7.** *For a finite set $M \subseteq \mathbb{Z}^2$, we say that $B = [\,i \mathinner{.\,.} i'\,] \times [\,j \mathinner{.\,.} j'\,]$ is a* bounding box *of M if $M \subseteq [\,i \mathinner{.\,.} i'\,) \times [\,j \mathinner{.\,.} j'\,)$. We then define an infinite matrix $D_{M,B} \in \mathbb{Z}^{\mathbb{Z} \times \mathbb{Z}}$ as follows:*

$$D_{M,B}[a, b] := \begin{cases} \mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{br}_b^B) & if\ a, b \in [\,i - j' \mathinner{.\,.} i' - j\,], \\ |a - b| & otherwise. \end{cases}$$
◾

■ **Lemma 8.8.** *For every finite set $M \subseteq \mathbb{Z}^2$, the distance matrix $D_{M,B}$ does not depend on the choice of the bounding box B.*

Proof. Consider an arbitrary bounding box $B = [\,i \mathinner{.\,.} i'\,] \times [\,j \mathinner{.\,.} j'\,]$ and the minimum bounding box $\hat{B} = [\,\hat{\imath} \mathinner{.\,.} \hat{\imath}'\,] \times [\,\hat{\jmath} \mathinner{.\,.} \hat{\jmath}'\,]$; if $M = \varnothing$, set $\hat{\imath} = \hat{\imath}' = i$ and $\hat{\jmath} = \hat{\jmath}' = j$. We prove $D_{M,B}[a, b] = D_{M,\hat{B}}[a, b]$ by analyzing several cases:

- $\boldsymbol{a, b \in [\,\hat{\imath} - \hat{\jmath}' \mathinner{.\,.} \hat{\imath}' - \hat{\jmath}\,]}.$ Lemma 8.5 implies $\mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{lt}_a^{\hat{B}}) = 0 = \mathsf{dist}_M(\mathsf{br}_b^{\hat{B}}, \mathsf{br}_b^B)$. Thus,

$$D_{M,B}[a, b] = \mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{br}_b^B) = \mathsf{dist}_M(\mathsf{lt}_a^{\hat{B}}, \mathsf{br}_b^{\hat{B}}) = D_{M,\hat{B}}[a, b].$$

- $\boldsymbol{a, b \in [\,i - j' \mathinner{.\,.} i' - j\,]}$ **and** $\boldsymbol{a \notin [\,\hat{\imath} - \hat{\jmath}' \mathinner{.\,.} \hat{\imath}' - \hat{\jmath}\,]}.$ Lemma 8.5 implies $\mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{br}_a^B) = 0$ and $\mathsf{dist}_M(\mathsf{br}_a^B, \mathsf{br}_b^B) = |a - b|$. Thus,

$$D_{M,B}[a, b] = \mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{br}_b^B) = \mathsf{dist}_M(\mathsf{br}_a^B, \mathsf{br}_b^B) = |a - b| = D_{M,\hat{B}}[a, b].$$

- $a, b \in [\, i - j' \,..\, i' - j \,]$ **and** $b \notin [\, \hat{\imath} - \hat{\jmath}' \,..\, \hat{\imath}' - \hat{\jmath} \,]$. Lemma 8.5 implies $\mathsf{dist}_M(\mathsf{lt}_b^B, \mathsf{br}_b^B) = 0$ and $\mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{lt}_b^B) = |a - b|$. Thus,

$$D_{M,B}[a, b] = \mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{br}_b^B) = \mathsf{dist}_M(\mathsf{lt}_a^B, \mathsf{lt}_b^B) = |a - b| = D_{M,\hat{B}}[a, b].$$

- **Otherwise,** $D_{M,B}[a, b] = |a - b| = D_{M,\hat{B}}[a, b].$ ◾

Definition 8.9. *For a finite set $M \subseteq \mathbb{Z}^2$, we define the* distance matrix $D_M \coloneqq D_{M,B}$ *(for an arbitrary bounding box B of M) and the* seaweed matrix $P_M \coloneqq {}^{1\!}/_{2}\, D_M^{\square}$. ◾

Fact 8.10. *For every finite set $M \subseteq \mathbb{Z}^2$ and vector $(u, v) \in \mathbb{Z}^2$, we have $P_{M+(u,v)} = P_M \searrow (u - v)$, where $M + (u, v) = \{(x + u, y + v) : (x, y) \in M\}$.* ◾

Lemma 8.11. *For every finite set $M \subseteq \mathbb{Z}^2$, the seaweed matrix P_M is a bounded permutation matrix. Moreover, $D_M[a, b] = 2P_M^{\Sigma}[a, b] + a - b$ holds for all $a, b \in \mathbb{Z}^2$.*

Proof. By [KS21, Lemma 19], for every bounding box $B = [\, i \,..\, i' \,] \times [\, j \,..\, j' \,]$, the matrix P_M restricted to entries $P_M[a, b]$ with $a, b \in [\, i - j' \,..\, i' - j \,)$ is a permutation matrix. Since the bounding box can be chosen arbitrarily large, the entire matrix P_M is a bounded permutation matrix. The second claim also follows from [KS21, Definition 18 and Lemma 19]. ◾

Definition 8.12. *We say that two boxes $B = [\, i \,..\, i' \,] \times [\, j \,..\, j' \,]$ and $\hat{B} = [\, \hat{\imath} \,..\, \hat{\imath}' \,] \times [\, \hat{\jmath} \,..\, \hat{\jmath}' \,]$ are*
- vertically adjacent *if $i = \hat{\imath}$, $i' = \hat{\imath}'$, and $j' = \hat{\jmath}$,*
- horizontally adjacent *if $i' = \hat{\imath}$, $j = \hat{\jmath}$, and $j' = \hat{\jmath}'$.* ◾

Lemma 8.13. *Consider finite sets $M, \hat{M} \subseteq \mathbb{Z}^2$ with (horizontally or vertically) adjacent bounding boxes B and $\hat{B}$, respectively. Then, $P_{M \cup \hat{M}} = P_M \boxdot P_{\hat{M}}$.*

Proof. It suffices to prove that $D_{M \cup \hat{M}} = D_M \odot D_{\hat{M}}$. By symmetry, we assume without loss of generality that the boxes $B, \hat{B}$ are vertically adjacent, that is, $B = [\, i \,..\, i' \,] \times [\, j \,..\, j' \,]$ and $\hat{B} = [\, i \,..\, i' \,] \times [\, j' \,..\, j'' \,]$ for some integers $i \leq i'$ and $j \leq j' \leq j''$. By Lemma 8.8, for every $a, c \in \mathbb{Z}$, we have $D_{M \cup \hat{M}}[a, c] = \mathsf{dist}_{M \cup \hat{M}}((a + j, j), (c + j'', j''))$, since

$$[\, \min\{i, a + j, c + j''\} \,..\, \max\{i', a + j, c + j''\} \,] \times [\, j \,..\, j'' \,]$$

is a bounding box for $M \cup \hat{M}$. In addition, since each path from $(a + j, j)$ to $(c + j'', j'')$ passes through a vertex of the form $(b + j', j')$ for some $b \in \mathbb{Z}$, we also have

$$\mathsf{dist}_{M \cup \hat{M}}((a + j, j), (c + j'', j'')) = \min_{b \in \mathbb{Z}}\{\mathsf{dist}_{M \cup \hat{M}}((a + j, j), (b + j', j')) + \mathsf{dist}_{M \cup \hat{M}}((b + j', j'), (c + j'', j''))\}.$$

Now, for every $a, b, c \in \mathbb{Z}$, Lemmas 8.5 and 8.8 yield

$$D_M[a, b] = \mathsf{dist}_M((a + j, j), (b + j', j')) = \mathsf{dist}_{M \cup \hat{M}}((a + j, j), (b + j', j')),$$

since $[\, a + j \,..\, b + j' \,) \times [\, j \,..\, j' \,)$ is disjoint from $\hat{M}$, and

$$D_{\hat{M}}[b, c] = \mathsf{dist}_{\hat{M}}((b + j', j'), (c + j'', j'')) = \mathsf{dist}_{M \cup \hat{M}}((b + j', j'), (c + j'', j'')),$$

since $[\, b + j' \,..\, c + j'' \,) \times [\, j' \,..\, j'' \,)$ is disjoint from M.

Thus, $D_{M \cup \hat{M}}[a, c] = \min_{b \in \mathbb{Z}}\{D_M[a, b] + D_{\hat{M}}[b, c]\}$ holds as claimed. ◾

Lemma 8.14. *Every bounded permutation matrix can be represented as P_M for some finite set $M \subseteq \mathbb{Z}^2$.*

Proof. Tiskin [Tis07] showed that every bounded permutation matrix can be represented as the seaweed product of transposition matrices, which can be defined as the unique permutation matrices G_d : $\{0,1\}^{\mathbb{Z} \times \mathbb{Z}}$ satisfying $\mathsf{span}(G_d) = \{d\}$ for $d \in \mathbb{Z}$. By Definitions 8.7 and 8.9, for any $y, d \in \mathbb{Z}_{\geq 0}$, we have $G_d = P_{\{(y+d,y)\}}$. By Lemma 8.13, any product $G_{d_1} \boxdot \cdots \boxdot G_{d_t}$ can be represented as P_M for $M = \{(y + d_y, y) : y \in [\,1 \mathinner{.\,.} t\,]\}$. ◾

For a set $M \subseteq \mathbb{Z}^2$, we define $\mathsf{span}(M) = \mathsf{span}(\{x - y : (x, y) \in M\})$.

Lemma 8.15. *For every finite set $M \subseteq \mathbb{Z}^2$, we have $\mathsf{span}(P_M) \subseteq \mathsf{span}(M)$.*

Proof. Let $a \in \mathbb{Z} \setminus {}^+\mathsf{span}(M)$. This means that $a, a+1 \notin \mathsf{span}(M)$ and thus $a, a+1 \notin \{x - y : (x, y) \in M\}$. By Lemma 8.5, this implies $D_M[a, a] = D_M[a+1, a+1] = 0$ and $D_M[a, a+1] = D_M[a+1, a] = 1$. Consequently,

$$P_M[a, a] = \tfrac{1}{2}(D_M[a, a+1] + D_M[a+1, a] - D_M[a, a] - D_M[a+1, a+1]) = 1.$$

Thus, $\{a \in \mathbb{Z} : P_M[a, a] = 0\} \subseteq {}^+\mathsf{span}(M)$. Given that $\mathsf{span}(M)$ is an interval, this yields $\mathsf{span}(P_M) = {}^-\mathsf{span}(\{a \in \mathbb{Z} : P_M[a, a] = 0\}) \subseteq \mathsf{span}(M)$. ◾

8.2 Restriction of Permutation Matrices

For a finite set $M \subseteq \mathbb{Z}^2$ and an integer interval I, we denote $M|_I := \{(x, y) \in M : x - y \in I\}$; observe that $\mathsf{span}(M|_I) = \mathsf{span}(\{x - y : (x, y) \in M|_I\}) \subseteq \mathsf{span}(I) = I$.

Lemma 8.16. *Consider a finite set $M \subseteq \mathbb{Z}^2$ and an integer interval $I = (\,\ell \mathinner{.\,.} r\,)$. The restricted seaweed matrix $P_{M|_I}$ can be characterized as follows based on the sets*

$$L = \{(a, b) \in [\,\ell \mathinner{.\,.} r\,]^2 : a + b \leq \ell + r \text{ and } D_M[a, b] \geq a + b - 2\ell\},$$
$$R = \{(a, b) \in [\,\ell \mathinner{.\,.} r\,]^2 : a + b \geq \ell + r \text{ and } D_M[a, b] \geq 2r - a - b\}.$$

For every $a, b \in {}^+I$, we have $P_{M|_I}[a, b] = 1$ if and only if at least one of the following cases holds:
(1) $(a + 1, b), (a, b + 1) \in L$ *and* $(a + 1, b + 1) \notin L$*, or*
(2) $(a + 1, b), (a, b + 1) \in R$ *and* $(a, b) \notin R$*, or*
(3) $P_M[a, b] = 1$*,* $(a, b) \notin L$*, and* $(a + 1, b + 1) \notin R$*.*

Moreover, $(a + 1, b), (a, b + 1) \notin L$ holds for every $(a, b) \in [\,\ell \mathinner{.\,.} r\,]^2 \setminus L$ and $(a - 1, b), (a, b - 1) \notin R$ holds for every $(a, b) \in [\,\ell \mathinner{.\,.} r\,]^2 \setminus R$.

Proof. Let us first prove an auxiliary claim characterizing the *restricted distance matrix $D_{M|_I}$*.

Claim 8.17. *For every $a, b \in \mathbb{Z}$, the restricted distance matrix satisfies*

$$D_{M|_I}[a, b] = \begin{cases} \min\{D_M[a, b], a + b - 2\ell, 2r - a - b\} & \text{if } a, b \in (\,\ell \mathinner{.\,.} r\,), \\ |a - b| & \text{otherwise.} \end{cases}$$

Proof. Consider the smallest bounding box B of $M|_I$. The second case follows from the fact that $\mathsf{span}(P_{M|_I}) \subseteq \mathsf{span}(M|_I) \subseteq (\,\ell \mathinner{.\,.} r\,)$. As for the first case, note that $\mathsf{AG}(M)$ is a subgraph of $\mathsf{AG}(M|_I)$, so $D_{M|_I}[a, b] \leq D_M[a, b]$. Moreover, the triangle inequality yields

$$D_{M|_I}[a, b] \leq D_{M|_I}[a, \ell] + D_{M|_I}[\ell, b] = |a - \ell| + |\ell - b| = a + b - 2\ell \quad \text{and}$$
$$D_{M|_I}[a, b] \leq D_{M|_I}[a, r] + D_{M|_I}[r, b] = |a - r| + |r - b| = 2r - a - b.$$

Finally, observe that if a path from lt_a^B to br_b^B in $\mathsf{AG}(M|_I)$ does not reach any vertex (x, y) with $x - y \notin I$, then path is also present in $\mathsf{AG}(M)$. Otherwise, the path must reach a vertex (x, y) with $x - y \in \{\ell, r\}$, so its length is at least $\min\{|a - \ell| + |\ell - b|, |a - r| + |r - b|\} = \min\{a + b - 2\ell, 2r - a - b\}$. ◾

Claim 8.17 implies

$$L = \{(a, b) \in [\,\ell\mathinner{\ldotp\ldotp}r\,]^2 : D_{M|_I}[a, b] = a + b - 2\ell\} \quad \text{and}$$
$$R = \{(a, b) \in [\,\ell\mathinner{\ldotp\ldotp}r\,]^2 : D_{M|_I}[a, b] = 2r - a - b\}.$$

By Definitions 8.7 and 8.9, the neighboring entries of $D_{M|_I}$ differ by at most one. Consequently, if $(a, b) \in [\,\ell\mathinner{\ldotp\ldotp}r\,]^2 \setminus L$, then $D_{M|_I}[a, b] < a + b - 2\ell$, so $D_{M|_I}[a + 1, b] < a + b - 2\ell + 1 = (a + 1) + b - 2\ell$ and $D_{M|_I}[a, b+1] < a + b - 2\ell + 1 = a + (b + 1) - 2\ell$. Similarly, if $(a, b) \in [\,\ell\mathinner{\ldotp\ldotp}r\,]^2 \setminus R$, then $D_{M|_I}[a, b] < 2r - a - b$, so $D_{M|_I}[a - 1, b] < 2r - a - b + 1 = 2r - (a - 1) - b$ and $D_{M|_I}[a, b - 1] < 2r - a - b + 1 = 2r - a - (b - 1)$.

Thus, it remains to prove the statement characterizing $P_{M|_I}$. In Case (1), we have

$$
\begin{aligned}
P_{M|_I}[a, b] &= D_{M|_I}[a + 1, b] + D_{M|_I}[a, b + 1] - D_{M|_I}[a, b] - D_{M|_I}[a + 1, b + 1] \\
&> (a + 1 + b - 2\ell) + (a + b + 1 - 2\ell) - (a + b - 2\ell) - (a + 1 + b + 1 - 2\ell) = 0.
\end{aligned}
$$

In Case (2), we have

$$
\begin{aligned}
P_{M|_I}[a, b] &= D_{M|_I}[a + 1, b] + D_{M|_I}[a, b + 1] - D_{M|_I}[a, b] - D_{M|_I}[a + 1, b + 1] \\
&> (2r - a - 1 - b) + (2r - a - b - 1) - (2r - a - b) - (2r - a - b - 1) = 0.
\end{aligned}
$$

In Case (3), we have

$$
\begin{aligned}
P_{M|_I}[a, b] &= D_{M|_I}[a + 1, b] + D_{M|_I}[a, b + 1] - D_{M|_I}[a, b] - D_{M|_I}[a + 1, b + 1] \\
&= D_M[a + 1, b] + D_M[a, b + 1] - D_M[a, b] - D_M[a + 1, b + 1] \\
&= P_M[a, b] = 1.
\end{aligned}
$$

Thus, if at least of the three cases holds, then indeed $P_{M|_I}[a, b] = 1$.

For a proof of the converse implication, suppose that $P_{M|_I}[a, b] = 1$ for some $a, b \in [\,\ell\mathinner{\ldotp\ldotp}r\,)$. Since the neighboring entries of $D_{M|_I}$ differ by at most one, the equality

$$1 = D_{M|_I}[a + 1, b] + D_{M|_I}[a, b + 1] - D_{M|_I}[a, b] - D_{M|_I}[a + 1, b + 1]$$

implies $D_{M|_I}[a + 1, b + 1] = D_{M|_I}[a, b]$ and $D_{M|_I}[a + 1, b] = D_{M|_I}[a, b + 1] = D_{M|_I}[a, b] + 1$. If $(a, b) \in L$, then we have $(a + 1, b), (a, b + 1) \in L$ and $(a + 1, b + 1) \notin L$, that is, case (1) holds. If $(a + 1, b + 1) \in R$, then we have $(a + 1, b), (a, b + 1) \in R$ and $(a, b) \notin R$, that is, case (2) holds Finally, if $(a, b) \notin L$ and $(a + 1, b + 1) \notin R$, then $(a, b), (a + 1, b), (a, b + 1), (a + 1, b + 1) \notin L \cup R$, so

$$
\begin{aligned}
P_M[a, b] &= D_M[a + 1, b] + D_M[a, b + 1] - D_M[a, b] - D_M[a + 1, b + 1] \\
&= D_{M|_I}[a + 1, b] + D_{M|_I}[a, b + 1] - D_{M|_I}[a, b] - D_{M|_I}[a + 1, b + 1] \\
&= P_{M|_I}[a, b] = 1,
\end{aligned}
$$

that is, Case (3) holds. ⌟

Corollary 8.18. *Consider two finite sets $M, \hat{M} \subseteq \mathbb{Z}^2$ with $P_M = P_{\hat{M}}$. For every integer interval $I \subseteq \mathbb{Z}$, we have $P_{M|_I} = P_{\hat{M}|_I}$.*

Proof. If I is a finite interval, then Lemma 8.16 provides a complete characterization of $P_{M|_I}$ in terms of D_M and P_M, whereas Lemma 8.11 provides a complete characterization of D_M in terms of P_M. Consequently, $P_{M|_I} = P_{\hat{M}|_I}$ holds in this case. If I is an infinite interval, then we define $J = \mathsf{span}(M \cup \hat{M}) \cap I$ and observe that $M|_I = M|_J$ as well as $\hat{M}|_I = \hat{M}|_J$, so $P_{M|_I} = P_{M|_J} = P_{\hat{M}|_J} = P_{\hat{M}|_I}$. ⌟

Corollary 8.18 combined with Lemma 8.14 let us define the restriction operation on bounded permutation matrices.

◨ **Definition 8.19**. *For a bounded permutation matrix A and an interval I, we define the* restriction of A *to* I*, denoted* $A|_I$*, as* $P_{M|_I}$*, where $M \subseteq \mathbb{Z}^2$ is an arbitrary finite set such that $A = P_M$.* ◪

◨ **Lemma 8.20**. *The following equalities hold for all bounded permutation matrices A, B, intervals $I, J \subseteq \mathbb{Z}$, and shifts $s \in \mathbb{Z}$:*
(a) $A|_I \boxdot B|_I = (A \boxdot B)|_I$*;*
(b) $(A|_I)|_J = A|_{I \cap J}$*;*
(c) $(A|_I) \searrow s = (A \searrow s)|_{I+s}$*.*

Proof. (a) Let $M, \hat{M} \subseteq \mathbb{Z}^2$ be finite sets such that $A = P_M$ and $B = P_{\hat{M}}$ (these sets exist by Lemma 8.14). By Definitions 8.7 and 8.9, we can shift M and $\hat{M}$ along the diagonals without influencing P_M and $P_{\hat{M}}$. In particular, this lets us assume that M and $\hat{M}$ admit vertically adjacent bounding boxes so that $A \boxdot B = P_{M \cup \hat{M}}$ holds by Lemma 8.13. Since $M|_I$ and $\hat{M}_I$ also admit the same vertically adjacent bounding boxes, Lemma 8.13 also yields $A|_I \boxdot B|_I = P_{M|_I \cup \hat{M}|_I} = P_{(M \cup \hat{M})|_I} = (A \boxdot B)|_I$.

(b) Let $M \subseteq \mathbb{Z}^2$ be a finite set such that $A = P_M$. Observe that $(M|_I)|_J = \{(x, y) \in M|_I : x - y \in J\} = \{(x, y) \in M : x - y \in I \cap J\} = M|_{I \cap J}$. By Definition 8.19, this implies $(A|_I)_J = A_{I \cap J}$.

(c) Let $M \subseteq \mathbb{Z}^2$ be a finite set such that $A = P_M$. Observe that $(M + (s, 0))|_{I+s} = \{(x, y) \in M + (s, 0) : x - y \in I + s\} = \{(x + s, y) \in M + (s, 0) : x + s - y \in I + s\} = \{(x + s, y) : (x, y) \in M \text{ and } x - y \in I\} = (M|_I) + (s, 0)$. Consequently, Fact 8.10 implies $(A \searrow s)|_{I+s} = P_{(M + (s, 0))|_{I+s}} = P_{(M|_I) + (s, 0)} = (A|_I) \searrow s$. ◪

Next, we turn the combinatorial characterization of Lemma 8.16 into an efficient algorithm for restricting permutation matrices.

◨ **Lemma 8.21** (Restrict(σ, I)). *Given a permutation matrix A (represented by a permutation σ) and an integer interval I, the permutation matrix $A|_I$ (represented by a permutation $\sigma|_I$) can be constructed in $\mathcal{O}(|\mathsf{span}(A)|)$ time.*

Proof. Write $(\ell \mathinner{.\,.} r) \coloneqq I \cap \mathsf{span}(A)$; observe that $A|_I = A|_{(\ell \mathinner{.\,.} r)}$, so we construct $\sigma|_I$ using the characterization of Lemma 8.16. For each $a \in [\ell \mathinner{.\,.} r)$, we initialize $\sigma|_I(a) \coloneqq \sigma(a)$; this covers case (3), so it suffices to assign the values $\sigma|_I(a)$ corresponding to cases (1) and (2). As for case (1), we traverse the *boundary* of the set L (that is, pairs $(a, b) \in L$ with $(a + 1, b + 1) \notin L$). A symmetric traversal of the boundary of the set R (that is, pairs $(a, b) \in L$ with $(a - 1, b - 1) \notin R$) covers case (2). The algorithm implementing this strategy is presented as Algorithm 3.

Let us focus on the traversal of the boundary of L. Observe that $(r, \ell), (\ell, r) \in L$ due to $D_M[r, \ell] \geq |r - \ell| = r + \ell - 2\ell$ and $D_M[\ell, r] \geq |\ell - r| = \ell + r - 2\ell$. Consequently, by the last part of Lemma 8.16, the boundary of L forms a path from (r, ℓ) to (ℓ, r) with individual steps going either down or to the left. In Lines 4–15, we traverse this path while maintaining $v = D_M[a, b]$. This value is initialized to $D_M[\ell, r] = 2P_M^\Sigma[\ell, r] + \ell - r$ according to Lemma 8.11 and the definition of the distribution matrix.

After visiting $(a, b) \in L$, we proceed to $(a + 1, b)$ or $(a, b - 1)$ depending on whether $(a + 1, b) \in L$ or not. According to the definition of L, this condition is equivalent to $a + 1 + b \leq \ell + r$ and $D_M[a + 1, b] \geq a + 1 + b - 2\ell$. The first part is verified explicitly, whereas for the second one, we observe that $D_M[a + 1, b] - D_M[a, b] = 1 - 2 \sum_{b' \in [\ell \mathinner{.\,.} b)} P_M[a, b']$, that is, $D_M[a + 1, b] = v + 1$ if $\sigma(a) \geq b$ and $D_M[a + 1, b] = v - 1$ otherwise. In the former case, $v \geq a + b - 2\ell$ yields $v + 1 \geq a + 1 + b - 2\ell$, whereas in the latter one, we explicitly check whether $v - 1 \geq a + 1 + b - 2\ell$. If this test, implemented in Line 7, reveals $(a + 1, b) \in L$, we move to $(a + 1, b)$ and update v as described above. Otherwise, we proceed to $(a, b - 1)$. In order to update v, we observe that $D_M[a, b - 1] - D_M[a, b] = 1 - 2 \sum_{a' \in [a \mathinner{.\,.} r)} P_M[a', b - 1]$,

■ **Algorithm 3** Restricting the seaweed matrix.

```
1 Restrict(σ, I)
2     ( ℓ . . r ) ← I ∩ span(A);
3     foreach d ∈ [ ℓ . . r ) do σ|_I(d) := σ(d);
4     (a, b) := (ℓ, r);
5     v := 2|{d ∈ dom(σ) : d ≥ ℓ and σ(d) < r}| + ℓ − r;
6     while (a, b) ≠ (r, ℓ) do
7         if a + 1 + b ≤ ℓ + r and (σ(a) ≥ b or v − 1 ≥ a + 1 + b − 2ℓ) then
8             if σ(a) ≥ b then v := v + 1;
9             else v := v − 1;
10            a := a + 1;
11        else
12            σ|_I(a) := b − 1;
13            if σ_M^{-1}(b − 1) < a then v := v + 1;
14            else v := v − 1;
15            b := b − 1;
16    while (a, b) ≠ (ℓ, r) do
17        if a − 1 + b ≥ ℓ + r and (σ(a − 1) < b or v − 1 ≥ 2r − a + 1 − b) then
18            if σ(a − 1) < b then v := v + 1;
19            else v := v − 1;
20            a := a − 1;
21        else
22            σ|_I(a − 1) := b;
23            if σ_M^{-1}(b) ≥ a then v := v + 1;
24            else v := v − 1;
25            b := b + 1;
26    return σ|_I;
```

that is, $D_M[a, b − 1] = v + 1$ if $\sigma(b − 1) < a$ and $D_M[a, b − 1] = v − 1$ otherwise. In this case, we also know that $\sigma|_I(a)$ is covered by case (1) of Lemma 8.16. Thus, we set $\sigma|_I(a) = b − 1$ in Line 12; this value is either assigned correctly (if $(a + 1, b − 1) \in L$) or overwritten in the next step (otherwise).

Lines 16–25 implement traversal of the boundary of R and cover case (2) of Lemma 8.16. Observe that $(\ell, r), (r, \ell) \in R$ due to $D_M[\ell, r] \geq |\ell − r| = 2r − \ell − r$ and $D_M[r, \ell] \geq |r − \ell| = 2r − \ell − r$. Consequently, by the last part of Lemma 8.16, the boundary of R forms a path from (ℓ, r) to (r, ℓ) with individual steps going either up or to the right. We traverse this path maintaining $v = D_M[a, b]$. Given that the traversal of the boundary of L terminated at $(a, b) = (\ell, r)$, we do not need to initialize v.

After visiting $(a, b) \in R$, we proceed to $(a − 1, b)$ or $(a, b + 1)$ depending on whether $(a − 1, b) \in R$ or not. According to the definition of R, this condition is equivalent to $a − 1 + b \geq \ell + r$ and $D_M[a − 1, b] \geq 2r − a + 1 − b$. The first part is verified explicitly, whereas for the second one, we observe that

$$D_M[a − 1, b] − D_M[a, b] = 2 \sum_{b' \in [\ell . . b)} P_M[a − 1, b'] − 1,$$

that is, $D_M[a − 1, b] = v + 1$ if $\sigma(a − 1) < b$ and $D_M[a − 1, b] = v − 1$ otherwise. In the former case, $v \geq 2r − a − b$ yields $v + 1 \geq 2r − a + 1 − b$, whereas in the latter one, we explicitly check whether $v − 1 \geq 2r − a + 1 − b$. If this test, implemented in Line 17, reveals $(a − 1, b) \in R$, we move to $(a − 1, b)$

and update v as described above. Otherwise, we proceed to $(a, b+1)$. In order to update v, we observe that $D_M[a, b+1] - D_M[a, b] = 2\sum_{a' \in [a .. r)} P_M[a', b] - 1$, that is, $D_M[a, b+1] = v+1$ if $\sigma(b) \geq a$ and $D_M[a, b+1] = v - 1$ otherwise. In this case, we also know that $\sigma|_I(a-1)$ is covered by case (2) of Lemma 8.16. Thus, we set $\sigma|_I(a-1) = b$ in Line 22; this value is either assigned correctly (if $(a-1, b+1) \in R$) or overwritten in the next step (otherwise).

The running time of the algorithm is $\mathcal{O}(|\mathsf{span}(A)|)$; this includes building the permutation σ_M^{-1}. ◪

8.3 Computing the Seaweed Matrix

Lemma 8.22. *Consider a finite set $M \subseteq \mathbb{Z}^2$. For every $(x, y) \in \mathbb{Z}$, there are unique thresholds $\phi(x, y), \psi(x, y) \in \mathbb{Z}$ such that the following holds for any bounding box $B \in [i .. i'] \times [j .. j']$ of M:*

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) - \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) = \begin{cases} 1 & \text{if } d \leq \psi(x, y), \\ -1 & \text{otherwise.} \end{cases}$$

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) - \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) = \begin{cases} -1 & \text{if } d \leq \phi(x, y), \\ 1 & \text{otherwise.} \end{cases}$$

Moreover,
- *If $(x, y) \in M$, then $\phi(x+1, y) = \min\{\psi(x, y), \phi(x, y)\}$ and $\psi(x, y+1) = \max(\psi(x, y), \phi(x, y))$.*
- *Otherwise, $\phi(x+1, y) = \psi(x, y)$ and $\psi(x, y+1) = \phi(x, y)$.*

Proof. Let us first consider points $(x, y) \in \mathbb{Z}^2$ with $x < i$ or $y < j$. In this case, we have $\mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y+1)) = |x-y-d|$, $\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) = |x+1-y-d|$, $\mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) = |x-y-1-d|$. Consequently, $\psi(x, y) = \phi(x+1, y) = x - y$ and $\phi(x, y) = \psi(x, y+1) = x - y - 1$.

For the remaining points, we proceed by induction on $x + y$. Suppose that the existence of unique values $\psi(x, y)$ and $\phi(x, y)$ has already been established. If $(x, y) \notin M$, then $\mathsf{dist}_M((x, y), (x+1, y+1)) = 0$. We consider four cases depending on whether $d \leq \psi(x, y)$ and $d \leq \phi(x, y)$.

- If $d \leq \psi(x, y)$ and $d \leq \phi(x, y)$, then

$$\begin{aligned} \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) - 1 \\ &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) + 1. \end{aligned}$$

- If $d \leq \psi(x, y)$ and $d > \phi(x, y)$, then

$$\begin{aligned} \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) - 1 \\ &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) - 1. \end{aligned}$$

- If $d > \psi(x, y)$ and $d \leq \phi(x, y)$, then

$$\begin{aligned} \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) + 1 \\ &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) + 1. \end{aligned}$$

- If $d > \psi(x, y)$ and $d > \phi(x, y)$, then

$$\begin{aligned} \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y)) &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1, y)) + 1 \\ &= \mathsf{dist}_M(\mathsf{lt}_d^B, (x, y+1)) - 1. \end{aligned}$$

Based on this case analysis, we see that $\phi(x+1,y) = \psi(x,y)$ and $\psi(x,y+1) = \phi(x,y)$ are unique and well-defined.

If $(x,y) \in M$, then Lemma 8.5 implies

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y+1)) = 1 + \min\{\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y)), \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y+1))\}.$$

We consider the same four cases as previously:

- If $d \leq \psi(x,y)$ and $d \leq \phi(x,y)$, then

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y)) - 1 = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y+1)) + 1,$$

so

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)).$$

- If $d \leq \psi(x,y)$ and $d > \phi(x,y)$, then

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y)) - 1 = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y+1)) - 1,$$

so

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)) + 2.$$

- If $d > \psi(x,y)$ and $d \leq \phi(x,y)$, then

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y)) + 1 = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y+1)) + 1,$$

so

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)).$$

- If $d > \psi(x,y)$ and $d > \phi(x,y)$, then

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y)) + 1 = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y+1)) - 1,$$

so

$$\mathsf{dist}_M(\mathsf{lt}_d^B, (x+1,y+1)) = \mathsf{dist}_M(\mathsf{lt}_d^B, (x,y)).$$

Based on this case analysis, we conclude that $\phi(x+1,y) = \min\{\psi(x,y), \phi(x,y)\}$ and $\psi(x,y+1) = \max(\psi(x,y), \phi(x,y))$ are unique and well-defined. ∎

Lemma 8.23 ($\mathtt{Seaweed}(I, \mathtt{successor}(M, \star, \star))$). *Consider a finite set $M \subseteq \mathbb{Z}^2$. Suppose that we are given an interval $I \supseteq \mathsf{span}(M)$ and function $\mathtt{successor}(M, \star, \star)$ that, given $d \in I$ and $y \in \{-\infty\} \cup \mathbb{Z}$, in $\mathcal{O}(1)$ time returns $\min\{y' \geq y : (y'+d, y') \in M\}$, where $\min \varnothing = \infty$. Then, a permutation $\sigma_M : {}^+I \to {}^+I$ representing P_M can be constructed in $\mathcal{O}(|I|^2 \log\log|I|)$ time.*

Proof. Write $I := (\ell .. r)$. Our solution, presented as Algorithm 4 maintains a permutation $\sigma : {}^+I \to {}^+I$ and implicitly iterates over all pairs $(y,d) \in \mathbb{Z} \times [\ell-1 .. r]$ in the lexicographic order. The main invariant is that, after processing (y,d), the permutation σ is as follows:

$$\sigma(d') = \begin{cases} \psi(y+1+d', y+1) & \text{for } d' \in [\ell .. d), \\ \phi(y+1+d, y) & \text{for } d' = d, \\ \psi(y+d', y) & \text{for } d' \in (d .. r). \end{cases}$$

Additionally, the algorithm maintains a priority queue $\mathcal{Q}$ that stores events $(y', d') \in \mathbb{Z} \times (\ell .. r)$ satisfying the following three invariants for every $d' \in (\ell .. r)$:

■ **Algorithm 4** Constructing the seaweed matrix.

---

1 $\text{Seaweed}(I, \text{successor}(M, \star, \star))$

2 $(\ell \mathinner{..} r) := I;$

3   **foreach** $d \in [\ell \mathinner{..} r)$ **do** $\sigma(d) := d;$

4   **foreach** $d \in (\ell \mathinner{..} r)$ **do** $\mathcal{Q}.\text{insert}(\text{successor}(M, d, -\infty), d);$

5   **while** $\mathcal{Q}$ *not empty* **do**

6 $(y, d) := \mathcal{Q}.\text{extractMin}();$

7   **if** $y \neq \infty$ **and** $\sigma(d - 1) < \sigma(d)$ **then**

8 $\text{swap}(\sigma(d - 1), \sigma(d));$

9   **if** $d - 1 \in (\ell \mathinner{..} r)$ **then** $\mathcal{Q}.\text{insert}(\text{successor}(M, d - 1, y + 1), d - 1);$

10   **if** $d + 1 \in (\ell \mathinner{..} r)$ **then** $\mathcal{Q}.\text{insert}(\text{successor}(M, d + 1, y), d + 1);$

11 **return** $\sigma^{-1};$

---

(1) If $(y', d') \in \mathcal{Q}$ with $y' \neq \infty$, then $(y', d') \succ_{lex} (y, d)$ and $(y' + d', y') \in M$.

(2) If $\sigma(d' - 1) < \sigma(d')$ and $d' \leq d$, then $(\text{successor}(M, d', y + 1), d') \in \mathcal{Q}$.

(3) If $\sigma(d' - 1) < \sigma(d')$ and $d' > d$, then $(\text{successor}(M, d', y), d') \in \mathcal{Q}$.

In the initialization phase (Lines 3–4), we set $\sigma(d) = d$ for $d \in {}^+I$ and $\mathcal{Q} = \{(\text{successor}(M, d, -\infty), d) : d \in (\ell \mathinner{..} r)\}$. The processing of (y, d) handled implicitly if $(y, d) \notin \mathcal{Q}$. Otherwise, we extract (y, d) from $\mathcal{Q}$ and check whether $\sigma(d - 1) < \sigma(d)$. If so, we swap $\sigma(d - 1)$ with $\sigma(d)$ and insert to $\mathcal{Q}$ pairs $(\text{successor}(M, d - 1, y + 1), d - 1)$ (provided that $d - 1 \in (\ell \mathinner{..} r)$) and $(\text{successor}(M, d - 1, y), d + 1)$ (provided that $d + 1 \in (\ell \mathinner{..} r)$). After processing all $(y, d) \in \mathbb{Z} \times [\ell - 1 \mathinner{..} r]$, the queue $\mathcal{Q}$ may still contain entries of the form $\{\infty\} \times (\ell \mathinner{..} r)$, which are extracted one by one until $\mathcal{Q}$ is empty.

Let us first prove that the invariants are satisfied after the initialization. For this, we select $d = \ell - 1$ and an arbitrary value y with $M \subseteq \mathbb{Z} \times [y \mathinner{..} \infty)$. Since $\psi(y + d', d') = d'$ holds for all $d' \in {}^+I$, the permutation σ satisfies the invariant. As for the invariants regarding $\mathcal{Q}$, note that $\mathcal{Q} = \{(\text{successor}(M, d', y), d') : d' \in (\ell \mathinner{..} r)\}$. In particular, $\mathcal{Q}$ contains all the required pairs. Moreover, if $(y', d') \in \mathcal{Q}$ with $y' \neq \infty$, then $y' \geq y$ (and thus $(y', d') \succ_{lex} (y, d)$) and $(y' + d', y') \in M$ (by definition of $\text{successor}(M, \star, \star)$).

Next, consider processing $(y, d) \in [\ell - 1 \mathinner{..} r]$. If $d = \ell - 1$, the invariants for $(y, \ell - 1)$ are exactly the same as the invariants for $(y - 1, r)$. If $d = \ell$, then the only difference between the invariants for $(y, \ell - 1)$ and (y, ℓ) is that $\sigma(\ell)$ needs to be updated from $\psi(y + \ell, y)$ to $\phi(y + 1 + \ell, y)$. However, due to $(y + \ell, y) \notin M$, Lemma 8.22 guarantees $\phi(y + 1 + \ell, y) = \psi(y + \ell, y)$. If $d = r$, then the only difference between the invariants for $(y, r - 1)$ and (y, r) is that $\sigma(r - 1)$ needs to be updated from $\phi(y + r, y)$ to $\psi(y + r, y + 1)$. However, due to $(y + r, y) \notin M$, Lemma 8.22 guarantees $\psi(y + r, y + 1) = \phi(y + r, y)$. Hence, in each of the above three cases, the void implementation (guaranteed by $\mathcal{Q} \subseteq \mathbb{Z} \times (\ell \mathinner{..} r)$) is correct.

In the main case of $d \in (\ell \mathinner{..} r)$, the values $\sigma(d - 1), \sigma(d)$ should be updated from $\phi(y + d, y), \psi(y + d, y)$ to $\psi(y + d, y + 1), \phi(y + d + 1, y)$. If $(y + d, y) \notin M$, then $(y + d, y) \notin \mathcal{Q}$ and our implementation is void. This is correct because Lemma 8.22 yields $\psi(y + d, y + 1) = \phi(y + d, y)$ and $\phi(y + d + 1, y) = \psi(y + d, y)$. If $(y + d, y) \in M$ and $\sigma(d - 1) < \sigma(d)$, then the invariant on $\mathcal{Q}$ guarantees $(y + d, y) \in \mathcal{Q}$. Hence, our algorithm swaps $\sigma(d - 1)$ with $\sigma(d)$. This is correct because Lemma 8.22 yields $\psi(y + d, y + 1) = \max(\psi(y + d, y), \phi(y + d, y)) = \psi(y + d, y)$ and $\phi(y + d + 1, y) = \min\{\psi(y + d, y), \phi(y + d, y)\} = \phi(y + d, y)$. If $(y + d, y) \in M$ and $\sigma(d - 1) \geq \sigma(d)$, the queue $\mathcal{Q}$ may contain $(y + d, y)$ or not. In both scenarios, our algorithm keeps $\sigma(d - 1)$ and $\sigma(d)$ intact This is correct because Lemma 8.22 yields $\psi(y + d, y + 1) = \max(\psi(y + d, y), \phi(y + d, y)) = \phi(y + d, y)$ and $\phi(y + d + 1, y) = \min\{\psi(y + d, y), \phi(y + d, y)\} = \psi(y + d, y)$.

Next, we shall prove that the invariants regarding $\mathcal{Q}$ remain satisfied. Since (y, d) is removed from $\mathcal{Q}$, all the remaining elements (y', d') satisfy $(y', d') \succ_{lex} (y, d)$ and $(y', d') \in M$. By definition of $\texttt{successor}(M, \star, \star)$, this is also true for the newly inserted elements, if any. Next, consider $d' \in (\ell .. r)$ with $\sigma(d' - 1) < \sigma(d)$. If $d' \in (\ell .. d - 1)$, the entries $\sigma(d' - 1)$ and $\sigma(d')$ were kept intact, so $(\texttt{successor}(M, d', y + 1), d')$ is still guaranteed to be contained in $\mathcal{Q}$. The same is true for $d' = d - 1$ if we did not swap $\sigma(d - 1)$ with $\sigma(d)$. If we did, however, then $(\texttt{successor}(M, d - 1, y + 1), d - 1)$ is contained in $\mathcal{Q}$ because it was inserted explicitly. Symmetrically, for $d' \in (d + 1 .. r)$, the entries $\sigma(d' - 1)$ and $\sigma(d')$ were kept intact, so $(\texttt{successor}(M, d', y), d')$ is still guaranteed to be contained in $\mathcal{Q}$. The same is true for $d' = d + 1$ if we did not swap $\sigma(d - 1)$ with $\sigma(d)$. If we did, however, then $(\texttt{successor}(M, d + 1, y), d)$ is contained in $\mathcal{Q}$ because it was inserted explicitly. It remains to consider $d' = d$. Due to $\psi(y + d, y + 1) < \phi(y + d, y)$, Lemma 8.22 guarantees $(y, y + d) \notin M$. Consequently, $(\texttt{successor}(M, d, y), d) = (\texttt{successor}(M, d, y + 1), d)$ is still guaranteed to be contained in $\mathcal{Q}$.

We conclude the correctness analysis by considering $d = r$ and an arbitrary value y with $M \cap \mathbb{Z} \times (-\infty .. y]$ The invariants show that $\sigma(d') = \psi(y + d', d')$ for $d' \in {}^+I$ and that $\mathcal{Q}$ does not contain any entry (y', d') with $y' \neq \infty$ (so no further iterations alter σ). Recall that $P_M[a, b] = 1$ implies $D_M[a, b] = D_M[a + 1, b + 1]$ and $D_M[a, b + 1] = D_M[a + 1, b] = D_M[a, b] + 1$, that is, $D_M[a, b + 1] - D_M[a, b] = 1$ and $D_M[a + 1, b + 1] - D_M[a + 1, b] = -1$. Using Definitions 8.7 and 8.9 and Lemma 8.22, we conclude that $a \leq \sigma(b)$ and $a + 1 > \sigma(b)$, that is, that $a = \sigma(b)$. Since P_M is a permutation matrix, this guarantees that σ^{-1} is the permutation representing P_M.

As for the running time, we observe that each swap of subsequent entries of σ increases the number of inversions in σ, so the total number of swaps is $\mathcal{O}(|I|^2)$. Consequently, the total number of operations on $\mathcal{Q}$ is $\mathcal{O}(|I| + |I|^2) = \mathcal{O}(|I|^2)$. Using the state-of-the-art priority queries with integer keys [Han04, Tho07], we can achieve $\mathcal{O}(\log \log |I|)$ time per operation, for a total running time of $\mathcal{O}(|I|^2 \log \log |I|)$. ∎

9 Applications of Seaweeds

For two strings $X, Y \in \Sigma^*$, we define $M(X, Y) = \{(x, y) \in [0 .. |X|) \times [0 .. |Y|) : X[x] \neq Y[y]\}$. Moreover, we denote $D_{X,Y} = D_{M(X,Y)}$ and $P_{X,Y} = P_{M(X,Y)}$. Lemma 8.5 and Definition 8.7 yields the following characterization:

Fact 9.1. *For $X, Y \in \Sigma^*$ and $0 \leq \ell \leq r \leq |X|$, we have $\delta_D(X[\ell .. r), Y) = D_{X,Y}[\ell, r - |Y|]$.* ∎

Lemma 9.2. *All strings $X, Y, Z \in \Sigma^*$ satisfy $P_{X,YZ} = P_{X,Y} \boxdot (P_{X,Z} \searrow (-|Y|))$ and $P_{XY,Z} = P_{X,Z} \boxdot (P_{Y,Z} \searrow |X|)$.*

Proof. Observe that $M(X, YZ) = M(X, Y) \cup (M(X, Z) + (0, |Y|))$ and the two sets of the right-hand side have vertically adjacent bounding boxes. Hence, Fact 8.10 and Lemma 8.13 imply $P_{X,YZ} = P_{X,Y} \boxdot P_{M(X,Z)+(0,|Y|)} = P_{X,Y} \boxdot (P_{X,Z} \searrow (-|Y|))$. Similarly, $M(XY, Z) = M(X, Z) \cup (M(Y, Z) + (|X|, 0))$, and the two sets of the right-hand side have horizontally adjacent bounding boxes. Hence, Fact 8.10 and Lemma 8.13 imply $P_{XY,Z} = P_{X,Z} \boxdot P_{M(Y,Z)+(|X|,0)} = P_{X,Z} \boxdot (P_{Y,Z} \searrow |X|)$. ∎

The principle behind Lemma 9.2 was used in the algorithm of [CKM20] to maintain $\delta_D(X, Y)$ subject to edit operations in X, Y. The algorithm of [CKM20] actually maintains the seaweed matrix $P_{X,Y}$, so we get the following result:

Fact 9.3 ([CKM20]). *There exists a dynamic algorithm that maintains $P_{X,Y}$ subject to character insertions and deletions in $X, Y \in \Sigma^*$. The initialization costs $\mathcal{O}(|X| \cdot |Y|)$ time and the updates cost $\mathcal{O}((|X| + |Y|) \log^2(|X| + |Y|))$ time.* ∎

Lemma 9.4. *Consider families $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^{\leq n}$. The seaweed matrices $P_{X,Y}$ for $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ can all be constructed in $\mathcal{O}(n^2 + (1 + d_X)(1 + d_Y)n \log^2 n)$ time, where $d_X = \delta_E(\mathcal{X})$, and $d_Y = \delta_E(\mathcal{Y})$.*

Proof. We iterate over $(X, Y) \in \mathcal{X} \times \mathcal{Y}$ while maintaining $P_{X,Y}$ via Fact 9.3. The initialization costs $\mathcal{O}(n^2)$ time whereas updating (X, Y) to (X', Y') costs $\mathcal{O}((\delta_E(X, X') + \delta_E(Y, Y'))n \log^2 n)$ time. To bound the total contribution of the update costs, let $\hat{X}$ and $\hat{Y}$ be such that $\delta_E(\mathcal{X}) = \sum_{X \in \mathcal{X}} \delta_E(X, \hat{X})$ and $\delta_E(\mathcal{Y}) = \sum_{Y \in \mathcal{Y}} \delta_E(Y, \hat{Y})$. Then, $\delta_E(X, X') + \delta_E(Y, Y') \leq \delta_E(X, \hat{X}) + \delta_E(X', \hat{X}) + \delta_E(Y, \hat{Y}) + \delta_E(Y', \hat{Y})$. Consequently, the total update time is $\mathcal{O}((|\mathcal{Y}| \cdot \delta_E(\mathcal{X}) + |\mathcal{X}| \cdot \delta_E(\mathcal{Y}))n \log^2 n) = \mathcal{O}((1 + d_X)(1 + d_Y)n \log^2 n)$ because $|\mathcal{Y}| \leq 1 + \delta_E(\mathcal{Y})$ and $|\mathcal{X}| \leq 1 + \delta_E(\mathcal{X})$. ∎

Lemma 9.5. *There exists a dynamic algorithm that maintains a matrix $\boxdot_{i=1}^{t}(A_i \searrow \sum_{j=1}^{i} \delta_i)$ subject to insertions and deletions of pairs (δ_i, A_i) consisting of an integer δ_i and a bounded permutation matrix A_i. The update cost is $\mathcal{O}(w \log w \log t)$ time and the initialization costs $\mathcal{O}(wt \log w)$ time, where $w = |\mathsf{span}(\bigcup_{i=1}^{t} \mathsf{span}(A_i \searrow (\sum_{j=1}^{i} \delta_j)))|$.*

Proof. We maintain a balanced binary tree with t leaves representing $[1 \mathinner{..} t]$. A node ν whose subtree contains leaves representing $[\ell \mathinner{..} r]$ maintains $A_\nu := \boxdot_{i=\ell}^{r}(A_i \searrow \sum_{j=\ell}^{i} \delta_i)$ as well as $s_\nu = \sum_{i=\ell}^{r} \delta_i$. If ν is the ith leaf, then $A_\nu = A_i \searrow \delta_i$ and $s_\nu = \delta_i$. If ν is an internal node with children ν_L, ν_R, then $A_\nu = A_{\nu_L} \boxdot (A_{\nu_R} \searrow s_{\nu_L})$ and $s_\nu = s_{\nu_L} + s_{\nu_R}$; this matrix can be constructed in $\mathcal{O}(w \log w)$ time using Theorem 8.2. Any insertion or deletion requires updating $\mathcal{O}(\log t)$ nodes, whereas the initialization requires computing A_ν for all $\mathcal{O}(t)$ nodes. Thus, the running times are $\mathcal{O}(w \log w \log t)$ and $\mathcal{O}(wt \log w)$, respectively. ∎

Lemma 9.6. *Consider strings X, Y, an integer interval $I = [\ell \mathinner{..} r]$, and a decomposition $Y = Y_1 \cdots Y_t$ such that $Y_i = Y[y_i \mathinner{..} y'_i)$ for $i \in [1 \mathinner{..} t]$. Moreover, for $i \in [1 \mathinner{..} t]$, let $X_i = X[x_i \mathinner{..} x'_i)$, where $x_i \leq \max(0, y_i + \ell)$ and $x'_i \geq \min\{|X|, y'_i + r\}$, and $I_i \supseteq I - x_i + y_i$ is an integer interval. Then, $P_{X,Y}|_I = (\boxdot_{i=1}^{t}((P_{X_i, Y_i}|_{I_i}) \searrow (x_i - y_i)))|_I$.*

Proof. We start with an auxiliary claim.

Claim 9.7. For every $i \in [1 \mathinner{..} t]$, we have $P_{X,Y_i}|_{I+y_i} = ((P_{X_i, Y_i}|_{I_i})|_{I-x_i+y_i}) \searrow x_i$.

Proof. Let $M_i = M(X, Y_i)$ so that $P_{M_i} = P_{X,Y_i}$. Observe that $M_i|_{I+y_i} = \{(x, y) \in M_i : x - y \in I + y_i\} = \{(x, y) \in M_i : x - y - y_i \in I\} = \{(x, y) \in M_i : x - y - y_i \in [\ell \mathinner{..} r]\} = \{(x, y) \in M_i : x \in [y + y_i + \ell \mathinner{..} y + y_i + r]\} \subseteq [0 \mathinner{..} |X_i|) \times [y_i + \ell \mathinner{..} y'_i + r)$. At the same time, $M_i|_{I+y_i} \subseteq [0 \mathinner{..} |X|) \times [0 \mathinner{..} |Y_i|)$, so $M_i|_{I+y_i} \subseteq [x_i \mathinner{..} x'_i) \times [0 \mathinner{..} |Y_i|)$. In particular, $M_i|_{I+y_i} = (M_i \cap ([x_i \mathinner{..} x'_i) \times [0 \mathinner{..} |Y_i|)))|_{I+y_i} = (M(X_i, Y_i) - (0, y_i))|_{I+y_i}$, so $P_{X,Y_i}|_{I+y_i} = (P_{X_i, Y_i} \searrow x_i)|_{I+y_i} = (P_{X_i, Y_i}|_{I-x_i+y_i}) \searrow x_i$ holds by Lemma 8.20(c). Moreover, due to Lemma 8.20(b) and $I_i \supseteq I - x_i + y_i$, the equality $P_{X,Y_i}|_{I+y_i} = (P_{X_i, Y_i}|_{I-x_i+y_i}) \searrow x_i = ((P_{X_i, Y_i}|_{I_i})|_{I-x_i+y_i}) \searrow x_i$ holds as claimed. ∎

$$
\begin{aligned}
P_{X,Y}|_I &= \left(\boxdot_{i=1}^{t}(P_{X,Y_i} \searrow (-y_i))\right)|_I && \text{Lemma 9.2} \\
&= \left(\left(\boxdot_{i=1}^{t}(P_{X,Y_i} \searrow (-y_i))\right)|_I\right)|_I && \text{Lemma 8.20(b)} \\
&= \left(\boxdot_{i=1}^{t}((P_{X,Y_i} \searrow (-y_i))|_I)\right)|_I && \text{Lemma 8.20(a)} \\
&= \left(\boxdot_{i=1}^{t}((P_{X,Y_i}|_{I+y_i}) \searrow (-y_i))\right)|_I && \text{Lemma 8.20(c)} \\
&= \left(\boxdot_{i=1}^{t}(((P_{X_i, Y_i}|_{I_i})|_{I-x_i+y_i}) \searrow (x_i - y_i))\right)|_I && \text{Claim 9.7} \\
&= \left(\boxdot_{i=1}^{t}(((P_{X_i, Y_i}|_{I_i}) \searrow (x_i - y_i))|_I)\right)|_I && \text{Lemma 8.20(c)} \\
&= \left(\left(\boxdot_{i=1}^{t}((P_{X_i, Y_i}|_{I_i}) \searrow (x_i - y_i))\right)|_I\right)|_I && \text{Lemma 8.20(a)} \\
&= \left(\boxdot_{i=1}^{t}((P_{X_i, Y_i}|_{I_i}) \searrow (x_i - y_i))\right)|_I. && \text{Lemma 8.20(b)}
\end{aligned}
$$

This completes the proof. ∎

◤ **Corollary 9.8**. *Consider strings $X, \hat{X}, Y, \hat{Y}$, alignments $\mathcal{A}_X : \hat{X} \rightsquigarrow X$ and $\mathcal{A}_Y : \hat{Y} \rightsquigarrow Y$, an integer $k \geq \delta_E^{\mathcal{A}_X}(X, \hat{X}) + \delta_E^{\mathcal{A}_Y}(Y, \hat{Y})$, an interval $[\ell \mathinner{.\,.} r]$, and a decomposition $\hat{Y} = \hat{Y}_1 \cdots \hat{Y}_t$ such that $\hat{Y}_i = \hat{Y}[\hat{y}_i \mathinner{.\,.} \hat{y}_i')$ for $i \in [1 \mathinner{.\,.} t]$. For $i \in [1 \mathinner{.\,.} t]$, let*

- $\hat{X}_i = \hat{X}[\hat{x}_i \mathinner{.\,.} \hat{x}_i')$, *where* $\hat{x}_i = \max(0, \hat{y}_i + \ell - k)$ *and* $\hat{x}_i' = \min\{|\hat{X}|, \hat{y}_i' + r + k\}$.
- $Y_i = Y[y_i \mathinner{.\,.} y_i') = \mathcal{A}_Y(\hat{Y}_i)$.
- $X_i = X[x_i \mathinner{.\,.} x_i') = \mathcal{A}_X(\hat{X}_i)$.
- $I_i = [\ell - k - \hat{x}_i + \hat{y}_i \mathinner{.\,.} r + k - \hat{x}_i + \hat{y}_i]$

Then, $P_{X,Y}|_I = (\boxdot_{i=1}^{t} ((P_{X_i,Y_i}|_{I_i}) \searrow (x_i - y_i)))|_I$.

Proof. By Fact 2.2, $Y = Y_1 \cdots Y_t$ is a decomposition of Y. Due to Lemma 9.6, it suffices to prove that $x_i \leq \max(0, y_i + \ell)$, $x_i' \geq \min\{|X|, y_i' + r\}$, and $I_i \supseteq I - x_i + y_i$. If $\hat{x}_i = 0$, then $x_i = 0$. Otherwise, $x_i \leq \hat{x}_i + \delta_E^{\mathcal{A}_X}(X, \hat{X}) = \hat{y}_i + \ell - k + \delta_E^{\mathcal{A}_X}(X, \hat{X}) \leq y_i + \delta_E^{\mathcal{A}_Y}(Y, \hat{Y}) + \ell - k + \delta_E^{\mathcal{A}_X}(X, \hat{X}) \leq y_i + \ell$ holds as claimed. If $\hat{x}_i' = |\hat{X}|$, then $x_i' = |X|$. Otherwise, $x_i' \geq \hat{x}_i' - \delta_E^{\mathcal{A}_X}(X, \hat{X}) = \hat{y}_i' + r + k - \delta_E^{\mathcal{A}_X}(X, \hat{X}) \leq y_i' - \delta_E^{\mathcal{A}_Y}(Y, \hat{Y}) + r + k - \delta_E^{\mathcal{A}_X}(X, \hat{X}) \geq y_i' + r$ holds as claimed. Finally, note that $\ell - k - \hat{x}_i + \hat{y}_i \leq \ell - k - x_i + \delta_E^{\mathcal{A}_X}(X, \hat{X}) + y_i + \delta_E^{\mathcal{A}_Y}(Y, \hat{Y}) \leq \ell - x_i + y_i$ and $r + k - \hat{x}_i + \hat{y}_i \geq r + k - x_i - \delta_E^{\mathcal{A}_X}(X, \hat{X}) + y_i - \delta_E^{\mathcal{A}_Y}(Y, \hat{Y}) \geq r - x_i + y_i$, so $I_i \supseteq I - x_i + y_i$ holds as claimed. ◾

◤ **Lemma 9.9**. *Given a string family $\mathcal{S}$, a string $\bar{S} \in \mathcal{S}$ such that $\sum_{S \in \mathcal{S}} \delta_E(S, \bar{S}) \leq 2\delta_E(\mathcal{S})$ can be constructed in $\mathcal{O}(1 + \delta_E(\mathcal{S})^3)$ time in the* `PILLAR` *model.*

Proof. The algorithm computes $\delta_E(S, S')$ for all pairs of distinct strings $S, S' \in \mathcal{S}$ and returns the string $\bar{S} \in \mathcal{S}$ minimizing $\sum_{S \in \mathcal{S}} \delta_E(S, \bar{S})$.

As for correctness, consider $\hat{S} \in \Sigma^*$ such that $\delta_E(\mathcal{S}) = \sum_{S \in \mathcal{S}} \delta_E(S, \hat{S})$. Then, there exists $\bar{S}' \in \mathcal{S}$ such that $|\mathcal{S}|\delta_E(\bar{S}', \hat{S}) \leq \sum_{S \in \mathcal{S}} \delta_E(S, \hat{S}) = \delta_E(\mathcal{S})$. Consequently, $\sum_{S \in \mathcal{S}} \delta_E(S, \bar{S}') \leq \sum_{S \in \mathcal{S}} (\delta_E(S, \hat{S}) + \delta_E(\bar{S}', \hat{S})) \leq \delta_E(\mathcal{S}) + |\mathcal{S}|\delta_E(\bar{S}', \hat{S}) \leq 2\delta_E(\mathcal{S})$. Now, by definition of $\bar{S}$, we have $\sum_{S \in \mathcal{S}} \delta_E(S, \bar{S}) \leq \sum_{S \in \mathcal{S}} \delta_E(S, \bar{S}') \leq 2\delta_E(\mathcal{S})$.

As for the running time note that computing $\delta_E(S, S')$ costs $\mathcal{O}(\delta_E(S, S')^2) = \mathcal{O}((\delta_E(S, \hat{S}) + \delta_E(S', \hat{S}))^2)$ time. Across all pairs (S, S') this is bounded by $\mathcal{O}(|\mathcal{S}| \sum_{S \in \mathcal{S}} \delta_E(S, \hat{S})^2 + \sum_{S, S' \in \mathcal{S}} \delta_E(S, \hat{S})\delta_E(S', \hat{S})) = \mathcal{O}(|\mathcal{S}|\delta_E(\mathcal{S})^2) = \mathcal{O}(1 + \delta_E(\mathcal{S})^3)$. ◾

◤ **Lemma 9.10**. *Consider families $\mathcal{X}, \mathcal{Y} \subseteq \Sigma^{\leq n}$, a positive integer $d \geq \delta_E(\mathcal{X}) + \delta_E(\mathcal{Y})$, as well as an integer interval I. The matrices $P_{X,Y}|_I$ for $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ can all be constructed in $\mathcal{O}((d^3 + d|I|^2) \log^2(d + |I|))$ time in the* `PILLAR` *model.*

Proof. First, we compute $\hat{X} \in \mathcal{X}$ and $\hat{Y} \in \mathcal{Y}$ such that $\delta_E(\mathcal{X}, \hat{X}) \leq 2\delta_E(\mathcal{X})$ and $\delta_E(\mathcal{Y}, \hat{Y}) \leq 2\delta_E(\mathcal{Y})$ (Lemma 9.9) as well as optimal alignments $\mathcal{A}_X : \hat{X} \rightsquigarrow X$ and $\mathcal{A}_Y : \hat{Y} \rightsquigarrow Y$ for all $X \in \mathcal{X}$ and $Y \in \mathcal{Y}$ (represented by the underlying breakpoints).

Next, we iteratively construct of a partition $\hat{Y} = \hat{Y}_1 \cdots \hat{Y}_t$ of $\hat{Y}$, where $\hat{Y}_i = \hat{Y}[\hat{y}_i \mathinner{.\,.} \hat{y}_i')$ for $i \in [1 \mathinner{.\,.} t]$. In the i-th iteration, we set $\hat{y}_i = 0$ (if $i = 1$) or $\hat{y}_i = \hat{y}_{i-1}'$ (otherwise). Next, we define $\hat{y}_i'$. If $\hat{y}_i > |\hat{Y}| - (|I| + 4d)$, we set $\hat{y}_i' = |\hat{Y}|$. Otherwise, we define $\hat{Y}_i = \hat{Y}[\hat{y}_i \mathinner{.\,.} \hat{y}_i')$ to be the longest possible fragment of $\hat{Y}_i$ starting position $\hat{y}_i$ such that either (a) $|\hat{Y}_i| = |I| + 4d$ or (b) $Y_i := \mathcal{A}_Y(\hat{Y}_i)$ matches $\hat{Y}_i$ for all $Y \in \mathcal{Y}$ and $X_i := \mathcal{A}_X(\hat{X}_i)$ matches $\hat{X}_i$ for all $X \in \mathcal{X}$, where $\hat{X}_i$ is defined as in Corollary 9.8 for $k = 2d$. This construction partitions $\hat{Y}$ into *perfect* fragments satisfying condition (b) and the remaining *imperfect* fragments.

Let us denote $\mathcal{U}_i = \{\mathcal{A}_X(\hat{X}_i) : X \in \mathcal{X}\}$ and $\mathcal{Y}_i = \{A_Y(\hat{Y}_i) : Y \in \mathcal{Y}\}$. Our goal is to compute $P_{X_i,Y_i}|_{I_i}$ (with I_i defined in Corollary 9.8 for $k = 2d$) for all $(X_i, Y_i) \in \mathcal{U}_i \times \mathcal{Y}_i$. If $\hat{Y}_i$ is a perfect fragment, then $|\mathcal{U}_i| = |\mathcal{Y}_i| = 1$, and we apply Lemma 8.23. Otherwise, we use Lemma 9.4 to first construct P_{X_i,Y_i}, and then we derive $P_{X_i,Y_i}|_{I_i}$ using Lemma 8.21.

Finally, we initialize the data structure of Lemma 9.5 with $A_i = P_{\hat{X}_i,\hat{Y}_i}|_{I_i}$ and $\delta_1 = 0$ and $\delta_i = \hat{x}_i - \hat{y}_i - \hat{x}_{i-1} + \hat{y}_{i-1}$ for $i \in [\,2\mathinner{.\,.}t\,]$. For each $(X,Y) \in \mathcal{X} \times \mathcal{Y}$, we substitute $A_i \coloneqq P_{X_i,Y_i}|_{I_i}$ whenever $(X_i,Y_i) \neq (\hat{X}_i,\hat{Y}_i)$ and $\delta_i \coloneqq x_i - y_i - x_{i-1} + y_{i-1}$ whenever $x_i - y_i - x_{i-1} + y_{i-1} \neq \hat{x}_i - \hat{y}_i - \hat{x}_{i-1} + \hat{y}_{i-1}$. Then, we retrieve the matrix $A \coloneqq \boxdot_{i=1}^t (A_i \searrow \sum_{j=1}^i \delta_i)$ from the dynamic algorithm of Corollary 9.8 and return $A|_I$, computed using Lemma 8.21. Finally, we undo all the substitutions applied for (X,Y).

Correctness of the algorithm follows directly from Corollary 9.8. It remains to analyze the running time. First, we note that the applications of Lemma 9.9 take $\mathcal{O}(d^3)$ time. Constructing the alignment $\mathcal{A}_X : \hat{X} \rightsquigarrow X$ costs $\mathcal{O}(1 + \delta_E(X,\hat{X})^2)$ time, which sums up to $\mathcal{O}(d^2)$ across all $X \in \mathcal{X}$. Symmetrically, all the alignments $\mathcal{A}_Y : \hat{Y} \rightsquigarrow Y$ are built in $\mathcal{O}(d^2)$.

In order to efficiently implement the partitioning of $\hat{X}$, we construct the sets

$$\mathsf{B}_\mathcal{X} = \{\hat{x} \in [\,0\mathinner{.\,.}|\hat{X}|\,] : (\hat{x},x) \in \mathsf{B}_{\mathcal{A}_X} \text{ for some } X \in \mathcal{X} \text{ and } x \in [\,0\mathinner{.\,.}|X|\,]\} \quad \text{and}$$
$$\mathsf{B}_\mathcal{Y} = \{\hat{y} \in [\,0\mathinner{.\,.}|\hat{Y}|\,] : (\hat{y},y) \in \mathsf{B}_{\mathcal{A}_Y} \text{ for some } Y \in \mathcal{Y} \text{ and } y \in [\,0\mathinner{.\,.}|Y|\,]\}.$$

For each $X \in \mathcal{X}$, the contribution of X to $\mathsf{B}_\mathcal{X}$ can be constructed in $\mathcal{O}(1 + \delta_E(X,\hat{X}))$ time by scanning the (breakpoints behind) $\mathcal{A}_X$. Consequently, the set $\mathsf{B}_\mathcal{X}$ is of size $\mathcal{O}(d)$ and can be constructed in $\mathcal{O}(d \log d)$ time (this includes sorting and removing duplicates). A symmetric argument shows that $\mathsf{B}_\mathcal{Y}$ is of size $\mathcal{O}(d)$ and can be constructed in $\mathcal{O}(d \log d)$ time.

Recall that for $\hat{Y}_i = \hat{Y}[\hat{y}_i\mathinner{.\,.}\hat{y}_i')$, we have $\hat{X}_i = \hat{X}[\hat{x}_i\mathinner{.\,.}\hat{x}_i')$, where $\hat{x}_i = \max(0,\hat{y}_i + \ell - k)$ and $\hat{x}_i' = \min\{|\hat{X}|, \hat{y}_i' + r + k\}$. Thus, $\hat{Y}[\hat{y}_i\mathinner{.\,.}\hat{y}_i')$ is a perfect fragment as long as $[\,\hat{y}_i + \ell - k\mathinner{.\,.}\hat{y}_i' + r + k\,) \cap \mathsf{B}_\mathcal{X} = \varnothing$ and $[\,\hat{y}_i\mathinner{.\,.}\hat{y}_i'\,) \cap \mathsf{B}_\mathcal{Y} = \varnothing$. In $\mathcal{O}(\log d)$ time (by binary search over $\mathsf{B}_\mathcal{X}$ and $\mathsf{B}_\mathcal{Y}$), we can verify this condition for $\hat{y}_i' = \hat{y}_i + |I| + 4d$ and, if satisfied, determine the maximum possible $\hat{y}_i'$. Thus, the partition is constructed in $\mathcal{O}((d + t) \log d)$ time.

Our next goal is to prove that $t = \mathcal{O}(d)$. For this, we note that we cannot simultaneously have $[\,\hat{x}_i\mathinner{.\,.}\hat{x}_i'\,) \cap \mathsf{B}_\mathcal{X} = \varnothing = [\,\hat{x}_{i+1}\mathinner{.\,.}\hat{x}_{i+1}'\,)$ and $[\,\hat{y}_i\mathinner{.\,.}\hat{y}_i'\,) \cap \mathsf{B}_\mathcal{Y} = \varnothing = [\,\hat{y}_{i+1}\mathinner{.\,.}\hat{y}_{i+1}'\,) \cap \mathsf{B}_\mathcal{Y}$ for any $i \in [\,1\mathinner{.\,.}t\,)$. At the same time, $[\,\hat{x}_i\mathinner{.\,.}\hat{x}_i'\,) \cap [\,\hat{x}_{i+2}\mathinner{.\,.}\hat{x}_{i+2}'\,) = \varnothing = [\,\hat{y}_i\mathinner{.\,.}\hat{y}_i'\,) \cap [\,\hat{y}_{i+2}\mathinner{.\,.}\hat{y}_{i+2}'\,)$ holds for $i \in [\,1\mathinner{.\,.}t - 2\,]$ (because $\hat{x}_{i+2} \geq \hat{y}_{i+1}' + \ell - k \geq \hat{y}_{i+1} + |I| + 4d + \ell - k = \hat{x}_{i+1} + (r - \ell) + 2k + r - k = \hat{x}_i' + r + k \leq \hat{y}_i'$). Consequently, the number t' of imperfect fragments is at most $2|\mathsf{B}_\mathcal{Y}| + 2|\mathsf{B}_\mathcal{X}| = \mathcal{O}(d)$ and the partition size t satisfies $t \leq 2t' + 1 = \mathcal{O}(d)$. Moreover, each $X \in \mathcal{X}$ satisfies $\sum_{i=1}^t \delta_E(\hat{X}_i, X_i) \leq 2\delta_E(\hat{X}, X)$ and each $Y \in \mathcal{Y}$ satisfies $\sum_{i=1}^t \delta_E(\hat{Y}_i, Y_i) \leq 2\delta_E(\hat{Y}, Y)$. In particular, we conclude that $\sum_{i=1}^t \delta_E(\mathcal{U}_i) = \mathcal{O}(d)$ and $\sum_{i=1}^t \delta_E(\mathcal{Y}_i) = \mathcal{O}(d)$.

Each application of Lemma 9.4 costs $\mathcal{O}((|I| + d)^2 + (1 + \delta_E(\mathcal{U}_i))(1 + \delta_E(\mathcal{Y}_i))(|I| + d) \log^2(|I| + d))$ time, which is $\mathcal{O}(d(|I| + d)^2 + d^2(|I| + d) \log^2(|I| + d)) = \mathcal{O}((d^3 + d|I|^2) \log^2(|I| + d))$ in total. On the other hand, each application of Lemma 8.23 costs $\mathcal{O}((d + |I|)^2 \log\log(|I| + d))$, which is $\mathcal{O}((d^3 + d|I|^2) \log\log(|I| + d))$ in total; the subsequent calls to Lemma 8.21 are dominated by this running time.

Observe that $|\mathsf{span}(\bigcup_{i=1}^t \mathsf{span}(A_i \searrow (\sum_{j=1}^i \delta_j)))| = \mathcal{O}(|I| + d)$ holds at all times, so the initialization of Corollary 9.8 costs $\mathcal{O}((|I| + d)t \log(|I| + d)) = \mathcal{O}((d^2 + d|I|) \log(|I| + d))$ time, whereas each update costs $\mathcal{O}((|I| + d) \log t \log(|I| + d)) = \mathcal{O}((|I| + d) \log^2(|I| + d))$ time. The total number of updates is $\mathcal{O}(|\mathcal{X}| \cdot \delta_E(\mathcal{Y}) + \delta_E(X) \cdot |\mathcal{Y}|) = \mathcal{O}(d^2)$, so this sums up to $\mathcal{O}((d^3 + d^2|I|) \log^2(|I| + d))$. Finally, the total time needed to restrict the returned matrices is $\mathcal{O}(|\mathcal{X}| \times |\mathcal{Y}| \times (|I| + d)) = \mathcal{O}(d^3 + d^2|I|)$. ⌟

⌑ **Lemma 9.11**. *Let $U_1, \ldots, U_z$ and $V_1, \ldots, V_z$ denote Δ-puzzles with values U and V, respectively. Moreover, consider $k, w \in \mathbb{Z}_{\geq 0}$ such that*

$$k + \sum_{i=1}^z \big||U_i| - |V_i|\big| \leq w \leq \Delta/2.$$

Further, define strings U_i' and intervals I_i so that

- $U_1' = U_1[\,0\mathbin{.\,.}|U_1| + w - \Delta\,)$ *and* $I_1 = [\,-k\mathbin{.\,.}w\,]$;
- $U_i' = U_i[\,w\mathbin{.\,.}|U_i| + w - \Delta\,)$ *and* $I_i = [\,0\mathbin{.\,.}\Delta\,]$ *for* $i \in [\,2\mathbin{.\,.}z-1\,]$;
- $U_z' = U_z[\,w\mathbin{.\,.}|U_z|\,)$ *and* $I_z = [\,0\mathbin{.\,.}\Delta\,]$.

Then, we have

$$P_{V,U}|_{[\,-k\mathbin{.\,.}|V|-|U|+k\,]} = \Big(\,\boxdot_{i=1}^{z}\Big((P_{V_i,U_i'}|_{I_i}) \searrow \Big(\sum_{j=1}^{i-1}(|V_j| - \Delta - |U_j'|)\Big)\Big)\Big)\Big|_{[\,-k\mathbin{.\,.}|V|-|U|+k\,]}.$$

Proof. Observe that $U_i' = U[\,u_i\mathbin{.\,.}u_i'\,)$, where $u_i = \sum_{j=1}^{i-1}|U_j'|$ and $u_i' = \sum_{j=1}^{i}|U_j'|$. Moreover, $V_i = V[\,v_i\mathbin{.\,.}v_i'\,)$, where $v_i = \sum_{j=1}^{i-1}(|V_i| - \Delta)$ and $v_i' = \Delta + \sum_{j=1}^{i}(|V_i| - \Delta)$. By Lemma 9.6, it suffices to prove that $v_i \le \max(0, u_i - k)$, $v_i' \ge \min\{|V|, u_i' + |V| - |U| + k\}$, and $I_i \supseteq [\,-k\mathbin{.\,.}|V| - |U| + k + 1\,) - v_i + u_i$ hold for $i \in [\,1\mathbin{.\,.}z\,]$.

As for the first inequality, we have $v_i = 0$ if $i = 1$. Otherwise, $u_i = w + \sum_{j=1}^{i-1}(|U_i| - \Delta)$ and $v_i = \sum_{j=1}^{i-1}(|V_i| - \Delta)$, so $v_i - u_i = \sum_{j=1}^{i-1}(|V_i| - |U_i|) - w \le w - k - w = -k$ holds as claimed. As for the second inequality, we have $v_i' = |V|$ if $i = z$. Otherwise, $u_i' = w + \sum_{j=1}^{i}(|U_i| - \Delta)$ and $v_i' = \Delta + \sum_{j=1}^{i}(|V_i| - \Delta)$, so $v_i' - u_i' = \sum_{j=1}^{i}(|V_i| - |U_i|) + \Delta - w = |V| - |U| - (\sum_{j=i+1}^{z}|V_i| - |U_i|) + \Delta - w \ge |V| - |U| - w + k + \Delta - w \ge |V| - |U| + k$ holds as claimed. Next, we need to prove that $I + u_i - v_i \subseteq I_i$. This is true for $i = 1$, when $u_i = v_i = 0$ and $I = [\,-k\mathbin{.\,.}|V| - |U| + k + 1\,] \subseteq [\,-k\mathbin{.\,.}w\,] = I$. Otherwise, $u_i - v_i = \sum_{j=1}^{i-1}(|U_i| - |V_i|) + w \ge k - w + w \ge k$ and

$$u_i - v_i = \sum_{j=1}^{i-1}(|U_i| - |V_i|) + w = |U| - |V| + \sum_{j=i}^{z}(|V_i| - |U_i|) + w \le |U| - |V| + w - k + w = |U| - |V| + 2w - k,$$

that is, $u_i - v_i \subseteq [\,k\mathbin{.\,.}|U| - |V| + 2w - k\,]$. Hence, so $[\,-k\mathbin{.\,.}|V| - |U| + k\,] + v_i - u_i \subseteq [\,0\mathbin{.\,.}2w\,] \subseteq [\,0\mathbin{.\,.}\Delta\,]$ holds as claimed. ∎

Lemma 9.12. *Consider an instance of the* DYNAMICPUZZLEMATCHING *problem. We can maintain a permutation matrix* A *such that* $A = P_{V,U}|_{[\,-k\mathbin{.\,.}|V|-|U|+k+1\,]}$ *holds whenever* $U_1, \ldots, U_z$ *and* $V_1, \ldots, V_z$ *are* Δ-*puzzles with values* U *and* V, *respectively.*

- *The preprocessing of each family* $\mathcal{S}_i$ *costs* $\mathcal{O}((d_i^3 + d_i\Delta^2)\log^2(d_i + \Delta))$ *time, where* $d_i = \min\{1, \delta_E(\mathcal{S}_i)\}$,
- *the initialization of* $\mathcal{I}$ *costs* $\mathcal{O}(z\Delta\log\Delta)$ *time, and*
- *the updates of* $\mathcal{I}$ *cost* $\mathcal{O}(\Delta\log z\log\Delta)$ *time.*

Proof. Set $w := \lfloor\Delta/2\rfloor$. The preprocessing consists of the following steps:
1. We build $\mathcal{U}_\beta = \{S[\,0\mathbin{.\,.}|S|+w-\Delta\,) : S \in \mathcal{S}_\beta\}$ and process $(\mathcal{U}_\beta, \mathcal{S}_\beta)$ using Lemma 9.10 for $I = [\,-k\mathbin{.\,.}w\,]$.
2. We build $\mathcal{U}_\mu = \{S[\,w\mathbin{.\,.}|S|+w-\Delta\,) : S \in \mathcal{S}_\mu\}$ and process $(\mathcal{U}_\mu, \mathcal{S}_\mu)$ using Lemma 9.10 for $I = [\,0\mathbin{.\,.}\Delta\,]$.
3. We build $\mathcal{U}_\varphi = \{S[\,w\mathbin{.\,.}|S|\,) : S \in \mathcal{S}_\varphi\}$ and process $(\mathcal{U}_\varphi, \mathcal{S}_\varphi)$ using Lemma 9.10 for $I = [\,0\mathbin{.\,.}\Delta\,]$.

At initialization time, we initialize the data structure of Lemma 9.5 with $A_i = P_{V_i,U_i'}|_{I_i}$ for $i \in [\,1\mathbin{.\,.}z\,]$, $\delta_1 = 0$, and $\delta_i = |U_{i-1}'| - |V_{i-1}| - \Delta$ for $i \in [\,2\mathbin{.\,.}z\,]$, where U_i' and I_i are defined in Lemma 9.11. At update time, we update the sequence A_i and δ_i accordingly. At query time, we retrieve the matrix $A := \boxdot_{i=1}^{t}(A_i \searrow \sum_{j=1}^{i}\delta_j)$ from the dynamic algorithm of Lemma 9.5, and return $A|_{[\,-k\mathbin{.\,.}|V|-|U|+k\,]}$, retrieved using Lemma 8.21.

As for correctness, we note that all the possible matrices $A_i = P_{V_i,U_i'}|_{I_i}$ have been constructed at during the preprocessing phase. Moreover, Lemma 9.11 guarantees that $A|_{[\,-k\mathbin{.\,.}|V|-|U|+k\,]} = P_{V,U}|_{[\,-k\mathbin{.\,.}|V|-|U|+k\,]}$ provided that $U_1, \ldots, U_z$ and $V_1, \ldots, V_z$ are Δ-puzzles with values U and V, respectively.

It remains to analyze the running time. Due to $\delta_E(\mathcal{U}_i) \le \delta_E(\mathcal{S}_i)$, preprocessing each family $\mathcal{S}_i$ using Lemma 9.10 costs $\mathcal{O}((d_i^3 + d_i\Delta^2)\log^2(d_i + \Delta))$ time.

Observe that $|\mathsf{span}(\bigcup_{i=1}^{t}\mathsf{span}(A_i \searrow (\sum_{j=1}^{i} \delta_j)))| = \mathcal{O}(\Delta)$ holds at all times, so the initialization of Corollary 9.8 costs $\mathcal{O}(\Delta z \log \Delta)$ time, whereas each update costs $\mathcal{O}(\Delta \log z \log \Delta)$ time. Each query costs $\mathcal{O}(\Delta)$ time. ◾

Lemma 9.13. *Consider strings U, V and an integer $k \in \mathbb{Z}_{\geq 0}$ such that $I := [-k \mathinner{..} |V| - |U| + k]$ is non-empty. Given $P_{V,U}|_I$, the set $\mathrm{Occ}_k^D(U, V)$ can be constructed in $\mathcal{O}(|I| \log |I| / \log \log |I|)$ time.*

Proof. By Fact 9.1, we have $\delta_D(V[i \mathinner{..} j), U) = D_{V,U}[i, j - |U|]$. Moreover, Claim 8.17 yields

$$D_{M(V,U)|_I}[i, j - |U|]$$
$$= \min\{D_{V,U}[i, j - |U|], \quad i + j - |U| + 2(k+1), \quad 2(|V| - |U| + k + 1) - i - j + |U|\}$$
$$= \min\{D_{V,U}[i, j - |U|], \quad 2k + 2 + i + j - |U|, \quad 2k + 2 + 2|V| - |U| - i - j\}.$$

If $j \geq |U| - k$, then

$$2k + 2 + i + j - |U| \geq k + 2 + i > k.$$

If $i \leq |V| - |U| + k$, then

$$2(|V| - |U| + k + 1) - i - j + |U| > |V| - j + k + 2 > k.$$

Hence, for $(i, j) \in [0 \mathinner{..} |V| - |U| + k] \times [|U| - k \mathinner{..} |V|]$, we have $\delta_D(U, V[i \mathinner{..} j)) \leq k$ if and only if $D_{M(V,U)|_I}[i, j - |U|] \leq k$. On the other hand $\delta_D(U, V[i \mathinner{..} j)) \geq |j - i - |U|| > k$ holds whenever $j < |U| - k$ or $i > |V| - |U| + k$. Hence, our task reduces to checking, for each $i \in [0 \mathinner{..} |V| - |U| + k]$, whether $D_{M(V,U)|_I}[i, j - |U|] \leq k$ holds for some $j \in [|U| - k \mathinner{..} |V|]$. For this, we recall that $D_{M(V,U)|_I}$ is a Monge matrix and note that it remains a Monge matrix when restricted to $[0 \mathinner{..} |V| - |U| + k] \times [-k \mathinner{..} |V| - |U|]$. The SMAWK algorithm [AKM+87] finds row-minima in an $m \times m$ Monge matrix using $\mathcal{O}(m)$ queries asking for values of the matrix entries. By Fact 8.1 and Lemma 8.11, after $\mathcal{O}(|I|\sqrt{\log |I|})$-time preprocessing, we have $\mathcal{O}(\log |I| / \log \log |I|)$-time access to entries of $D_{M(V,U)|_I}$. Since $m = |V| - |U| + k + 1 \leq |I|$, the final running time is $\mathcal{O}(|I| \log |I| / \log \log |I|)$. ◾

Proposition 9.14. *There is a data structure for a DYNAMICPUZZLEMATCHING$(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ problem variant, reporting $\mathrm{Occ}_k^D(U, V)$ instead of $\mathrm{Occ}_k^E(U, V)$, with $\mathcal{O}(\Delta \log z \log \Delta)$-time updates and queries, $\mathcal{O}(\Delta z \log \Delta)$-time initialization, and $\mathcal{O}((d^3 + \Delta^2 d) \log^2(d + \Delta))$-time preprocessing, where $d = \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi)$.*

Proof. Let $w = \lfloor \Delta/2 \rfloor$. We use Lemma 9.12 to maintain a permutation matrix A such that $A = P_{V,U}|_{[-k \mathinner{..} |V| - |U| + k]}$ holds whenever $U_1, \ldots, U_z$ and $V_1, \ldots, V_z$ are Δ-puzzles with values U, V, respectively.

It remains to implement queries. If $|V| < |U| - k$, then we report that $\mathrm{Occ}^D(U, V) = \varnothing$. Otherwise, we derive $A = P_{V,U}|_{[-k \mathinner{..} |V| - |U| + k]}$. Next, we compute $\mathrm{Occ}_k^D(U, V)$ using Lemma 9.13. These two steps add $\mathcal{O}(w \log w / \log \log w)$ to the query time, which is dominated by the update time $\mathcal{O}(w \log z \log w)$. ◾

Theorem 1.5. *There is a data structure for DYNAMICPUZZLEMATCHING$(k, \Delta, \mathcal{S}_\beta, \mathcal{S}_\mu, \mathcal{S}_\varphi)$ with $\mathcal{O}(\Delta \log z \log \Delta)$-time updates and queries, $\mathcal{O}(\Delta z \log \Delta)$-time initialization, and $\mathcal{O}((d^3 + \Delta^2 d) \log^2(d + \Delta))$-time preprocessing, where $d = \delta_E(\mathcal{S}_\beta) + \delta_E(\mathcal{S}_\mu) + \delta_E(\mathcal{S}_\varphi)$.[24]*

---

[24] Recall that z is the length of the DPM-sequence that we maintain in the data structure.

Proof. We apply Proposition 9.14 using the distortion-free embedding from δ_E to δ_D [Tis07]. This embedding is defined by mapping every string $S \in \Sigma^*$ to a string $f(S) \coloneqq \bigodot_{i=0}^{|S|-1} S[i]\$$, where $\$ \notin \Sigma$. Observe that $\delta_E(S,T) = \frac{1}{2}\delta_D(f(S), f(T))$ and $\mathrm{Occ}_k^E(P,T) = \{i \in [0 \mathbin{..} |T| - |P| + k] : 2i \in \mathrm{Occ}_{2k}^D(f(P), f(T))\}$. Thus, we use Proposition 9.14 we the strings mapped through f and the integer parameters k and Δ doubled. Note that PILLAR operations on the family $f(\mathcal{X})$ can be easily implemented using the PILLAR operations on $\mathcal{X}$. ∎

References

Abr87 Karl R. Abrahamson. Generalized string matching. *SIAM Journal on Computing*, 16(6):1039–1051, 1987. `doi:10.1137/0216067`.

AHNR98 Arne Andersson, Torben Hagerup, Stefan Nilsson, and Rajeev Raman. Sorting in linear time? *Journal of Computer and System Sciences*, 57(1):74–93, 1998. `doi:10.1006/jcss.1998.1580`.

AKM⁺87 Alok Aggarwal, Maria M. Klawe, Shlomo Moran, Peter W. Shor, and Robert E. Wilber. Geometric applications of a matrix-searching algorithm. *Algorithmica*, 2:195–208, 1987. `doi:10.1007/BF01840359`.

ALP04 Amihood 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`.

BF00 Michael A. Bender and Martin Farach-Colton. The LCA problem revisited. In *LATIN 2000: Theoretical Informatics, 4th Latin American Symposium, Proceedings*, pages 88–94, 2000. `doi:10.1007/10719839_9`.

BGK⁺16 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`.

BI18 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`.

BLR⁺15 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`.

CDK19 Diptarka Chakraborty, Debarati Das, and Michal Koucký. Approximate online pattern matching in sublinear time. In *39th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2019*, pages 10:1–10:15, 2019. `doi:10.4230/LIPIcs.FSTTCS.2019.10`.

CFP⁺16 Raphaël Clifford, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. The k-mismatch problem revisited. In *27th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 2039–2052, 2016. `doi:10.1137/1.9781611974331.ch142`.

CGK⁺20 Timothy M. Chan, Shay Golan, Tomasz Kociumaka, Tsvi Kopelowitz, and Ely Porat. Approximating text-to-pattern hamming distances. In *52nd Annual ACM Symposium on Theory of Computing, STOC 2020*, pages 643–656, 2020. `doi:10.1145/3357713.3384266`.

CH02 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`.

Cha20 Panagiotis Charalampopoulos. *Data Structures for Strings in the Internal and Dynamic Settings*. PhD thesis, King's College London, 2020.

CKM20 Panagiotis Charalampopoulos, Tomasz Kociumaka, and Shay Mozes. Dynamic string alignment. In *31st Annual Symposium on Combinatorial Pattern Matching, CPM 2020*, pages 9:1–9:13, 2020. `doi:10.4230/LIPIcs.CPM.2020.9`.

CKW20 Panagiotis Charalampopoulos, Tomasz Kociumaka, and Philip Wellnitz. Faster approximate pattern matching: A unified approach. In *61st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2020*, pages 978–989, 2020. Full version: arXiv:2004.08350v2. `arXiv:2004.08350v2`, `doi:10.1109/FOCS46700.2020.00095`.

CP10 Timothy M. Chan and Mihai Pătraşcu. Counting inversions, offline orthogonal range counting, and related problems. In *21st Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010*, pages 161–173, 2010. `doi:10.1137/1.9781611973075.15`.

Far97 Martin Farach. Optimal suffix tree construction with large alphabets. In *38th Annual IEEE Symposium on Foundations of Computer Science, FOCS 1997*, pages 137–143, 1997. `doi:10.1109/SFCS.1997.646102`.

GG86 Zvi Galil and Raffaele Giancarlo. Improved string matching with k mismatches. *SIGACT News*, 17(4):52–54, 1986. `doi:10.1145/8307.8309`.

GKK⁺18 Paweł Gawrychowski, Adam Karczmarz, Tomasz Kociumaka, Jakub Łącki, and Piotr Sankowski. Optimal dynamic strings. In *29th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018*, pages 1509–1528, 2018. `arXiv:1511.02612`, `doi:10.1137/1.9781611975031.99`.

GU18 Pawel Gawrychowski and Przemyslaw Uznański. Towards unified approximate pattern matching for Hamming and L_1 distance. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, 2018. `doi:10.4230/LIPIcs.ICALP.2018.62`.

Han04 Yijie Han. Deterministic sorting in $O(n \log \log n)$ time and linear space. *J. Algorithms*, 50(1):96–105, 2004. `doi:10.1016/j.jalgor.2003.09.001`.

I17 Tomohiro I. Longest common extensions with recompression. In *28th Annual Symposium on Combinatorial Pattern Matching, CPM 2017*, volume 78, 2017. `doi:10.4230/LIPIcs.CPM.2017.18`.

Jeż15 Artur Jeż. Faster fully compressed pattern matching by recompression. *ACM Transactions on Algorithms*, 11(3):20:1–20:43, 2015. `doi:10.1145/2631920`.

Jeż16 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`.

KK20 Dominik Kempa and Tomasz Kociumaka. Resolution of the Burrows–Wheeler transform conjecture. In *61st Annual IEEE Symposium on Foundations of Computer Science, FOCS 2020*, 2020.

KK22 Dominik Kempa and Tomasz Kociumaka. Dynamic suffix array with polylogarithmic queries and updates. In *54th Annual ACM Symposium on Theory of Computing, STOC 2022*, 2022. `arXiv:2201.01285`.

Koc18 Tomasz Kociumaka. *Efficient Data Structures for Internal Queries in Texts*. PhD thesis, University of Warsaw, October 2018.

Kos87 S.R. Kosaraju. Efficient string matching. Manuscript, 1987.

KPS21 Tomasz Kociumaka, Ely Porat, and Tatiana Starikovskaya. Small space and streaming pattern matching with k edits. In *62nd Annual Symposium on Foundations of Computer Science, FOCS 2021*, 2021. `arXiv:2106.06037`.

KRRW15 Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, and Tomasz Walen. Internal pattern matching queries in a text and applications. In *26th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2015*, pages 532–551, 2015. `doi:10.1137/1.9781611973730.36`.

KS21 Tomasz Kociumaka and Saeed Seddighin. Improved dynamic algorithms for longest increasing subsequence, 2021. `arXiv:2011.10874`.

LMS98 Gad M. Landau, Eugene W. Myers, and Jeanette P. Schmidt. Incremental string comparison. *SIAM Journal on Computing*, 27(2):557–582, 1998. `doi:10.1137/S0097539794264810`.

LV86 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`.

LV88 Gad M. Landau and Uzi Vishkin. Fast string matching with k differences. *Journal of Computer and System Sciences*, 37(1):63–78, 1988. `doi:10.1016/0022-0000(88)90045-1`.

LV89 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`.

MP80 William J. Masek and Mike Paterson. A faster algorithm computing string edit distances. *Journal of Computer and System Sciences*, 20(1):18–31, 1980. `doi:10.1016/0022-0000(80)90002-1`.

Nav01 Gonzalo Navarro. A guided tour to approximate string matching. *ACM Comput. Surv.*, 33(1):31–88, 2001. `doi:10.1145/375360.375365`.

Sel80 Peter H. Sellers. The theory and computation of evolutionary distances: Pattern recognition. *J. Algorithms*, 1(4):359–373, 1980. `doi:10.1016/0196-6774(80)90016-4`.

Sta17 Tatiana Starikovskaya. Communication and streaming complexity of approximate pattern matching. In *CPM 2017*, pages 13:1–13:11, 2017. `doi:10.4230/LIPIcs.CPM.2017.13`.

SV96 Süleyman Cenk Sahinalp and Uzi Vishkin. Efficient approximate and dynamic matching of patterns using a labeling paradigm (extended abstract). In *37th Annual IEEE Symposium on Foundations of Computer Science, FOCS 1996*, pages 320–328, 1996. `doi:10.1109/SFCS.1996.548491`.

Tho07 Mikkel Thorup. Equivalence between priority queues and sorting. *J. ACM*, 54(6):28, 2007. `doi:10.1145/1314690.1314692`.

Tis07 Alexander Tiskin. Semi-local string comparison: algorithmic techniques and applications, 2007. `arXiv:0707.3619`.

Tis08 Alexandre Tiskin. Semi-local string comparison: Algorithmic techniques and applications. *Math Comput Sci*, 1(4):571–603, 2008. `doi:10.1007/s11786-007-0033-3`.

Tis15 Alexander Tiskin. Fast distance multiplication of unit-monge matrices. *Algorithmica*, 71(4):859–888, 2015. `doi:10.1007/s00453-013-9830-z`.

A   Notation Overview for Part I

Notation	Explanation
i, j	Integer indices.
s, v, w, x, y	Positions in strings, typically occurring in pairs as boundaries of fragments.
F, G, S, U, V, X, Y	Strings or fragments of strings; possibly with additional meaning in their respective contexts.
(a_X, a_Y)	Elements of an alignment of (fragments of) $X \rightsquigarrow Y$. We may deviate from this notation in the special case of alignments of (fragments of) a string to (fragments of) itself.
$T, \|T\| = n$	Text T of length n.
$P, \|P\| = m$	Pattern P of length m.
k	A non-negative threshold k used to denote the allowed number of edits when searching for occurrences of P in T.
$\mathrm{Occ}_k^E(P, T)$	The set of all starting positions of k-edit occurrences of P in T.
The Standard Trick	We may assume $n < \tfrac{3}{2} m + k$ with a running-time overhead of $\mathcal{O}(n/m)$ (see Section 2.2).
$Q, \|Q\| = q$	Primitive string Q of length q.
$\mathcal{A}_T : T \rightsquigarrow Q^\infty[\, x_T \mathbin{..} y_T \,), d_T$	An alignment of T onto a fragment of Q^∞ with cost d_T. We have $x_T \in [\, 0 \mathbin{..} q \,)$.
$\mathcal{A}_P : P \rightsquigarrow Q^\infty[\, 0 \mathbin{..} y_P \,), d_P$	An alignment of P onto a fragment of Q^∞ with cost d_P.
$\kappa := d_P + k + d_T$	For an alignment $\mathcal{A} : P \rightsquigarrow T[\, v \mathbin{..} w \,)$ of cost at most k, the bound κ is an upper bound on the number of edits of the induced alignment $\mathcal{A}_T \circ \mathcal{A} \circ \mathcal{A}_P^{-1} : Q^\infty[\, 0 \mathbin{..} y_P \,) \rightsquigarrow Q^\infty[\, v' \mathbin{..} w' \,)$.
$\tau := q \lceil \kappa/2q \rceil$	Auxiliary parameter.
d	Positive integer that satisfies $d_T/3 \le d$ and $d_P \le d$ and $2k \le d$ and $8q/m \le d$.
J	See Definition 3.6; we have $J = [\, \lceil (x_T - \kappa)/\tau \rceil \mathbin{..} \lfloor (y_T - y_P + \kappa)/\tau \rfloor \,]$.
$R_j := T[\, r_j \mathbin{..} r_j' \,)$	See Definition 3.6; we have $r_j := \min\{a_T : (a_T, a_Q) \in \mathcal{A}_T \text{ for } j\tau - \kappa \le a_Q\}$ and $r_j' := \max\{a_T : (a_T, a_Q) \in \mathcal{A}_T \text{ for } a_Q \le j\tau + y_P + \kappa\}$. Furthermore, $r_{\max J+1} := n$.
$S = \bigodot_{i=1}^{\beta_S} S[\, s_{i-1} \mathbin{..} s_i \,)$	Tile partition of a string S with respect to some alignment $\mathcal{A}_S : S \rightsquigarrow Q^\infty[\, x_S \mathbin{..} y_S \,)$ with $S[\, s_{i-1} \mathbin{..} s_i \,) = \mathcal{A}_S^{-1}(Q^\infty[\, \max\{x_S, (j-1)\tau\} \mathbin{..} \min\{y_S, j\tau\} \,))$.
$\beta_S := \lceil y_S/\tau \rceil$	The number of the last non-zero tile of S.
$P = \bigodot_{i=1}^{\beta_P} P[\, p_{i-1} \mathbin{..} p_i \,)$	τ-tile partition of P. We assume that $\beta_P \ge 20$ (by Lemma 4.5).
$T = \bigodot_{i=1}^{\beta_T} T[\, t_{i-1} \mathbin{..} t_i \,)$	τ-tile partition of T.
$\Delta := 6\kappa$	Overlap of neighboring puzzle pieces.
$z := \beta_P - 17$	Number of puzzle pieces.
$P_1, \ldots, P_z$	A Δ-puzzle with value P, where $P_1 := P[\, p_0 \mathbin{..} p_2 + \Delta \,)$ and
$\mathsf{val}_\Delta(P_1, \ldots, P_z) = P$	$P_i := P[\, p_i \mathbin{..} p_{i+1} + \Delta \,)$ for $i \in (\, 1 \mathbin{..} z \,)$ and $P_z := P[\, p_z \mathbin{..} \|P\| \,)$.
$T_{j,1}, \ldots, T_{j,z}$	A Δ-puzzle with value R_j, where $T_{j,1} := T[\, r_j \mathbin{..} t_{j+2} + \Delta \,)$ and
$\mathsf{val}_\Delta(T_{j,1}, \ldots, T_{j,z}) = R_j$	$T_{j,i} := T[\, t_{j+i} \mathbin{..} t_{j+i+1} + \Delta \,)$ for $i \in (\, 1 \mathbin{..} z \,)$ and $T_{j,z} := T[\, t_{j+z} \mathbin{..} r_j' \,)$.
$T_{\min J+2}, \ldots, T_i, \ldots, T_{\max J+z-1}$	Shorthand for the pieces $T_{j,i'}$ with $j + i' = i$ and $i \in (\, 1 \mathbin{..} z \,)$; that is $T_i := T[\, t_i \mathbin{..} t_{i+1} + \Delta \,)$.

$\mathcal{S}_\beta := \{P_1\} \cup \{T_{j,1} : j \in J\}$	The set of leading puzzle pieces.
$\mathcal{S}_\mu := \{P_i : i \in (1 .. z)\} \cup$ $\{T_i : i \in (\min J + 1 .. \max J + z)\}$	The set of internal puzzle pieces.
$\mathcal{S}_\varphi := \{P_z\} \cup \{T_{j,z} : j \in J\}$	The set of trailing puzzle pieces.
$\delta_E(\mathcal{S}) := \min_{\hat{S} \in \Sigma^*} \sum_{S \in \mathcal{S}} \delta_E(S, \hat{S})$	The median edit distance of the family $\mathcal{S}$.
$\mathsf{Special}(T), \mathsf{Special}(P)$	The special internal pieces in T and P, that is, the pieces in $\mathcal{S}_\mu$ that differ from $Q^\infty[0 .. \tau + \Delta)$.
$\mathcal{I} = (U_1, V_1)(U_2, V_2) \ldots (U_z, V_z)$	Sequence of ordered pairs of strings, a *DPM-sequence*. Occasionally, we call the elements (U_i, V_i) a *DPM-pair*.
$\mathrm{Occ}_k^E(\mathcal{I}) := \mathrm{Occ}_k^E(U, V)$	Set of k-error occurrences of the string $U := \mathsf{val}_\Delta(U_1, \ldots, U_z)$ in the string $V := \mathsf{val}_\Delta(V_1, \ldots, V_z)$.
$\mathrm{tor}(\mathcal{I}) := \sum_{i=1}^z \big| |U_i| - |V_i| \big|$	The torsion of the DPM-sequence $\mathcal{I}$.
$\mathcal{I}_j := (P_1, T_{j,1})(P_2, T_{j,2}) \cdots (P_z, T_{j,z})$	DPM-sequence representing the strings $P = \mathsf{val}_\Delta(P_1, \ldots, P_z)$ and $R_j = \mathsf{val}_\Delta(T_{j,1}, \ldots, T_{j,z})$; we have $\mathrm{tor}(\mathcal{I}_j) \leq 3\kappa - k = \Delta/2 - k$ by Lemma 4.10.
$\mathcal{Q} := (Q^\infty[0 .. \tau + \Delta),$ $Q^\infty[0 .. \tau + \Delta))$	The plain (internal) DPM-pair.
$\mathcal{L}^P, \partial_P, L_1^P$	The set of locked fragments of P that contains a ∂_P-locked prefix L_1^P, computed using $\mathtt{Locked}(P, Q, d_P, \partial_P)$.
$\mathcal{L}^T$	The set of locked fragments of T, computed using $\mathtt{Locked}(T, Q, d_T, 0)$.
$\hat{\kappa}$	A slack allowance for computing overlaps of fragments for the purpose of marking.
$\mathrm{mk}(v, \hat{\kappa}, L^P, L^T)$	The number of marks given to a position v due to the (possible) overlap of $L^P \in \mathcal{L}^P$ with $L^T \in \mathcal{L}^T$.
$\mathrm{mk}(v) = \mathrm{mk}(v, \hat{\kappa}, \mathcal{L}^P, \mathcal{L}^T)$	The total number of marks given to a position v of T.
$\mathrm{mk}(v, L)$	The number of marks given to a position v of T due to overlaps of pairs of locked fragments containing locked fragment L.
$\eta, \mathbb{H}, \mathbb{L}$	A positive integer threshold used in the partition of the positions $[0 .. n - m + k]$ of T to a set $\mathbb{H}$ of heavy positions and a set $\mathbb{L}$ of light positions; see Definition 5.8.
$\rho(\star)$	A mapping from positions of T to integers in $[x_T .. y_T)$ such that $\mathcal{A}_T(T[v .. n)) = Q^\infty[\rho(v) .. y_T)$.
$\mathcal{D}(v)$	For a light position v of T, the set $\{L_1^P\} \cup \{L \in \mathcal{L}^P \cup \mathcal{L}_{[v .. v+m)}^T : \mathrm{mk}(v, L) < \delta_E(L, {}^*Q^*)\}$.

B   Notation Overview for Part II

Notation	Explanation
$I^+, {}^+I, I^-, {}^-I$	An interval I extended or shrunk by one element to the left or right.
$\mathsf{span}(S)$	The smallest integer interval containing a set $S \subseteq \mathbb{Z}$.
$A \odot B$	The min-plus product of matrices A, B.
$A^\square$	The density matrix of a matrix A.
$\mathsf{span}(A)$	The span of a permutation a matrix A, defined as ${}^-\mathsf{span}(\{i \in I : A[i,i] = 0\})$.
A^Σ	The distribution matrix of a matrix A.
$A \boxdot B$	The seaweed product of permutation matrices A, B; defined as $(A^\Sigma \odot B^\Sigma)^\square$.
$A \searrow s$	The diagonal shift of a matrix A by s units; $(A \searrow s)[i+s, j+s] = A[i,j]$.
M	A finite subset of $\mathbb{Z}^2$.
$\mathsf{AG}(M)$	The alignment graph of M.
dist_M	The distance function on $\mathbb{Z}^2$ induced by distances in $\mathsf{AG}(M)$.
$\prec$	A partial order on $\mathbb{Z}^2$ defined so that $(x,y) \prec (x',y')$ if and only if $x < x'$ and $y < y'$.
$\mathsf{LIS}(S)$	The maximum length of a $\prec$-chain within $S \subseteq \mathbb{Z}^2$.
$\mathsf{lt}^B, \mathsf{br}^B$	The left-top and the bottom-right boundary of a bounding box B.
$D_{M,B}, D_M$	The distance matrix of M with respect to a bounding box B (shown to be independent of B).
P_M	The seaweed matrix of M, defined as $\frac{1}{2} D_M^\square$.
$\mathsf{span}(M)$	The span of M, defined as $\mathsf{span}(\{x - y : (x,y) \in M\})$.
M_I	The restriction of M to interval I, defined as $\{(x,y) \in M : x - y \in I\}$.
A_I	The restriction of a permutation matrix A to interval I, defined as $P_{M|_I}$ if $A = P_M$ (shown to be independent of M).
$M(X,Y)$	The set of mismatches between two strings X, Y, defined as $\{(x,y) : X[x] \neq Y[y]\}$.
$D_{X,Y}, P_{X,Y}$	The distance matrix and the permutation matrix of $M(X,Y)$.
$\delta_D(X,Y)$	The deletion distance between strings X, Y.
$\delta_D(\mathcal{X})$	The median deletion distance of a finite string family $\mathcal{X}$.