

String Synchronizing Sets*

Sublinear-Time BWT Construction and Optimal LCE Data Structure

Dominik Kempa

Department of Computer Science
Centre for Discrete Mathematics and its Applications
University of Warwick
Coventry, United Kingdom
dominik.kempa@warwick.ac.uk

Tomasz Kociumaka

Department of Computer Science
Bar-Ilan University
Ramat Gan, Israel
Institute of Informatics
University of Warsaw
Warsaw, Poland
kociumaka@mimuw.edu.pl

ABSTRACT

Burrows–Wheeler transform (BWT) is an invertible text transformation that, given a text T of length n , permutes its symbols according to the lexicographic order of suffixes of T . BWT is one of the most heavily studied algorithms in data compression with numerous applications in indexing, sequence analysis, and bioinformatics. Its construction is a bottleneck in many scenarios, and settling the complexity of this task is one of the most important unsolved problems in sequence analysis that has remained open for 25 years. Given a binary string of length n , occupying $O(n/\log n)$ machine words, the BWT construction algorithm due to Hon et al. (SIAM J. Comput., 2009) runs in $O(n)$ time and $O(n/\log n)$ space. Recent advancements (Belazzougui, STOC 2014, and Munro et al., SODA 2017) focus on removing the alphabet-size dependency in the time complexity, but they still require $\Omega(n)$ time. Despite the clearly suboptimal running time, the existing techniques appear to have reached their limits.

In this paper, we propose the first algorithm that breaks the $O(n)$ -time barrier for BWT construction. Given a binary string of length n , our procedure builds the Burrows–Wheeler transform in $O(n/\sqrt{\log n})$ time and $O(n/\log n)$ space. We complement this result with a conditional lower bound proving that any further progress in the time complexity of BWT construction would yield faster algorithms for the very well studied problem of counting inversions: it would improve the state-of-the-art $O(m\sqrt{\log m})$ -time solution by Chan and Pătraşcu (SODA 2010). Our algorithm is based on a novel concept of string synchronizing sets, which is of independent interest. As one of the applications, we show that this technique lets us design a data structure of the optimal size $O(n/\log n)$ that answers Longest Common Extension queries (LCE queries) in $O(1)$ time and, furthermore, can be deterministically constructed in the optimal $O(n/\log n)$ time.

*Dominik Kempa was supported by the EPSRC award EP/N011163/1. Tomasz Kociumaka was supported by ISF grants no. 824/17 and 1278/16 and by an ERC grant MPM under the EU's Horizon 2020 Research and Innovation Programme (grant no. 683064).

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

STOC '19, June 23–26, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6705-9/19/06...\$15.00

<https://doi.org/10.1145/3313276.3316368>

CCS CONCEPTS

• **Theory of computation** → **Pattern Matching; Data compression; Data structures design and analysis**; Cell probe models and lower bounds; Problems, reductions and completeness.

KEYWORDS

Burrows–Wheeler transform, Longest Common Extension queries, Longest Common Prefix queries, packed strings

ACM Reference Format:

Dominik Kempa and Tomasz Kociumaka. 2019. String Synchronizing Sets: Sublinear-Time BWT Construction and Optimal LCE Data Structure. In *Proceedings of the 51st Annual ACM SIGACT Symposium on the Theory of Computing (STOC '19)*, June 23–26, 2019, Phoenix, AZ, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3313276.3316368>

1 INTRODUCTION

The problem of text indexing is to preprocess an input text T so that given any query pattern P , we can quickly find the occurrences of P in T (typically in $O(|P| + \text{occ})$ time, where $|P|$ is the length of P and occ is the number of reported occurrences). Two classical data structures for this task are the suffix tree [42] and the suffix array [33]. The suffix tree is a trie containing all suffixes of T with each unary path compressed into a single edge labeled by a text substring. The suffix array is a list of suffixes of T in the lexicographic order, with each suffix encoded using its starting position. Both data structures take $\Theta(n)$ words of space, where n is the length of T . In addition to indexing, they underpin dozens of applications in bioinformatics, data compression, and information retrieval [2, 17]. While the suffix tree is slightly faster for some operations, the suffix array is often preferred due to its simplicity and lower space usage.

Nowadays, however, indexing datasets of size close to the capacity of available RAM is often required. Even the suffix arrays are then prohibitively large, particularly in applications where the text consists of symbols from some alphabet Σ of small size $\sigma = |\Sigma|$ (e.g., $\Sigma = \{A, C, G, T\}$ and so $\sigma = 4$ in bioinformatics). For such collections, the classical indexes are $\Theta(\log_{\sigma} n)$ times larger than the text, which takes only $\Theta(n \log \sigma)$ bits, i.e., $\Theta(n/\log_{\sigma} n)$ machine words, and thus they prevent many sequence analysis tasks to be performed without a significant penalty in space consumption.

This situation changed dramatically in early 2000's, when Ferragina and Manzini [13], as well as Grossi and Vitter [16], independently proposed indexes with the capabilities of the suffix array

(incurring only a $O(\log^\epsilon n)$ slowdown in the query time) that take a space asymptotically equal to that of the text (and with very small constant factors). These indexes are known as the FM-index and the compressed suffix array (CSA). The central component and the time and space bottleneck in the construction of both the FM-index and CSA¹ is the *Burrows–Wheeler transform* (BWT) [8]. BWT is an invertible permutation of the text that consists of symbols preceding suffixes of text in the lexicographic order. Almost immediately after their discovery, the BWT-based indexes replaced suffix arrays and suffix trees and the BWT itself has become the basis of almost all space-efficient algorithms for sequence analysis. Modern textbooks spend dozens of pages describing its applications [32, 38, 39], and BWT-indexes are widely used in practice; in bioinformatics, they are the central component of many read-aligners [30, 31].

BWT Construction. Given the practical importance of BWT, its efficient construction emerged as one of the most important open problems in the field of indexing and sequence analysis. The first breakthrough was the algorithm of Hon et al. [21], who reduced the time complexity of BWT construction for binary strings from $O(n \log n)$ to $O(n)$ time using working space of $O(n)$ bits. This bound has been recently generalized to any alphabet size σ . More precisely, Belazzougui [4] described a (randomized) $O(n)$ -time construction working in optimal space of $O(n/\log_\sigma n)$ words. Munro et al. [34] then proposed an alternative (and deterministic) construction. These algorithms achieve the optimal construction space, but their running time is still $\Omega(n)$, which is up to $\Theta(\log n)$ times more than the lower bound of $\Omega(n/\log_\sigma n)$ time (required to read the input and write the output). Up until now, all $o(n)$ -time algorithms required additional assumptions, such as that the BWT is highly compressible using run-length encoding [25].

In this paper, we propose the first algorithm that always breaks the $O(n)$ -time barrier for BWT construction. Given a binary string of length n , our algorithm builds the Burrows–Wheeler transform in $O(n/\sqrt{\log n})$ time and $O(n/\log n)$ space. We complement this result with a conditional lower bound proving that any further progress in the time complexity of BWT construction would imply faster algorithms for the very well studied problem of counting inversions: it would improve the state-of-the-art $O(m\sqrt{\log m})$ -time solution by Chan and Pătraşcu [9]. We also generalize our construction to larger alphabets whose size σ satisfies $\log \sigma \leq \sqrt{\log n}$. In this case, the running time is $O(n \log \sigma / \sqrt{\log n})$ and the space complexity is $O(n \log \sigma / \log n)$, proportional to the input and output size.

LCE Queries. The *Longest Common Extension* queries $LCE(i, j)$ (also known as the Longest Common Prefix queries), given two positions in a text T , return the length of the longest common prefix of the suffixes $T[i..n]$ and $T[j..n]$ starting at positions i and j , respectively. These queries were introduced by Landau and Vishkin [29] in the context of approximate pattern matching. Since then, they became one of the most commonly used tools in text processing. Standard data structures answer LCE queries in constant time and take linear space. The original construction algorithm [20, 29, 42] works in linear time for constant alphabets only, but it has been subsequently generalized to larger integer alphabets [12] and simplified

substantially [6, 24]. Thus, LCE queries are completely resolved in the classic setting where the text T is stored in $O(n)$ space.

However, if T is over a small alphabet of size σ , then it can be stored in $O(n \log \sigma)$ bits. Yet, until very recently, even for the binary alphabet there was no data structure of $o(n \log n)$ bits supporting LCE queries in constant time. The first such solutions are by Tanimura et al. [41] and Munro et al. [35], who showed that constant-time queries can be implemented using data structures of size $O(n \log \sigma / \sqrt{\log n})$ and $O(n\sqrt{\log \sigma} / \sqrt{\log n})$, respectively. The latter result admits an $O(n/\sqrt{\log_\sigma n})$ -time construction from the packed representation of T . In yet another study, Birenzweige et al. [7] considered LCE queries in a model where T is available for read-only random access, but not counted towards the data structure size. Constant-time LCE queries in the optimal space of $O(n \log \sigma)$ bits can be deduced as a corollary of their results, but the construction algorithm is randomized and takes $O(n)$ time.

Our contribution in the area of LCE queries is a data structure of the optimal size $O(n/\log_\sigma n)$ that answers LCE in $O(1)$ time and, furthermore, can be deterministically constructed in the optimal $O(n/\log_\sigma n)$ time. This significantly improves the state of the art and essentially closes the LCE problem also in the packed setting.

Our Techniques. Our main innovation and the key tool behind both our results is a novel notion of *string synchronizing sets*, which relies on *local consistency*—the idea to make symmetry-breaking decisions involving a position i of the text T based on the characters at the nearby positions. This way, we can guarantee that equal fragments of the text are handled in the same way. The classic implementations of local consistency involve parsing the text; see e.g. [23, 40]. Unfortunately, the context size at a given level of the parsing is expressed in terms of the number of phrases, whose lengths may vary significantly between regions of the text. To overcome these limitations, Kociumaka et al. [28] introduced *samples assignments* with fixed context size. Birenzweige et al. [7] then applied the underlying techniques to define *partitioning sets*, which they used for answering LCE queries. Moreover, they obtained an alternative construction of partitioning sets (with slightly inferior properties) by carefully modifying the parsing scheme of [40]. In his PhD thesis [27], the second author introduced *synchronizing functions*, an improved version of *samples assignments* with stronger properties and efficient deterministic construction procedures. He also used synchronizing functions to develop the optimal LCE data structure in a packed text. In this work, we reproduce the latter result using *synchronizing sets*, which are closely related to synchronizing functions, but enjoy a much simpler and cleaner interface.

Organization of the Paper. After introducing the basic notation and tools in Section 2, we start by defining the main concept of the paper—the string synchronizing set—and proving some of its properties (Section 3). Next, we show how to sort suffixes in such a set (Section 4) and extend these ideas into an optimal LCE data structure (Section 5). We then describe how to build the BWT given a small string synchronizing set (Section 6) and prove the conditional optimality of our construction (Section 7). We conclude by showing efficient algorithms for the construction of string synchronizing set (Section 8). Due to space limitations, proofs of the claims marked with ♠ are presented only in the full version of the paper [26].

¹Although originally formulated in terms of the so-called “ Ψ function” [16], it is now established (see, e.g., [21, 34]) that the CSA is essentially dual to the FM-index.

2 PRELIMINARIES

Let $T \in \Sigma^*$ be a string over alphabet $\Sigma = [0 \dots \sigma - 1]$. Unless explicitly stated otherwise, we assume $\sigma = n^{O(1)}$, where $n = |T|$. For $1 \leq i \leq j \leq n$, we write $T[i \dots j]$ to denote the *substring* $T[i]T[i+1] \dots T[j]$. Throughout, we use $[i \dots j]$ as a shorthand for $[i \dots j - 1]$. The length of the longest common prefix of $X, Y \in \Sigma^*$ is denoted $\text{lcp}(X, Y)$.

An integer $p \in [1 \dots |X|]$ is a *period* of X if $X[i] = X[i + p]$ for $i \in [1 \dots |X| - p]$. The shortest period of X is denoted as $\text{per}(X)$.

LEMMA 2.1 (PERIODICITY LEMMA [14]). *If a string X has periods p, q such that $p + q - \text{gcd}(p, q) \leq |X|$, then $\text{gcd}(p, q)$ is also its period.*

2.1 Suffix Array and BWT

The *suffix array* [33] $\text{SA}[1 \dots n]$ of a text T is a permutation defining the lexicographic order on suffixes: $T[\text{SA}[i] \dots n] < T[\text{SA}[j] \dots n]$ if $i < j$. It takes $O(n)$ space and can be constructed in $O(n)$ time [24].

Given positions i, j in T , the *Longest Common Extension* query $\text{LCE}(i, j)$ asks for $\text{lcp}(T[i \dots n], T[j \dots n])$. The standard solution consists of the suffix array SA , the inverse permutation SA^{-1} (defined so that $\text{SA}[\text{SA}^{-1}[i]] = i$), the LCP table $\text{LCP}[2 \dots n]$ (whose entries are $\text{LCP}[i] = \text{LCE}(\text{SA}[i - 1], \text{SA}[i])$), and a data structure for range minimum queries built on top of the LCP table; see [6, 12, 20, 24].

PROPOSITION 2.2. *LCE queries in a text $T \in [0 \dots \sigma]^n$ with $\sigma = n^{O(1)}$ can be answered in $O(1)$ time after $O(n)$ -time preprocessing.*

The *Burrows–Wheeler transform* (BWT) [8] of $T[1 \dots n]$ is defined as $\text{BWT}[i] = T[\text{SA}[i] - 1]$ if $\text{SA}[i] > 1$ and $\text{BWT}[i] = T[n]$ otherwise. To ensure the correct handling of boundary cases, it is often assumed that $\text{BWT}[\text{SA}^{-1}[1]]$ contains a sentinel $\$ \notin \Sigma$. In this paper, we avoid this to make sure that $\text{BWT}[1 \dots n] \in [0 \dots \sigma]^n$. Our construction also returns $\text{SA}^{-1}[1]$, though, so that the corresponding value can be set as needed.

2.2 Word RAM Model

Throughout the paper, we use the standard word RAM model of computation [18] with w -bit *machine words*, where $w \geq \log n$.

In the word RAM model, strings are typically represented as arrays, with each character occupying a single memory cell. Nevertheless, a single character can be represented using $\lceil \log \sigma \rceil$ bits, which might be much less than w . Consequently, one may store a text $T \in [0 \dots \sigma]^n$ in $O\left(\left\lceil \frac{n \log \sigma}{w} \right\rceil\right)$ consecutive memory cells. In the *packed representation* of T , we assume that the first character corresponds to the $\lceil \log \sigma \rceil$ least significant bits of the first cell.

PROPOSITION 2.3 (♣). *Suppose that $T \in [0 \dots \sigma]^n$ is stored in the packed representation. The packed representation of any length- ℓ substring can be retrieved in $O\left(\left\lceil \frac{\ell \log \sigma}{w} \right\rceil\right)$ time. The longest common prefix of two length- ℓ fragments can be identified in the same time.*

A particularly important case is that of $\sigma = 2$. In many applications, these *bitvectors* are equipped with a data structure answering *rank* queries: for $B[1 \dots n]$, $\text{rank}_B(i) = |\{j \in [1 \dots i] : B[j] = 1\}|$. Jacobson [22] proved that rank_B queries can be answered in $O(1)$ time using an additional component of $o(n)$ extra bits. However, an efficient construction of such a component is much more recent.

PROPOSITION 2.4 ([3, 36]). *A packed bitvector $B[1 \dots n]$ can be extended in $O\left(\frac{n}{\log n}\right)$ time with a data structure of size $o\left(\frac{n}{\log n}\right)$ which answers rank_B queries in $O(1)$ time.*

2.3 Wavelet Trees

Wavelet trees, invented by Grossi, Gupta, and Vitter [15] for space-efficient text indexing, are important data structures with a vast number of applications far beyond text processing (see [37]).

The wavelet tree of a string $W \in [0 \dots 2^b]^n$ is recursively defined as follows. First, we create the root node v_ε . This completes the construction for $b = 0$. If $b > 0$, we attach to v_ε a bitvector $B_\varepsilon[1 \dots n]$ in which $B_\varepsilon[i]$ is the most significant bit of $W[i]$ (interpreted as a b -bit number). Next, we partition W into subsequences W_0 and W_1 by scanning W and appending $W[i]$, with the most significant bit d removed, to the subsequence W_d . Finally, we attach the recursively created wavelet trees of W_0 and W_1 (over alphabet $[0 \dots 2^{b-1}]$) to v_ε . The result is a perfect binary tree with 2^b leaves.

Assuming that we label edges 0 (resp. 1) if they go to the left (resp. right) child, we define the *label* of a node to be the concatenation of the labels on the root-to-node path. If B_X denotes the bitvector of a node v_X labeled $X \in \{0, 1\}^{<b}$, then B_X contains one bit (following X as a prefix) from each $W[i]$ whose binary encoding has prefix X . Importantly, the bits in the bitvector B_X occur in the same order as the corresponding elements $W[i]$ occur in W .

It is easy to see that the space occupied by the bitvectors is $O(nb)$ bits, i.e., $O\left(\frac{nb}{\log n}\right)$ words. We need one extra machine word per node for pointers and due to word alignment, which sums up to $O(2^b)$. Thus, the total size of a wavelet tree is $O\left(2^b + \frac{nb}{\log n}\right)$ machine words, which is $O\left(\frac{nb}{\log n}\right)$ if $b \leq \log n$. As shown recently, a wavelet tree can be constructed efficiently from the packed representation of W .

THEOREM 2.5 ([3, 36]). *Given the packed representation of a string W of length n over $[0 \dots 2^b]$ for $b \leq \log n$, we can construct its wavelet tree in $O(nb/\sqrt{\log n})$ time using $O(nb/\log n)$ space.*

3 STRING SYNCHRONIZING SETS

In this section, we introduce string synchronizing sets, the central novel concept underlying both main results of this paper.

DEFINITION 3.1. *Let T be a string of length n and let $\tau \leq \frac{1}{2}n$ be a positive integer. We say that a set $S \subseteq [1 \dots n - 2\tau + 1]$ is a τ -synchronizing set of T if it satisfies the following conditions:*

- (1) *if $T[i \dots i + 2\tau] = T[j \dots j + 2\tau]$, then $i \in S$ holds if and only if $j \in S$ (for $i, j \in [1 \dots n - 2\tau + 1]$), and*
- (2) *$S \cap [i \dots i + \tau] = \emptyset$ if and only if $i \in R$ (for $i \in [1 \dots n - 3\tau + 2]$), where*

$$R = \{i \in [1 \dots n - 3\tau + 2] : \text{per}(T[i \dots i + 3\tau - 2]) \leq \frac{1}{3}\tau\}.$$

Intuitively, the above definition requires that the decision on whether $i \in S$ depends entirely on $T[i \dots i + 2\tau]$, i.e., it is made consistently across the whole text (the first *consistency condition*) and that S contains densely distributed positions within (and only within) non-periodic regions of T (the second *density condition*).

The properties of a τ -synchronizing set S allow for symmetry-breaking decisions that let us individually process only positions $i \in S$, compared to the classic $O(n)$ -time algorithms handling all positions one by one. Thus, we are interested in minimizing the size of S . Since $R = \emptyset$ is possible in general, the smallest τ -synchronizing set we can hope for is of size $\Omega\left(\frac{n}{\tau}\right)$ in the worst case. Our deterministic construction in Section 8 matches this lower bound.

Note that the notion of a τ -synchronizing set is valid for every positive integer $\tau \leq \frac{1}{2}n$. Some applications make use of many

synchronizing sets with parameters τ spread across the whole domain; see [1, 27]. However, in this paper we only rely on τ -synchronizing sets for $\tau = \varepsilon \log_\sigma n$ (where ε is a sufficiently small positive constant), because this value turns out to be the suitable for processing the packed representation of a text $T \in [0 \dots \sigma]^n$ stored in $\Theta(n/\log_\sigma n)$ machine words. This is because our generic construction algorithm (Proposition 8.10) runs in $O(n)$ time, whereas a version optimized for packed strings (Theorem 8.11) takes $O(\frac{n}{\tau})$ time only for $\tau \leq \varepsilon \log_\sigma n$ with $\varepsilon < \frac{1}{5}$. (Note that an $O(\frac{n}{\tau})$ -time construction is feasible for $\tau = O(\log_\sigma n)$ only, because we need to spend $\Omega(n/\log_\sigma n)$ time already to read the whole packed text.) Moreover, the running time of our BWT construction procedure involves a term $\sigma^{O(\tau)}$, which would dominate if we set τ too large.

We conclude this section with two properties of τ -synchronizing sets useful across all our applications. To formulate them, we define the *successor* in S for each $i \in [1 \dots n - 2\tau + 1]$:

$$\text{succ}_S(i) := \min\{j \in S \cup \{n - 2\tau + 2\} : j \geq i\}.$$

The sentinel $n - 2\tau + 2$ guarantees that the set on the right-hand side is non-empty. Our first result applies the density condition to relate $\text{succ}_S(i)$ for $i \in R$ with maximal periodic regions of T .

FACT 3.2. *Let T be a text and let S be its τ -synchronizing set for a positive integer $\tau \leq \frac{1}{2}$. If $i \in R$ and $p = \text{per}(T[i \dots i + 3\tau - 2])$, then $T[i \dots \text{succ}_S(i) + 2\tau - 2]$ is the longest prefix of $T[i \dots n]$ with period p .*

PROOF. Let us define $s = \text{succ}_S(i)$ and observe that $[i \dots s] \cap S = \emptyset$. Consequently, $[j \dots j + \tau] \cap S = \emptyset$ holds for every $j \in [i \dots s - \tau]$. By the density condition, this implies $[i \dots s - \tau] \subseteq R$, i.e., that $\text{per}(T[j \dots j + 3\tau - 2]) \leq \frac{1}{3}\tau$ for $j \in [i \dots s - \tau]$. We shall prove by induction on j that $p = \text{per}(T[i \dots j + 3\tau - 2])$ holds for $j \in [i \dots s - \tau]$. The base case of $j = i$ follows from the definition of p . For $j > i$, on the other hand, let us denote $p' = \text{per}(T[j \dots j + 3\tau - 2])$ and assume $p = \text{per}(T[i \dots j + 3\tau - 3])$ by the inductive hypothesis. We observe that $j + 3\tau - 2 - p' - p > j + 2\tau - 2 \geq j$, so $T[j + 3\tau - 2] = T[j + 3\tau - 2 - p'] = T[j + 3\tau - 2 - p' - p] = T[j + 3\tau - 2 - p]$. This shows that $p = \text{per}(T[i \dots j + 3\tau - 2])$ and completes the inductive step. We conclude that p is the shortest period of $T[i \dots s + 2\tau - 2]$. We now need to prove that this is the longest prefix of $T[i \dots n]$ with period p . The claim is trivially true if $s = n - 2\tau + 2$. Otherwise, $s \in S$, so $[s - \tau + 1 \dots s] \cap S \neq \emptyset$. By the density condition, this means that $s - \tau + 1 \notin R$, i.e., $\text{per}(T[s - \tau + 1 \dots s + 2\tau - 1]) > \frac{1}{3}\tau$. As a result, $\text{per}(T[i \dots s + 2\tau - 1]) > \frac{1}{3}\tau \geq p$. This completes the proof. \square

The second result applies the consistency condition to relate $\text{succ}_S(i)$ with $\text{succ}_S(j)$ for two common starting positions $i, j \in [1 \dots n - 2\tau + 1]$ of a sufficiently long substring.

FACT 3.3. *Let T be a text and let S be its τ -synchronizing set for a positive integer $\tau \leq \frac{1}{2}$. If a substring X of length $|X| \geq 2\tau$ occurs in T at positions i and j , then either*

- (i) $\text{succ}_S(i) - i = \text{succ}_S(j) - j \leq |X| - 2\tau$, or
- (ii) $\text{succ}_S(i) - i > |X| - 2\tau$ and $\text{succ}_S(j) - j > |X| - 2\tau$.

Moreover, (i) holds if $|X| \geq 3\tau - 1$ and $\text{per}(X) > \frac{1}{3}\tau$.

PROOF. First, we shall prove that (i) holds if (ii) does not. Without loss of generality, we assume that $\text{succ}_S(i) - i \leq \text{succ}_S(j) - j$, which yields $\text{succ}_S(i) - i \leq |X| - 2\tau$. In particular, $T[i \dots \text{succ}_S(i) + 2\tau] = T[j \dots j - i + \text{succ}_S(i) + 2\tau]$ is a prefix of X . Moreover, $\text{succ}_S(i) \leq$

$n - 2\tau + 1$, so $\text{succ}_S(i) \in S$. The consistency condition therefore implies $j - i + \text{succ}_S(i) \in S$. Hence, $\text{succ}_S(j) \leq j - i + \text{succ}_S(i)$, and $\text{succ}_S(i) - i = \text{succ}_S(j) - j$ thus holds as claimed.

Next, we shall prove that $\text{succ}_S(i) - i \leq |X| - 2\tau$ if $|X| \geq 3\tau - 1$ and $\text{per}(X) > \frac{1}{3}\tau$. From this, we shall conclude that (ii) does not hold (whereas (i) holds) in that case. If $i \notin R$, then $[i \dots i + \tau] \cap S \neq \emptyset$ by the density condition, so $\text{succ}_S(i) - i \leq \tau - 1 \leq |X| - 2\tau$. Otherwise, let us define $p = \text{per}(T[i \dots i + 3\tau - 2])$ and note that $T[i \dots \text{succ}_S(i) + 2\tau - 2]$ has period p by Fact 3.2. Since $p \leq \frac{1}{3}\tau$, this means that $|X| > |T[i \dots \text{succ}_S(i) + 2\tau - 2]|$, which is equivalent to the desired inequality $\text{succ}_S(i) - i \leq |X| - 2\tau$. \square

4 SORTING SUFFIXES STARTING IN SYNCHRONIZING SETS

Let $T \in [0 \dots \sigma]^n$ be a text stored in the packed representation and let S be its τ -synchronizing set of size $O(\frac{n}{\tau})$ for $\tau = O(\log_\sigma n)$. In this section, we show that given the above as input, the suffixes of T starting at positions in S can be sorted lexicographically in the optimal $O(\frac{n}{\tau})$ time. We assume that the elements of S are stored in an array in the left-to-right order so that we can access the i th smallest element, denoted s_i , in constant time for $i \in [1 \dots |S|]$. The presented algorithm is the first step in our BWT construction. It also reveals the key ideas behind our LCE data structure.

4.1 The Nonperiodic Case

Consider first a case when $R = \emptyset$. The density condition then simplifies to the following statement:

$$(2') \quad S \cap [i \dots i + \tau] \neq \emptyset \text{ for every } i \in [1 \dots n - 3\tau + 2].$$

We introduce a string T' of length $n' := |S|$ defining it so that $T'[i] = T[s_i \dots \min(n, s_i + 3\tau - 1)]$. All characters of T' are strings over $[0 \dots \sigma]$ of length up to 3τ . Hence, they can be encoded using $O(\tau \log \sigma) = O(\log n)$ -bit integers so that the lexicographic order is preserved.² Furthermore, the lexicographic order of the suffixes of T' coincides with that of the corresponding suffixes of T .

LEMMA 4.1. *Assume $R = \emptyset$ holds for a text T . If positions i, j of T' satisfy $T'[i \dots n'] < T'[j \dots n']$, then $T[s_i \dots n] < T[s_j \dots n]$.*

PROOF. We proceed by induction on $\text{lcp}(T'[i \dots n'], T'[j \dots n'])$. The base case is that $T'[i] < T'[j]$. If $T'[i]$ is a proper prefix of $T'[j]$, then $T[s_i \dots n] = T'[i] < T'[j] \leq T[s_j \dots n]$. Otherwise, $T'[i] \cdot W < T'[j]$ holds for any string W , so $T[s_i \dots n] < T'[j] \leq T[s_j \dots n]$.

Henceforth, we assume that $T'[i] = T'[j]$. Since $i \neq j$, this implies $|T'[i]| = |T'[j]| = 3\tau$ and $s_i, s_j \leq n - 3\tau + 1$. The density condition yields $s_{n'} \geq n - 3\tau + 2$ (due to $n - 3\tau + 2 \notin R$), so we further have $i, j \in [1 \dots n']$ and $T'[i + 1 \dots n'] < T'[j + 1 \dots n']$. Moreover, $X := T[s_i + 1 \dots s_i + 3\tau] = T[s_j + 1 \dots s_j + 3\tau]$ occurs in T at positions $s_i + 1$ and $s_j + 1$. As $\text{per}(X) > \frac{1}{3}\tau$ (due to $s_i + 1 \notin R$), Fact 3.3 implies

$$\text{succ}_S(s_i + 1) - (s_i + 1) = \text{succ}_S(s_j + 1) - (s_j + 1) \leq \tau - 1,$$

i.e., $s_{i+1} - s_i = s_{j+1} - s_j \leq \tau$. Furthermore, $T[s_i \dots s_{i+1}] = T[s_j \dots s_{j+1}]$ because $T[s_i \dots s_i + 3\tau] = T'[i] = T'[j] = T[s_j \dots s_j + 3\tau]$. Due to $T[s_{i+1} \dots n] < T[s_{j+1} \dots n]$ (which we derive from the inductive hypothesis), this implies $T[s_i \dots n] < T[s_j \dots n]$ and completes the proof of the inductive step. \square

²For example, we may append $6\tau - 2|T'[i]|$ zeroes and $|T'[i]|$ ones to $T'[i]$. The result can then be interpreted as the base- σ representation of an integer in $[0 \dots \sigma^{6\tau}]$.

By Lemma 4.1, the suffix array of T' can be used to retrieve the lexicographic order of the suffixes $T[s_i \dots n]$ for $s_i \in S$. Recall that each symbol of T' takes $O(\tau \log \sigma) = O(\log n)$ bits, so the suffix array of T' can be computed in $O(|T'|) = O(\frac{n}{\tau})$ time [24].

4.2 The General Case

We now show how to adapt the approach from the previous section so that it also works if $R \neq \emptyset$. As before, we construct a string T' of length $n' = |S|$ over a polynomially bounded integer alphabet, and we sort its suffixes. However, the definition of T' becomes more involved. To streamline the formulae, we set $s_{n'+1} = n - 2\tau + 2$. For each $i \in [1 \dots n']$, we define $T'[i] = (T[s_i \dots \min(n, s_i + 3\tau - 1)], d_i)$, where d_i is an integer specified as follows:

- (a) If $s_{i+1} - s_i \leq \tau$ (in particular, if $s_i > n - 3\tau + 1$), then $d_i = 0$.
- (b) Otherwise, we set $p_i = \text{per}(T[s_i + 1 \dots s_i + 3\tau])$ and

$$d_i = \begin{cases} n - s_{i+1} + s_i & \text{if } T[s_{i+1} + 2\tau - 1] > T[s_{i+1} + 2\tau - 1 - p_i], \\ s_{i+1} - s_i - n & \text{otherwise (if } s_{i+1} = n - 2\tau + 2 \text{ in particular).} \end{cases}$$

Note that each $T'[i]$ can be encoded in $O(\tau \log \sigma + \log n) = O(\log n)$ bits so that the comparison of the resulting integers is equivalent to the lexicographic comparison of the corresponding symbols.

LEMMA 4.2. *If positions i, j of T' satisfy $T'[i \dots n'] < T'[j \dots n']$, then $T[s_i \dots n] < T[s_j \dots n]$.*

PROOF. Induction on $\text{lcp}(T'[i \dots n'], T'[j \dots n'])$. If $\text{lcp}(T[s_i \dots n], T[s_j \dots n]) < 3\tau$, then we proceed as in the proof of Lemma 4.1.

Otherwise, the string $X = T[s_i + 1 \dots s_i + 3\tau] = T[s_j + 1 \dots s_j + 3\tau]$ occurs in T at positions $s_i + 1$ and $s_j + 1$. If $\min(s_{i+1} - s_i, s_{j+1} - s_j) \leq \tau$, then Fact 3.3 yields $s_{i+1} - s_i = s_{j+1} - s_j \leq \tau$, so $d_i = d_j = 0$ and $T'[i] = T'[j]$. Moreover, $i, j \in [1 \dots n']$ due to $s_i, s_j \leq n - 3\tau + 1$. Consequently, the claim follows from $T[s_i \dots s_{i+1}] = T[s_j \dots s_{j+1}]$ because the inductive hypothesis yields $T[s_{i+1} \dots n] < T[s_{j+1} \dots n]$.

On the other hand, $\min(s_{i+1} - s_i, s_{j+1} - s_j) > \tau$ yields $d_i, d_j \neq 0$. Moreover, the density condition implies $s_i + 1, s_j + 1 \in R$ with $p_i = p_j = \text{per}(X) \leq \frac{1}{3}\tau$. By Fact 3.2, the longest prefix of $T[s_i + 1 \dots n]$ with period p_i is $P_i := T[s_i + 1 \dots s_i + 1 + 2\tau - 2]$ and the longest prefix of $T[s_j + 1 \dots n]$ with period p_j is $P_j := T[s_j + 1 \dots s_j + 1 + 2\tau - 2]$. Both P_i and P_j start with X , so one of them is a prefix of the other. We consider three cases based on how their lengths, $|P_i| = n + 2\tau - 2 - |d_i|$ and $|P_j| = n + 2\tau - 2 - |d_j|$, compare to each other.

- If $|d_i| > |d_j|$, then P_i is a proper prefix of P_j . If $i = n'$, then $T[s_i \dots n] = P_i < P_j \leq T[s_j \dots n]$. Otherwise, we note that $d_i < 0$ due to $d_i < d_j$, so $T[s_i + 1 + |P_i|] < P_i[|P_i| - p_i + 1] = P_j[|P_i| - p_j + 1] = T[s_j + 1 + |P_i|]$, which yields the claim.
- If $|d_i| = |d_j|$, then $P_i = P_j$. If $i = n'$, then $T[s_i \dots n] = P_i = P_j < T[s_j \dots n]$. Otherwise, we consider two subcases:
 - If $d_i = -d_j$, then $d_i < 0 < d_j$, so $T[s_i + 1 + |P_i|] < P_i[|P_i| - p_i + 1] = P_j[|P_j| - p_j + 1] < T[s_j + 1 + |P_j|]$, which also yields the claim.
 - Finally, if $d_i = d_j$, then $T'[i] = T'[j]$ and $i, j \in [1 \dots n']$, so the inductive hypothesis gives $T[s_{i+1} \dots n] < T[s_{j+1} \dots n]$. The claim follows due to $T[s_i \dots s_{i+1}] = T[s_j \dots s_{j+1}]$.
- If $|d_i| < |d_j|$, then P_j is a proper prefix of P_i . Moreover, $d_j > 0$ due to $d_i < d_j$, so $T[s_i + 1 + |P_j|] = P_i[|P_j| - p_i + 1] = P_j[|P_j| - p_j + 1] < T[s_j + 1 + |P_j|]$ and the claim holds. \square

We now prove that efficient construction of T' is indeed possible. The only difficulty is computing the values p_i in case $s_{i+1} - s_i > \tau$. To achieve this in constant time, we observe that $p_i \leq \frac{1}{3}\tau$ holds by the density condition due to $s_i < n - 3\tau + 2$ and $s_i + 1 \notin R$. Consequently, p_i is also the shortest period of every prefix of $T[s_i + 1 \dots s_i + 3\tau]$ of length $2p_i$ or more. By the synchronizing property of primitive strings [11, Lemma 1.11], this means that the leftmost occurrence of $T[s_i + 1 \dots s_i + \tau]$ in $T[s_i + 2 \dots s_i + 2\tau]$ starts at position p_i . We can find it in $O(1)$ time (after $O(n^\epsilon)$ -time preprocessing) using the packed string matching algorithm [5].

THEOREM 4.3. *Given the packed representation of a text $T \in [0 \dots \sigma]^n$ and its τ -synchronizing set S of size $O(\frac{n}{\tau})$ for $\tau = O(\log_\sigma n)$, we can compute in $O(\frac{n}{\tau})$ time the lexicographic order of all suffixes of T starting at positions in S .*

5 DATA STRUCTURE FOR LCE QUERIES

In Section 4, for a text $T \in [0 \dots \sigma]^n$ and its τ -synchronizing set S with $\tau = O(\log_\sigma n)$, we constructed a string T' such the lexicographic order of the suffixes of T' coincides with the order of suffixes of T starting at positions in S . In this section, we show how to reduce LCE queries in T to LCE queries in T' . Our approach results in a data structure with $O(1)$ -time LCE queries and $O(\frac{n}{\tau})$ -time construction provided that $|S| = O(\frac{n}{\tau})$. Recall that $n' = |S| = |T'|$, s_i is the i th smallest element of S , and $s_{n'+1} = n - 2\tau + 2$.

5.1 The Nonperiodic Case

Analogously to Section 4.1, we start with the case of $R = \emptyset$, which makes the definition of T' simpler: $T'[i] = T[s_i \dots \min(n, s_i + 3\tau - 1)]$.

Consider an LCE query in the text T . If $\text{LCE}(i, j) < 3\tau$, we can retrieve it in $O(1)$ time from the packed representation of T . Otherwise, Fact 3.3 yields $\text{succ}_\sigma(i) - i = \text{succ}_\sigma(j) - j < \tau$. Hence, $\text{LCE}(i, j) = s_{i'} - i + \text{LCE}(s_{i'}, s_{j'})$, where $s_{i'} = \text{succ}_\sigma(i)$ and $s_{j'} = \text{succ}_\sigma(j)$. A similar reasoning can be repeated to determine $\text{LCE}(s_{i'}, s_{j'})$, which must be smaller than 3τ or equal to $s_{i'+1} - s_{i'} + \text{LCE}(s_{i'+1}, s_{j'+1})$. The former condition can be verified by checking whether $T'[i'] = T'[j']$. A formal recursive application of this argument results in the following characterization:

FACT 5.1. *Consider a string $T \in [0 \dots \sigma]^n$ which satisfies $R = \emptyset$. For positions i, j in T such that $\text{LCE}(i, j) \geq 3\tau - 1$, let us define $s_{i'} = \text{succ}_\sigma(i)$ as well as $s_{j'} = \text{succ}_\sigma(j)$. If $\ell = \text{LCE}_{T'}(i', j')$, then*

$$\text{LCE}(i, j) = s_{i'+\ell} - i + \text{LCE}(s_{i'+\ell}, s_{j'+\ell}) < s_{i'+\ell} - i + 3\tau.$$

PROOF. The proof is by induction on ℓ . Due to $i, j \notin R$, Fact 3.3 yields $s_{i'} - i = s_{j'} - j < \tau$, and therefore $T[i \dots s_{i'}] = T[j \dots s_{j'}]$. Hence, $\text{LCE}(i, j) = s_{i'} - i + \text{LCE}(s_{i'}, s_{j'})$. If $\ell = 0$, it just remains to prove that $\text{LCE}(s_{i'}, s_{j'}) < 3\tau$, which follows from $T'[i'] \neq T'[j']$.

For $\ell > 0$, we note that $T'[i'] = T'[j']$, so $\text{LCE}(s_{i'}, s_{j'}) \geq 3\tau$ and $\text{LCE}(s_{i'} + 1, s_{j'} + 1) \geq 3\tau - 1$. The inductive hypothesis now yields $\text{LCE}(s_{i'} + 1, s_{j'} + 1) = s_{i'+\ell} - s_{i'} - 1 + \text{LCE}(s_{i'+\ell}, s_{j'+\ell})$ and $\text{LCE}(s_{i'+\ell}, s_{j'+\ell}) < 3\tau$. Since $\text{LCE}(i, j) = s_{i'} - i + \text{LCE}(s_{i'}, s_{j'}) = s_{i'} + 1 - i + \text{LCE}(s_{i'} + 1, s_{j'} + 1)$, this completes the proof. \square

Fact 5.1 leads to a data structure for LCE queries that consists of the packed representation of T (Proposition 2.3), a τ -synchronizing set S of size $O(\frac{n}{\tau})$, a component for LCE queries in T' (Proposition 2.2; the alphabet size is $\sigma^{3\tau} = n^{O(1)}$), and a bitvector $B[1 \dots n]$,

with $B[i] = 1$ if and only if $i \in S$, equipped with a component for $O(1)$ -time rank queries (Proposition 2.4).

To compute $\text{LCE}(i, j)$, we first use the packed representation to retrieve the answer in $O(1)$ time provided that $\text{LCE}(i, j) < 3\tau$. Otherwise, we obtain i' and j' such that $s_{i'} = \text{succ}_S(i)$ and $s_{j'} = \text{succ}_S(j)$ using rank_B queries, and we compute $\ell = \text{LCE}_{T'}(i', j')$. By Fact 5.1, $\text{LCE}(i, j) = s_{i'+\ell} - i + \text{LCE}(s_{i'+\ell}, s_{j'+\ell})$. Since $\text{LCE}(s_{i'+\ell}, s_{j'+\ell}) < 3\tau$, we finalize the algorithm using the packed representation again.

5.2 The General Case

In this section, we generalize the results of Section 5.1 so that the case of $R \neq \emptyset$ is also handled. Our data structure consists of the same components; the only difference is that the string T' is now defined as in Section 4.2 rather than as in Section 4.1.

The query algorithm needs more changes but shares the original outline. If $\text{LCE}(i, j) < 3\tau$, then we determine the answer using Proposition 2.3. Otherwise, we apply the following lemma as a reduction to computing $\text{LCE}(\text{succ}_S(i), \text{succ}_S(j))$.

LEMMA 5.2 (♣). *For positions i, j in T such that $\text{LCE}(i, j) \geq 3\tau - 1$, let us define $s_{i'} = \text{succ}_S(i)$ and $s_{j'} = \text{succ}_S(j)$. Then*

$$\text{LCE}(i, j) = \begin{cases} \min(s_{i'} - i, s_{j'} - j) + 2\tau - 1 & \text{if } s_{i'} - i \neq s_{j'} - j, \\ s_{i'} - i + \text{LCE}(s_{i'}, s_{j'}) & \text{if } s_{i'} - i = s_{j'} - j. \end{cases}$$

We are left with determining the values $\text{LCE}(s_i, s_j)$ for $i, j \in [1..n'] + 1$, i.e., handling LCE queries for positions in $S \cup \{n - 2\tau + 2\}$. The next result reduces this task to answering LCE queries in T' .

LEMMA 5.3 (♣). *If $\ell = \text{LCE}_{T'}(i, j)$ for positions $i, j \in [1..n'] + 1$, then $\text{LCE}(s_i, s_j) = s_{i+\ell} - s_i + \text{LCE}(s_{i+\ell}, s_{j+\ell})$. Moreover, $\text{LCE}(s_i, s_j) < 3\tau$ or $\text{LCE}(s_i, s_j) = \min(s_{i+1} - s_i, s_{j+1} - s_j) + 2\tau - 1$ holds if $\ell = 0$.*

We are now ready to describe the complete query algorithm determining $\text{LCE}(i, j)$ for two positions i, j in T . We start by using Proposition 2.3 to compare the first 3τ symbols of $T[i..n]$ and $T[j..n]$. If we detect a mismatch, the procedure is completed. Otherwise, we compute the indices $i', j' \in [1..n'] + 1$ of $s_{i'} = \text{succ}_S(i)$ and $s_{j'} = \text{succ}_S(j)$ using rank queries on the bitvector B . If $s_{i'} - i \neq s_{j'} - j$, then we answer the query $\text{LCE}(i, j) = \min(s_{i'} - i, s_{j'} - j) + 2\tau - 1$ according to Lemma 5.2. Otherwise, Lemma 5.2 yields $\text{LCE}(i, j) = s_{i'} - i + \text{LCE}(s_{i'}, s_{j'})$, and it remains to compute $\text{LCE}(s_{i'}, s_{j'})$. For this, we query for $\ell = \text{LCE}_{T'}(i', j')$ and note that $\text{LCE}(s_{i'}, s_{j'}) = s_{i'+\ell} - s_{i'} + \text{LCE}(s_{i'+\ell}, s_{j'+\ell})$ by Lemma 5.3. Finally, we are left with determining the latter LCE value. We start by comparing the first 3τ symbols of $T[s_{i'+\ell}..n]$ and $T[s_{j'+\ell}..n]$. If we detect a mismatch, the procedure is finished. Otherwise, we compute $\text{LCE}(s_{i'+\ell}, s_{j'+\ell}) = \min(s_{i'+\ell+1} - s_{i'+\ell}, s_{j'+\ell+1} - s_{j'+\ell}) + 2\tau - 1$ according to Lemma 5.3. This completes the algorithm.

Before we conclude, note that given a synchronizing set of size $O(\frac{n}{\tau})$ for $\tau = O(\log_\sigma n)$, the data structure can be constructed in $O(\frac{n}{\tau})$ time. This follows from Theorem 4.3 (building T'), Proposition 2.2 (LCE queries in T'), and Proposition 2.4 (rank $_B$ queries). If $\tau \leq \varepsilon \log_\sigma n$ for a positive constant $\varepsilon < \frac{1}{5}$, then Theorem 8.11 also lets us compute an appropriate τ -synchronizing set in $O(\frac{n}{\tau})$ time. The overall construction time, $O(\frac{n}{\tau})$, is minimized by $\tau = \Theta(\log_\sigma n)$.

THEOREM 5.4. *LCE queries in a text $T \in [0.. \sigma]^n$ with $\sigma = n^{O(1)}$ can be answered in $O(1)$ time after $O(n/\log_\sigma n)$ -time preprocessing of the packed representation of T .*

6 BWT CONSTRUCTION

Let $T \in [0.. \sigma]^n$, for $\log \sigma \leq \sqrt{\log n}$, be a text given in the packed representation, and let S be a τ -synchronizing set of T of size $O(\frac{n}{\tau})$, where $\tau = \varepsilon \log_\sigma n$ for some sufficiently small constant $\varepsilon > 0$. We assume that τ is a positive integer and that $3\tau - 1 \leq n$.

In this section, we show how to construct the BWT of T in $O(n \log \sigma / \sqrt{\log n})$ time and $O(n/\log_\sigma n)$ space. For simplicity, we first restrict ourselves to a binary alphabet. The time and space complexities then simplify to $O(n/\sqrt{\log n})$ and $O(n/\log n)$.

6.1 Binary Alphabet

Similarly as in previous sections, we first assume $R = \emptyset$ (note that this implies $S \neq \emptyset$ due to $3\tau - 1 \leq n$). In Section 6.1.2, we consider the general case and describe the remaining parts of our construction.

6.1.1 *The Nonperiodic Case.* To compensate for the lack of a sentinel $T[n] = \$$ (see Section 2), let us choose $b_\$ \in \{0, 1\}$ such that $\text{per}(X) > \frac{1}{3}\tau$ holds for $X = b_\$ T[1..2\tau]$. Using packed string matching [5], we add to S all positions where X occurs in T . This increases $|S|$ by $O(\frac{n}{\tau})$ and does not violate Definition 3.1. Denote by $(s'_i)_{i \in [1..|S|]}$ the set S , sorted using Theorem 4.3 according to the order of the corresponding suffixes, i.e., $T[s'_i..n] < T[s'_j..n]$ if $i < j$. Define a sequence W of length $|S|$ so that $W[i] \in [0..2^{3\tau}]$ is an integer whose base-2 representation is $\overline{T[s'_i - \tau..s'_i + 2\tau]}$, where \overline{X} denotes the string-reversal operation.³ In the word RAM model with word size $w = \Omega(\log n)$, reversing any $O(\log n)$ -bit string takes $O(1)$ time after $O(n^\delta)$ -time ($\delta < 1$) preprocessing. Thus, $W[1..|S|]$ can be constructed in $O(|S| + n^\delta) = O(n/\log n)$ time.

Recall that the density condition simplifies to $S \cap [i..i + \tau] \neq \emptyset$ for $i \in [1..n - 3\tau + 2]$ if $R = \emptyset$. Thus, except for $O(\tau)$ rightmost symbols, every symbol of T is included in at least one character of W . In principle, it suffices to rearrange these bits to obtain the BWT. For this, we utilize as a black box the wavelet tree of W and prove that its construction performs the necessary permuting. We are then left with a task of copying the bits from the wavelet tree.

More precisely, we show how to extract (almost) all bits of the BWT of T from the bitvectors B_X in the wavelet tree of W in $2^{\Theta(\tau)} + O(n/\log n)$ time. Intuitively, we partition the BWT into $2^{\Theta(\tau)}$ blocks that appear as bitvectors B_X .

A similar string W was constructed in [10] for an evenly distributed set of positions. In that case, however, the bitvectors in the wavelet tree form non-contiguous subsequences of the BWT.

Distinguishing Prefixes. To devise the announced partitioning of the BWT, for $j \in [1.. \max S]$ let $D_j = T[j.. \text{succ}_S(j) + 2\tau]$ be the *distinguishing prefix* of $T[j..n]$. By the density condition for $R = \emptyset$, $\text{succ}_S(j) - j < \tau$ and thus $|D_j| \leq 3\tau - 1$. Let $\mathcal{D} = \{D_j : j \leq \max S\}$.

Recall that B_X is the bitvector associated with the node v_X whose root-to-node label in the wavelet tree of W is X . By definition of the wavelet tree (applied to W), for any $X \in \{0, 1\}^{\leq 3\tau - 1}$, $B_{\overline{X}}$ contains the bit preceding X from each string $T[s'_i - \tau..s'_i + 2\tau]$ that has X as a suffix. (The order of these bits matches the sequence $(s'_i)_{i \in [1..|S|]}$.)

LEMMA 6.1. (1) *If $T[j..j + |X|] = X$ for $X \in \mathcal{D}$, then $D_j = X$.*
 (2) *If $\text{SA}[b..e]$ includes all suffixes of T having $X \in \mathcal{D}$ as a prefix, then $\text{BWT}[b..e] = B_{\overline{X}}$.*

³Whenever $T[k]$ is out of bounds, we let $T[k] = b_\$$ if $k = 0$ and $T[k] = 0$ otherwise.

PROOF. (1) Let $X = D_i$, i.e., $X = T[i.. \text{succs}(i) + 2\tau]$. Since X also occurs at position j and $\text{succs}(i) - i \leq |X| - 2\tau$, we have $\text{succs}(j) - j = \text{succs}(i) - i$ by Fact 3.3. Consequently, $D_j = T[j.. \text{succs}(j) + 2\tau] = X$.

(2) By the above discussion, $B_{\overline{X}}$ contains the bits preceding X as suffixes in $(T[s'_i - \tau.. s'_i + 2\tau])_{i \in [1..|S|]}$. From (1), there is a bijection between the occurrences of X in T and such suffixes (importantly, $b_{\S} T[1.. s'_i + 2\tau]$ for $s'_i \in S$ is not a suffix of X due to the modification of S , so X is never compared against out-of-bounds symbols of T in W). By definition of $\text{BWT}[b..e]$ and $|X| \leq 3\tau - 1$, $B_{\overline{X}}$ and $\text{BWT}[b..e]$ indeed contain the same (multisets of) bits.

To show that the bits of $\text{BWT}[b..e]$ occur in $B_{\overline{X}}$ in the same order, observe that $T[s'_i + 2\tau - |X|..n] < T[s'_j + 2\tau - |X|..n]$ holds if $T[s'_i - \tau.. s'_i + 2\tau]$ and $T[s'_j - \tau.. s'_j + 2\tau]$ have X as a suffix for $i < j$. This is because $T[s'_i + 2\tau - |X|.. s'_i] = T[s'_j + 2\tau - |X|.. s'_j]$ is a prefix of X , and we have $T[s'_i..n] < T[s'_j..n]$ by $i < j$. \square

The Algorithm. We start by building the string W and its wavelet tree. By Theorem 2.5, this takes $O((n/\tau) \log(2^{3\tau})/\sqrt{\log(n/\tau)}) = O(n/\sqrt{\log n})$ time and $O(n/\log n)$ space.

Next, we create a lookup table that, for any $X \in \{0, 1\}^{2\tau}$, tells whether X occurs at a position $j \in S$ (by the consistency condition, $j \in S$ for every position j where X occurs in T). It needs $O(n^{2\epsilon})$ space and is easily filled: set “yes” for each $T[j..j + 2\tau]$ with $j \in S$.

Initialize the output to an empty string. Consider the preorder traversal of a complete binary tree of depth $3\tau - 1$ with each edge to a left child labeled “0” and to a right child—“1”. Whenever we visit a node with root-to-node path X such that $|X| \geq 2\tau$, we check if the length- 2τ suffix of X is a “yes” substring. If so, we report X and skip the traversal of the current subtree. Otherwise, we descend into the subtree. This procedure enumerates \mathcal{D} in the lexicographic order in $O(2^{3\tau}) = O(n^{3\epsilon})$ time. For each reported substring X , we append $B_{\overline{X}}$ to the output string. Locating $v_{\overline{X}}$ takes $O(|X|) = O(\log n)$ time; hence, we spend $O(n/\log n + n^{3\epsilon} \log n) = O(n/\log n)$ time in total.

The above traversal outputs a BWT subsequence containing the symbols preceding positions in $[1.. \max S]$. To include the missing symbols, we make the following adjustment: while visiting a node with label $X \notin \mathcal{D}$, we check if X occurs as a suffix of T . If so, then before descending into the subtree, we append the preceding character $T[n - |X|]$ to the output string.

The algorithm runs in $O(n/\sqrt{\log n})$ time and uses $O(n/\log n)$ space. For correctness, observe that the set \mathcal{D} is prefix-free by Lemma 6.1. Thus, no symbol is output twice.

To complete the construction, we need $\text{SA}^{-1}[1]$ (see Section 2.1). Let D_1 be the distinguishing prefix of $T[1..n]$ and let i_1 be the index of $\min S$ in $(s'_i)_{i \in [1..|S|]}$. Observe that the symbol $T[0] = b_{\S}$ occurs in $B_{\overline{D_1}}$ at position $|i \leq i_1 : \overline{D_1} \text{ is a prefix of } W[i]|$, which can be determined in $O(|S|)$ time. Appending $B_{\overline{D_1}}$ to the constructed BWT, we map this position in $B_{\overline{D_1}}$ to the corresponding one in the BWT. Finally, we overwrite b_{\S} by setting $\text{BWT}[\text{SA}^{-1}[1]] = T[n]$.

6.1.2 *The General Case.* Let us define

$$\mathcal{F} = \{X \in \{0, 1\}^{3\tau-1} : X' \notin \mathcal{D} \text{ for every prefix } X' \text{ of } X\}.$$

Observe that if $T[j..j + 3\tau - 1] \in \mathcal{F}$, then $j \in R$. Conversely, whenever $j \in R$, then $T[j..j + 3\tau - 1] \in \mathcal{F}$. Thus, R contains precisely the starting positions of all strings in \mathcal{F} . Hence, in the general case with R possibly non-empty, the algorithm of Section 6.1.1 outputs the BWT subsequence missing exactly the symbols $T[j - 1]$ for $j \in R$.

The crucial property of R that allows handling the general case is that R cannot have many “gaps”. Moreover, whenever $X \in \mathcal{F}$ occurs at a position $j \in R$ with $j - 1 \in R$, then $T[j - 1]$ depends on X .

LEMMA 6.2. *Let $R' = \{j \in R : j - 1 \notin R\}$ be a subset of R . Then:*

- (1) $|R'| \leq |S| + 1$.
- (2) If $X = T[j..j + |X|] \in \mathcal{F}$ and $j \notin R'$, then $T[j - 1] = X[\text{per}(X)]$.

PROOF. (1) By density condition, $j \in R'$ implies $j - 1 \in S$ if $j > 1$.

(2) Note that $\text{per}(X) = \text{per}(T[j - 1.. \text{succs}(j - 1) + 2\tau - 2]) \leq \frac{1}{3}\tau$ due to Fact 3.2 and because $\text{succs}(j - 1) = \text{succs}(j) \geq j + \tau$. Hence, $T[j - 1] = T[j - 1 + \text{per}(X)] = X[\text{per}(X)]$. \square

Consider thus the following modification: whenever we reach $X \in \mathcal{F}$ during the enumeration of \mathcal{D} , we append to the output a unary string of f_X symbols $X[\text{per}(X)]$, where f_X is the number of occurrences of X in T . By Lemma 6.2, the number of mistakes in the resulting BWT, over all $X \in \mathcal{F}$, is only $|R'| = O(\frac{n}{\tau})$.

To implement the above modification (excluding BWT correction), we need to compute $\text{per}(X)$ and f_X for every $X \in \{0, 1\}^{3\tau-1}$. The period is determined using a lookup table.

Computing Frequencies of Length- ℓ Substrings. Consider $\lfloor |T|/\ell \rfloor$ blocks of length $2\ell - 1$ starting in T at positions of the form $1 + k\ell$ (the last block might be shorter). Sort all blocks in $O(\frac{n}{\ell})$ time into a list L . Then, scan the list and for each *distinct* block B in L , consider the multiset of all its length- ℓ substrings X . For each such X , increase its frequency by the frequency of B in L . The correctness follows by noting that $T[i..i + \ell]$ is contained in the $\lceil i/\ell \rceil$ th block only.

There are at most $1 + 2^{2\ell-1}$ distinct blocks and we spend $O(\ell)$ time for each. The total running time is therefore $O(\frac{n}{\ell} + \ell 2^{2\ell-1})$.

In our application, $\ell = 3\tau - 1 < 3\epsilon \log n$, so it suffices to choose $\epsilon < \frac{1}{6}$ so that $O(\frac{n}{\ell} + \ell 2^{2\ell-1}) = O(n/\log n + n^{6\epsilon} \log n) = O(n/\log n)$.

Correcting BWT. We will now show how to compute the rank (i.e., the position in the suffix array) of every suffix of T starting in R' . This will let us correct the mistakes in the BWT produced within the previous step. If r_j is the rank of $T[j..n]$, where $j \in R'$, we set $\text{BWT}[r_j] = T[j - 1]$. To compute r_j , we only need to know r'_j : the *local rank* of $T[j..n]$ among the suffixes of T starting with $T[j..j + 3\tau - 1]$ (note that any such suffix starts at $j' \in R$) since the rank among other suffixes is known during the enumeration of \mathcal{D} . Formally, for $j \in R'$, define $\text{pos}(j) = \{j' \in R : \text{LCE}_T(j, j') \geq 3\tau - 1 \text{ and } T[j'..n] \leq T[j..n]\}$ so that $r'_j = |\text{pos}(j)|$.

Motivated by Lemma 6.2, we focus on the properties of *runs* of consecutive positions in R . We start by partitioning such runs into classes, where the computation of local ranks is easier and can be done independently. For $X \in \mathcal{F}$, we define the *Lyndon root* $\text{L-root}(X) = \min\{X[t..t + p] : t \in [1..p]\} \in \{0, 1\}^{\leq \tau/3}$, where $p = \text{per}(X)$. We further set $\text{L-root}(j) = \text{L-root}(T[j..j + 3\tau - 1])$ for every $j \in R$. It is easy to see that if $j \in R \setminus R'$, then $\text{L-root}(j - 1) = \text{L-root}(j)$. Thus, to compute r'_j for some $j \in R'$, it suffices to look at the runs starting at $j' \in R'$ such that $\text{L-root}(j) = \text{L-root}(j')$.

Further, for $j \in R$, let us define $e_j = \min\{j' \geq j : j' \notin R\} + 3\tau - 1$. We define $\text{type}(j) = +1$ if $T[e_j] > T[e_j - p]$ and $\text{type}(j) = -1$ otherwise, where $p = \text{per}(T[j..e_j])$. Similarly as for the L-root , if $j \in R \setminus R'$, then $\text{type}(j - 1) = \text{type}(j)$. Furthermore, if $\text{type}(j) = -1$ holds for $j \in R'$, then $\text{type}(j') = -1$ holds for all $j' \in \text{pos}(j)$. Let $R^- = \{j \in R : \text{type}(j) = -1\}$, $R^+ = R \setminus R^-$, $R'^- = R' \cap R^-$, and

$R'^+ = R' \cap R^+$. In the rest of this section, we focus on computing r_j^+ for $j \in R'^-$. The set R'^+ is processed symmetrically.

To efficiently determine local ranks for a group of runs with the same L-root, we refine the classification further into individual elements of R . Let $U = \text{L-root}(j)$ for $j \in R$. It is easy to see that the following L-decomposition $T[j..e_j] = U'U^kU''$ (where $k \geq 1$, U' is a proper suffix of U , and U'' is a proper prefix of U) is unique. We call the triple $(|U'|, k, |U''|)$ the L-signature of j and the value $\text{L-exp}(j) = k$ its L-exponent. By the uniqueness of L-decompositions, given $j \in R'^-$, we have $\text{L-exp}(j') \leq \text{L-exp}(j)$ for all $j' \in \text{pos}(j)$.

Note that, letting $S = (s_i)_{i \in [1..|S|]}$ where $s_i < s_{i'}$ if $i < i'$ and $s_0 = 0, s_{|S|+1} = n - 2\tau + 2$, we have, by the density condition: $R' = \{s_{i+1} : i \in [0..|S|] \text{ and } s_{i+1} - s_i > \tau\}$. Furthermore, whenever $j - 1 = s_i$ for $j \in R'$, then $e_j = s_{i+1} + 2\tau - 1$. Thus, computing R' and $\{e_j\}_{j \in R'}$ (and also the type of each $j \in R'$) takes $O(|S|)$ time.

For $j \in R'^-$, let $r_j^+ = |\{j' \in \text{pos}(j) : \text{L-exp}(j') = \text{L-exp}(j)\}|$ and $r_j^- = |\{j' \in \text{pos}(j) : \text{L-exp}(j') < \text{L-exp}(j)\}|$.

To compute r_j^- for each $j \in R'^-$, consider sorting the list of all $j \in R'^-$ first according to $\text{L-root}(j)$, second according to $|U''|$ in its L-signature, and third according to $T[e_j..n]$. Such ordering can be obtained in $O(|R'^-|)$ time by utilizing the sequence $(s'_i)_{i \in [1..|S|]}$ and the fact that $e_j - 2\tau + 1 \in S$. The L-root and $|U''|$ are computed using lookup tables. Then, to determine r_j^- , we count $j' \in R'^-$ with the same L-root that are not later than j in the list and the factor $U'U^k$ in their L-decomposition is at least as long as for j . To this end, we issue a 3-sided orthogonal 2D range counting query on an input instance containing, for every $j \in R'^-$, a point with $|U'U^k|$ from its L-decomposition as the first coordinate and its position in the above list as the second coordinate. Answering a batch of m orthogonal 2D range counting queries takes $O(m\sqrt{\log m})$ time and $O(m)$ space [9]. Since $m = |R'^-|$, this step takes $O(n/\sqrt{\log n})$ time.

To compute r_j^+ for each $j \in R'^-$, we sort $j \in R'^-$ first by L-root(j), and then by L-exp(j). For a fixed $U \in \{0, 1\}^{\leq \tau/3}$, let us define

$$R_U^- = \{j \in R^- : \text{L-root}(j) = U\} \subseteq R^-.$$

Let also $R_U^- = R' \cap R_U^-$. The key observation is that there are only $|U|$ different prefixes of length $3\tau - 1$ in $\{T[j..n] : j \in R_U^-\}$. We will incrementally compute the frequency of each of these prefixes $X \in \{0, 1\}^{3\tau-1}$ and keep the count in an array $C[0..|U|]$, indexed by $t = |U'|$ in the L-decomposition of every $j \in R_U^-$ with $T[j..j+3\tau-1] = X$ (denote this set as $R_{U,t}^-$). A single round of the algorithm handles $H_k = \{j \in R_U^- : \text{L-exp}(j) = k\}$. We execute the rounds for increasing k and maintain the invariant that $C[t]$ contains $|\{j \in R_{U,t}^- : \text{L-exp}(j) \leq k\}|$ at the end of round k . At the beginning of round k , we use the values of C to first compute r_j^+ for each $j \in H_k$ and then update C to maintain the invariant. The update consists of increasing some entries in C for each $j \in H_k$, and then increasing all entries in C by the total number q of yet unprocessed positions (having higher L-exponent), i.e., $q = |R_U^-| - \sum_{i=1}^k |H_i|$. It is easy to see that for each $j \in H_k$, the update can be expressed as a constant number of increments in contiguous ranges of C . Additionally, if $H_{k-1} = \emptyset$ and $k_{\text{prev}} = \max\{k' < k : H_{k'} \neq \emptyset\}$, right at the beginning of round k (before computing r_j^+ for $j \in H_k$), we increment all of C by $q \cdot (k - k_{\text{prev}} - 1)$, to account for the skipped L-exponents. Each update of C takes $O(\log |U|) = O(\log \log n)$ time if we implement C as a balanced BST. Thus, the algorithm takes $O(|R'^-| \log \log n) = O(n/\sqrt{\log n})$

time. Note, that there are $O(2^{\tau/3}) = O(n^{\epsilon/3})$ different L-roots; hence, we can afford to initialize C in $O(\log n)$ time for each U .

THEOREM 6.3. *Given the packed representation of a text $T \in \{0, 1\}^n$ and its τ -synchronizing set S of size $O(n/\tau)$ for $\tau = \epsilon \log n$, where $\epsilon > 0$ is a sufficiently small constant, the Burrows–Wheeler transform of T can be constructed in $O(n/\sqrt{\log n})$ time and $O(n/\log n)$ space.*

6.2 Large Alphabets

Note that our BWT construction does not immediately generalize to larger alphabets since for binary strings it already relies on wavelet tree construction for sequences over alphabets of polynomial size. More precisely, in the binary case, we combined the bitvectors B_X in the wavelet tree of a large-alphabet sequence W to retrieve fragments of the Burrows–Wheeler transform of the text T . For an alphabet $\Sigma = [0.. \sigma]$, the BWT consists of $(\log \sigma)$ -bit characters. Thus, instead of standard binary wavelet trees, we use wavelet trees of degree σ . For simplicity, we assume that σ is a power of two.

6.2.1 High-Degree Wavelet Trees. To construct a degree- σ wavelet tree, we consider a string W of length n over an alphabet $[0.. \sigma^b]$ and think of every symbol $W[i]$ as of a number in base σ with exactly b digits (including leading zeros). First, we create the root node v_ϵ and construct its string D_ϵ of length n setting as $D_\epsilon[i]$ the most significant digit of $W[i]$. We then partition W into σ subsequences $W_0, W_1, \dots, W_{\sigma-1}$ by scanning through W and appending $W[i]$ with the most significant digit removed to W_c , where c is the removed digit of $W[i]$. We recursively repeat the construction for every W_c and attach the resulting tree as the c th child of the root. The nodes of the resulting degree- σ wavelet tree are labeled with strings $Y \in [0.. \sigma]^{\leq b}$. For $|Y| < b$, the string D_Y at node v_Y labeled with Y contains the next digit following the prefix Y in the σ -ary representation of $W[i]$ for each element $W[i]$ of W whose σ -ary representation contains Y as a prefix. (The digits in D_Y occur in the order as the corresponding entries $W[i]$.) The total size of a wavelet tree is $O(\sigma^b + \frac{nb \log \sigma}{\log n})$ words, which is $O(\frac{nb \log \sigma}{\log n})$ if $b \leq \log_\sigma n$.

As observed in [3], the binary wavelet tree construction procedure can be used as a black box to build degree- σ wavelet trees. Here, we present a more general version of this reduction.

LEMMA 6.4 (\spadesuit ; SEE [3, LEMMA 2.2]). *Given the packed representation of a string $W \in [0.. \sigma^b]^n$ with $b \leq \log_\sigma n$, we can construct its degree- σ wavelet tree in $O(nb \log \sigma / \sqrt{\log n} + nb \log^2 \sigma / \log n)$ time using $O(nb / \log_\sigma n)$ space.*

6.2.2 BWT Construction Algorithm. Our construction algorithm for $T \in [0.. \sigma]^n$ uses a τ -synchronizing set $(s'_i)_{i \in [1..|S|]}$ (where $T[s'_i..n] < T[s'_j..n]$ if $i < j$) for $\tau = \frac{\epsilon \log_\sigma n}{\sigma}$. We then build a sequence $W \in [0.. \sigma^{3\tau}]^{|S|}$ with $W[i] = T[s'_i - \tau.. s'_i + 2\tau]$, i.e., we reverse $T[s'_i - \tau.. s'_i + 2\tau]$ and then interpret it as a 3τ -digit integer in base σ . Next, we construct a degree- σ wavelet tree of W and combine the strings D_Y , where $Y \in [0.. \sigma]^{\leq 3\tau-1}$, to obtain the BWT of T . The procedure is analogous to the binary case.

THEOREM 6.5. *Given the packed representation of a text $T \in [0.. \sigma]^n$ with $\log \sigma \leq \sqrt{\log n}$ and its τ -synchronizing set S of size $O(\frac{n}{\tau})$ with $\tau = \epsilon \log_\sigma n$, where $\epsilon > 0$ is a sufficiently small constant, the Burrows–Wheeler transform of T can be constructed in $O(n \log \sigma / \sqrt{\log n})$ time and $O(n/\log_\sigma n)$ space.*

7 CONDITIONAL OPTIMALITY OF THE BINARY BWT CONSTRUCTION

Given an array $A[1..m]$ of integers, the task of counting inversions is to compute the number of pairs (i, j) such that $i < j$ and $A[i] > A[j]$. The currently fastest solution for the above problem, due to Chan and Pătraşcu [9], runs in $O(m\sqrt{\log m})$ time and $O(m)$ space.

Without loss of generality [19], we assume $A[i] \in [0..m)$. In this section, we show that improving the BWT construction from Section 6.1 also yields an improvement over [9]. More precisely, we show that computing the BWT of a packed text $T \in \{0, 1\}^n$ in time $O(f(n))$ implies an $O(m + f(m \log m))$ -time construction of the wavelet tree for A ; hence, improving over Theorem 6.3 implies an $o(m\sqrt{\log m})$ -time wavelet tree construction. The main result then follows from the next observation since it is easy to count inversions in a length- n bitvector in $O(n/\log n)$ time.

OBSERVATION 7.1. *The number of inversions in any integer sequence A is equal to the total number of inversions in the bitvectors of the wavelet tree of A .*

We first consider the case when $A[i] < m^\varepsilon$ for all $i \in [1..m)$ and for some sufficiently small constant $\varepsilon < 1$; it is almost a direct reversal of the reduction from Section 6.1.1.

7.1 The Case $A[i] < m^\varepsilon$

Let $\text{bin}_k(x) \in \{0, 1\}^k$ be the base-2 representation of $x \in [0..2^k)$ and let $\text{pad}_k : \{0, 1\}^k \rightarrow \{0, 1\}^{2k}$ be a padding function that, given $X \in \{0, 1\}^k$, inserts a 0-bit before each bit of X .

Assume $\varepsilon \log m$ and $\log m$ are integers. Given $A[1..m]$, let

$$T_A = \prod_{i=1}^m \overline{\text{bin}_{\varepsilon \log m}(A[i])} \cdot 01 \cdot 1^{\varepsilon \log m} \cdot \text{pad}_{\log m}(\text{bin}_{\log m}(i-1)) \cdot 0.$$

The text T_A is of length $m(3+2(1+\varepsilon)\log m)$ and constructing it takes $O(m)$ time. Recall that B_X denotes the bitvector corresponding to the node of the wavelet tree of A with root-to-node label X .

LEMMA 7.2. *For $X \in \{0, 1\}^{<\varepsilon \log m}$, let $\text{SA}[b..e]$ be the range containing all suffixes of T_A having $\overline{X} \cdot 01 \cdot 1^{\varepsilon \log m}$ as a prefix. Then, $B_X = \text{BWT}[b..e]$.*

PROOF. Let $g : i \mapsto (i-1)(3+2(1+\varepsilon)\log m) + \varepsilon \log m - |X| + 1$ restricted to $A_X := \{i \in [1..m) : X \text{ is a prefix of } \text{bin}_{\varepsilon \log m}(A[i])\}$ be a one-to-one map between A_X and the set $\text{SA}[b..e]$. It remains to observe that $T_A[g(i)..|T_A|] < T_A[g(j)..|T_A|]$ holds if $i < j$ for $i, j \in A_X$, and $\text{bin}_{\varepsilon \log m}(A[i])[|X|+1] = T_A[g(i)-1]$ for $i \in A_X$. \square

To compute the BWT ranges corresponding to the bitvectors B_X for $X \in \{0, 1\}^{<\varepsilon \log m}$, we proceed as in Section 6: perform a pre-order traversal of the complete binary tree of height $2\varepsilon \log m + 1$ corresponding to lexicographic enumeration of $X' \in \{0, 1\}^{\leq 2\varepsilon \log m + 1}$. For $X' = X011^{\varepsilon \log m}$, where $X \in \{0, 1\}^{<\varepsilon \log m}$, copy the bits from BWT into $B_{\overline{X}}$. The number of bits to copy is given by the number $f_{X'}$ of occurrences of X' in T_A . After copying the bits (or when we reach $X' \in \{0, 1\}^{2\varepsilon \log m + 1}$), advance the position in BWT by $f_{X'}$.

To determine the frequencies of all $\{0, 1\}^{\leq 2\varepsilon \log m + 1}$, compute $f_{X'}$ for all $X' \in \{0, 1\}^{2\varepsilon \log m + 1}$ in $O(m)$ time using the algorithm in Section 6.1.2. The remaining frequencies are then easily derived.

The algorithm runs in $O(m + f(m \log m))$ time, where $f(n)$ is the runtime of the BWT construction for a packed text from $\{0, 1\}^n$.

7.2 The Case $A[i] < m$

Given an array $A[1..m]$ with values $A[i] \in [0..m)$, let

$$T_A = \prod_{i=1}^m \overline{\text{bin}_{\log m}(A[i])} \cdot 01 \cdot 1^{\log m} \cdot 0 \cdot \text{bin}_{\log m}(i-1) \cdot 0 \cdot 1^{\log m} \cdot 0.$$

LEMMA 7.3. *For $X \in \{0, 1\}^{<\log m}$, let $\text{SA}[b..e]$ be the range containing all suffixes of T_A having $\overline{X} \cdot 01 \cdot 1^{\log m}$ as a prefix. Then, $B_X = \text{BWT}[b..e]$.*

The main challenge lies in obtaining the BWT ranges, as the approach of Section 7.1 no longer works. Consider instead partitioning the suffixes in SA according to the length- $\log m$ prefix (separately handling shorter suffixes). Observe that there is a unique bijection $\text{ext} : \{0, 1\}^{\log m} \setminus \{1^{\log m}\} \rightarrow \{0, 1\}^{<\log m} \cdot 01 \cdot 1^{\log m}$ such that X is a prefix of $\text{ext}(X)$. Furthermore, if the range $\text{SA}[b..e]$ contains suffixes of T_A prefixed by $X \in \{0, 1\}^{\log m} \setminus \{1^{\log m}\}$, then the analogous range $\text{SA}[b'..e']$ for $\text{ext}(X)$ satisfies $e' = e$. Thus, it suffices to precompute frequencies $F := \{(X, f_X) : X \in \{0, 1\}^{\log m} \setminus \{1^{\log m}\}\}$ and $F' := \{(X, f'_X) : X \in \{0, 1\}^{<\log m} \cdot 01 \cdot 1^{\log m}\}$.

To this end, construct $F_{\text{pref}} := \{(X, f'_X) : X \in \{0, 1\}^{\leq \log m}\}$, with f'_X defined as the number occurrences of X as prefixes of strings in $\mathcal{A} = (\text{bin}_{\log m}(A[i]))_{i \in [1..m)}$. To compute F_{pref} , we first in $O(m)$ time build $\{(X, f'_X) : X \in \{0, 1\}^{\log m}\} \subseteq F_{\text{pref}}$ by scanning \mathcal{A} . The remaining elements of F_{pref} are derived by the equality $f'_X = f'_{X0} + f'_{X1}$. Analogously prepare suffix frequencies F_{suf} . Given F_{pref} and F_{suf} , we can compute F in $O(m)$ time: the number of occurrences of $X \in \{0, 1\}^{\log m}$ overlapping factors \mathcal{A} is obtained from F_{pref} and F_{suf} ; the number of other occurrences is easily determined as it only depends on m . Finally, F' is computed directly from F_{suf} .

THEOREM 7.4. *If there exists an algorithm that, given the packed representation of a text $T \in \{0, 1\}^n$, constructs its BWT in $O(f(n))$ time, then we can compute the number of inversions in an array of m integers in $O(m + f(m \log m))$ time. In particular, if $f(n) = o(n/\sqrt{\log n})$, then the algorithm runs in $o(m\sqrt{\log m})$ time.*

8 SYNCHRONIZING SETS CONSTRUCTION

Throughout this section, we fix a text T of length n and a positive integer $\tau \leq \frac{1}{2}n$. We also introduce a partition \mathcal{P} of the set $[1..n - \tau + 1]$ so that positions i, j belong to the same class if and only if $T[i..i + \tau] = T[j..j + \tau]$. In other words, each class contains the starting positions of a certain length- τ substring of T . We represent \mathcal{P} using an identifier function $\text{id} : [1..n - \tau + 1] \rightarrow [0..|\mathcal{P}|)$ such that $\text{id}(i) = \text{id}(j)$ if and only if $T[i..i + \tau] = T[j..j + \tau]$.

The local consistency of a string synchronizing set S means that the decision on whether $i \in S$ should be made solely based on $T[i..i + 2\tau]$ or, equivalently, on $\text{id}(i), \dots, \text{id}(i + \tau)$. The density condition, on the other hand, is formulated in terms of a set $R = \{i \in [1..n - 3\tau + 2] : \text{per}(T[i..i + 3\tau - 2]) \leq \frac{1}{3}\tau\}$. Here, we introduce its superset $Q = \{i \in [1..n - \tau + 1] : \text{per}(T[i..i + \tau]) \leq \frac{1}{3}\tau\}$. The periodicity lemma (Lemma 2.1) lets us relate these two sets:

$$R = \{i \in [1..n - 3\tau + 2] : [i..i + 2\tau] \subseteq Q\}. \quad (1)$$

Based on an identifier function id and the set Q , we define a synchronizing set S as follows. Consider a window of size $\tau + 1$ sliding over the identifiers $\text{id}(j)$. For any position i of the window, we compute the smallest identifier $\text{id}(j)$ for $j \in [i..i + \tau] \setminus Q$. We insert i to S if the minimum is attained for $\text{id}(i)$ or $\text{id}(i + \tau)$.

CONSTRUCTION 8.1. For an identifier function id , we define

$$S = \{i \in [1..n - 2\tau + 1] : \min\{\text{id}(j) : j \in [i..i + \tau] \setminus Q\} \in \{\text{id}(i), \text{id}(i + \tau)\}\}.$$

Let us formally argue that this results in a synchronizing set.

LEMMA 8.2. Construction 8.1 always yields a τ -synchronizing set.

PROOF. As for the local consistency of S , observe that if positions i, i' satisfy $T[i..i + 2\tau] = T[i'..i' + 2\tau]$, then $\text{id}(i + \delta) = \text{id}(i' + \delta)$ for $\delta \in [0.. \tau]$. Moreover, $i + \delta \in Q$ if and only if $i' + \delta \in Q$. Consequently, $\{\text{id}(i), \text{id}(i + \tau)\} = \{\text{id}(i'), \text{id}(i' + \tau)\}$ and

$$\min\{\text{id}(j) : j \in [i..i + \tau] \setminus Q\} = \min\{\text{id}(j) : j \in [i'..i' + \tau] \setminus Q\}.$$

Thus, $i \in S$ if and only if $i' \in S$.

To prove the density condition, first assume that $i \notin R$. As (1) yields $[i..i + 2\tau] \setminus Q \neq \emptyset$, we can choose a position j in the latter set with minimum identifier. If $j < i + \tau$, then $j \in S$ due to $\text{id}(j) = \min\{\text{id}(j') : j' \in [j..j + \tau] \setminus Q\}$. Otherwise, $j - \tau \in S$ due to $\text{id}(j) = \min\{\text{id}(j') : j' \in [j - \tau..j] \setminus Q\}$. In either case, $S \cap [i..i + \tau] \neq \emptyset$.

The converse implication is easy: if $i \in R$, then $[i..i + 2\tau] \subseteq Q$ by (1), so $[j..j + \tau] \setminus Q = \emptyset$ and therefore $j \notin S$ for $j \in [i..i + \tau]$. Consequently, $S \cap [i..i + \tau] = \emptyset$. \square

The main challenge in building a τ -synchronizing set with Construction 8.1 is to choose an appropriate identifier function id so that resulting synchronizing set S is small. We first consider only texts satisfying $Q = \emptyset$. Note that this is a stronger assumption than $R = \emptyset$, which we made in the nonperiodic case of Sections 4 to 6.

8.1 The Nonperiodic Case

The key feature of strings without small periods is that their occurrences cannot overlap too much. We say that a set $A \subseteq \mathbb{Z}$ is d -sparse if every two distinct elements $i, i' \in A$ satisfy $|i - i'| > d$.

FACT 8.3. An equivalence class $P \in \mathcal{P}$ is $\frac{1}{3}\tau$ -sparse if $P \cap Q = \emptyset$.

PROOF. Suppose that positions $i, i' \in P$ satisfy $i < i' \leq i + \frac{1}{3}\tau$. We have $T[i..i + \tau] = T[i'..i' + \tau]$, so $\text{per}(T[i..i + \tau]) \leq i' - i \leq \frac{1}{3}\tau$. In particular, $\text{per}(T[i..i + \tau]) \leq \frac{1}{3}\tau$, so $i \in Q$ and $P \cap Q \neq \emptyset$. \square

8.1.1 *Randomized Construction.* It turns out that choosing id uniformly at random leads to satisfactory results if $Q = \emptyset$.

FACT 8.4. Let $\pi : \mathcal{P} \rightarrow [0..|\mathcal{P}|]$ be a uniformly random bijection, and let id be an identifier function such that $\text{id}(j) = \pi(P)$ for each $j \in P$ and $P \in \mathcal{P}$. If $Q = \emptyset$, then a τ -synchronizing set S defined with Construction 8.1 based on such a function id satisfies $\mathbb{E}[|S|] \leq \frac{6n}{\tau}$.

PROOF. Observe that for every position $i \in [1..n - 2\tau + 1]$, we have $|[i..i + \tau] \setminus Q| = |[i..i + \tau]| = \tau + 1$. Moreover, Fact 8.3 guarantees that $|[i..i + \tau] \cap P| \leq 3$ for each class $P \in \mathcal{P}$. Thus, positions in $[i..i + \tau]$ belong to at least $\frac{\tau}{3}$ distinct classes. Each of them has the same probability of having the smallest identifier, so $\mathbb{P}[i \in S] = \mathbb{P}[\min\{\text{id}(j) : j \in [i..i + \tau]\} \in \{\text{id}(i), \text{id}(i + \tau)\}] \leq 2 \cdot \frac{3}{\tau}$ holds for every $i \in [1..n - 2\tau + 1]$. By linearity of expectation, we conclude that $\mathbb{E}[|S|] \leq \frac{6n}{\tau}$. \square

8.1.2 *Deterministic Construction.* Our next goal is to provide an $O(n)$ -time deterministic construction of a synchronizing set S of size $|S| = O(\frac{n}{\tau})$. The idea is to gradually build an identifier function id assigning consecutive identifiers to classes $P \in \mathcal{P}$ one at a time. Our choice of the subsequent classes is guided by a carefully designed scoring function inspired by Fact 8.4.

PROPOSITION 8.5. If $Q = \emptyset$ for a text T and a positive integer $\tau \leq \frac{1}{2}n$, then in $O(n)$ time one can construct a τ -synchronizing set of size at most $\frac{18n}{\tau}$.

PROOF. First, we build the partition \mathcal{P} . A simple $O(n)$ -time implementation is based on the suffix array SA and the LCP table of T (see Section 2.1). We cut the suffix array before every position i such that $\text{LCP}[i] < \tau$, and we remove positions i with $SA[i] > n - \tau + 1$. The values $SA[\ell..r]$ in each remaining maximal region form a single class $P \in \mathcal{P}$. We store pointers to P for all positions $j \in P$.

Next, we iteratively construct the function id represented in a table $\text{id}[1..n - \tau + 1]$. Initially, each value $\text{id}[j]$ is undefined (\perp). In the k th iteration, we choose a class $P_k \in \mathcal{P}$ and set $\text{id}[j] = k$ for $j \in P_k$. Finally, we build a τ -synchronizing set as in Construction 8.1.

Our choice of subsequent classes P_k is guided by a scoring function. To define it, we distinguish *active blocks* which are maximal blocks $\text{id}[\ell..r]$ consisting of undefined values only and satisfying $r - \ell \geq \tau$. Each position $j \in [\ell..r]$ lying in an active block $\text{id}[\ell..r]$ is called an *active position* and assigned a score: -1 if it is among the leftmost or the rightmost $\lfloor \frac{1}{3}\tau \rfloor$ positions within an active block and $+2$ otherwise. Note that the condition $r - \ell \geq \tau$ implies that the total score of positions within an active block is non-negative. Hence, the global score of all active positions is also non-negative.

Our algorithm explicitly remembers whether each position is active and what its score is. Moreover, we store the aggregate score of active positions in each class $P \in \mathcal{P}$ and maintain a collection $\mathcal{P}^+ \subseteq \mathcal{P}$ of unprocessed classes with non-negative scores. Note that the aggregate score is 0 for the already processed classes, so \mathcal{P}^+ is non-empty until there are no unprocessed classes. Hence, we can always choose the subsequent class P_k from \mathcal{P}^+ .

Having chosen the class P_k , we need to set $\text{id}[j] = k$ for $j \in P_k$. If the position $j \in P_k$ was active, then some nearby positions j' may cease to be active or change their scores. This further modifies the aggregate scores of other classes $P \in \mathcal{P}$, some of which may enter or leave \mathcal{P}^+ . Nevertheless, the affected positions j' satisfy $|j - j'| \leq \tau$, so we can implement the changes caused by setting $\text{id}[j] = k$ in $O(\tau)$ time. The total cost of processing P_k is $O(|P_k| + \tau|A_k|)$, where $A_k \subseteq P_k$ consists of positions which were active initially.

Finally, we build the synchronizing set S according to Construction 8.1. Recall that a position $i \in [1..n - 2\tau + 1]$ is inserted to S if and only if $\min\{\text{id}[i..i + \tau]\} \in \{\text{id}[i], \text{id}[i + \tau]\}$. Hence, it suffices to slide a window of fixed width $\tau + 1$ over the table $\text{id}[1..n - \tau + 1]$ computing the sliding-window minima. This takes $O(n)$ time.

To bound the size $|S|$, consider a position i inserted to S , let $k = \min\{\text{id}[i..i + \tau]\}$, and note that $k = \text{id}[i]$ or $k = \text{id}[i + \tau]$. Observe that prior to processing P_k , we had $\text{id}[j] = \perp$ for $j \in [i..i + \tau]$. Consequently, $\text{id}[i..i + \tau]$ was contained in an active block and $i, i + \tau$ were active positions. At least one of them belongs to P_k , so it also belongs to A_k . Hence, if $i \in S$, then $i \in \bigcup_k A_k$ or $i + \tau \in \bigcup_k A_k$, which yields $|S| \leq 2|\bigcup_k A_k| = 2\sum_k |A_k|$. In order to bound this quantity, let us introduce sets $A_k^+ \subseteq A_k$ formed by active positions

that had score +2 prior to processing P_k . The choice of P_k as a class with non-negative aggregate score guarantees that $|A_k| \leq 3|A_k^+|$.

Finally, we shall prove that the set $A^+ := \bigcup_k A_k^+$ is $\frac{1}{3}\tau$ -sparse. Consider two distinct positions $j, j' \in A^+$ such that $j \in A_k^+, j' \in A_{k'}^+$. If $k = k'$, then $|j - j'| > \frac{1}{3}\tau$ because $A_k^+ \subseteq P_k$ is $\frac{1}{3}\tau$ -sparse by Fact 8.3. Hence, we may assume without loss of generality that $k' < k$. Prior to processing P_k , the position j' was not active (the value $\text{id}[j'] = k'$ was already set) whereas j was active and had score +2. Hence, there must have been at least $\lfloor \frac{1}{3}\tau \rfloor$ active positions with score -1 in between, so $|j - j'| \geq \lfloor \frac{1}{3}\tau \rfloor + 1 > \frac{1}{3}\tau$.

Consequently, $|A^+| \leq \lceil \frac{3(n-\tau+1)}{\tau} \rceil \leq \frac{3n}{\tau}$, which implies $|S| \leq 2 \sum_k |A_k| \leq 6 \sum_k |A_k^+| = 6|A^+| \leq \frac{18n}{\tau}$. Moreover, the overall running time is $O(n + \tau \sum_k |A_k|) = O(n + \tau \frac{n}{\tau}) = O(n)$. \square

8.1.3 Efficient Implementation for Small τ . Next, we shall implement our construction in $O(\frac{n}{\tau})$ time if $\tau \leq \varepsilon \log_\sigma n$ and $\varepsilon < \frac{1}{5}$.

Our approach relies on local consistency of the procedure in Proposition 8.5. More specifically, we note that the way this procedure handles a position i depends only on the classes of the nearby positions j with $|j - i| \leq \tau$. In particular, these classes determine how the score of i changes and whether i is inserted to S .

Motivated by this observation, we partition $[1 \dots n - \tau + 1]$ into $O(\frac{n}{\tau})$ blocks so that the b th block contains positions i with $\lceil \frac{i}{\tau} \rceil = b$. We define $T[1 + (b-2)\tau \dots (b+2)\tau]$ to be the *context* of the b th block (assuming that $T[i] = \# \notin \Sigma$ for out-of-bound positions $i \notin [1 \dots n]$), and we say that blocks are *equivalent* if they share the same context.

Based on the initial observation, we note that if two blocks b, b' are equivalent, then the corresponding positions $i = b\tau - \delta$ and $i' = b'\tau - \delta$ (for $\delta \in [0 \dots \tau)$) are processed in the same way by the procedure of Proposition 8.5. This means that we need to process just one *representative block* in each equivalence class.

We can retrieve each context in $O(1)$ time using Proposition 2.3. Consequently, it takes $O(\frac{n}{\tau})$ time to partition the blocks into equivalence classes and construct a family \mathcal{B} of representative blocks. Our choice of τ guarantees that $|\mathcal{B}| = O(1 + \sigma^{4\tau}) = O(n^{4\varepsilon})$. Similarly, the class $P \in \mathcal{P}$ of a position j is determined by $T[j \dots j + \tau]$, so $|\mathcal{P}| = O(n^\varepsilon)$ and the substring can also be retrieved in $O(1)$ time. Hence, the construction procedure has $O(n^\varepsilon)$ iterations. If we spent $O(\tau^{O(1)})$ time for each representative block at each iteration, the overall running time would be $O(n^\varepsilon n^{4\varepsilon} \tau^{O(1)}) = O(n^{5\varepsilon + o(1)}) = O(n^{1 - \Omega(1)}) = O(\frac{n}{\tau})$. This allows for a quite simple approach.

We maintain classes $P \in \mathcal{P}$ indexed by the underlying substrings. For each class, we store the identifier $\text{id}(j)$ assigned to the positions $j \in P$ and a list of positions $j \in P$ contained in the representative blocks. To initialize these components (with the identifier $\text{id}(j)$ set to \perp at first), we scan all representative blocks in $O(\tau^{O(1)})$ time per block, which yields $O(|\mathcal{B}|\tau^{O(1)}) = o(\frac{n}{\tau})$ time in total.

Choosing a class to be processed in every iteration involves computing scores. To determine the score of a particular class $P \in \mathcal{P}$, we iterate over all positions $j \in P$ contained in the representative blocks. We retrieve the class of each fragment j' with $|j - j'| \leq \tau$ in order to compute the score of j . We add this score, multiplied by the number of equivalent blocks, to the aggregate score of P . Having computed the score of each class, we take an arbitrary class P_k with a non-negative score (and no value assigned yet), and we assign the subsequent identifier k to this class. As announced above,

the running time of a single iteration is $O(|\mathcal{B}|\tau^{O(1)})$ as we spend $O(\tau^{O(1)})$ time for each position contained in a representative block.

In the post-processing, we compute S restricted to positions in representative blocks: For every position i contained in a representative block, we retrieve the classes of the nearby positions $j \in [i \dots i + \tau]$ to check whether i should be inserted to S . This takes $O(\tau^{O(1)})$ time per position i , which is $O(|\mathcal{B}|\tau^{O(1)}) = o(\frac{n}{\tau})$ in total.

Finally, we build the whole set S : For each block, we copy the positions of the corresponding representative block inserted to S (shifting the indices accordingly). The running time of this final phase is $O(|S| + \frac{n}{\tau})$, which is $O(\frac{n}{\tau})$ due to $|S| \leq \frac{18n}{\tau}$.

PROPOSITION 8.6. *For every constant $\varepsilon < \frac{1}{5}$, given the packed representation of a text $T \in [0 \dots \sigma]^n$ and a positive integer $\tau \leq \varepsilon \log_\sigma n$ such that $Q = \emptyset$, one can construct in $O(\frac{n}{\tau})$ time a τ -synchronizing set of size $O(\frac{n}{\tau})$.*

8.2 The General Case

In this section, we adapt our constructions so that they work for arbitrary strings. For this, we first study the structure of the set Q .

8.2.1 Structure of Highly Periodic Fragments. The probabilistic argument in the proof of Fact 8.4 relies on the large size of each set $[i \dots i + \tau] \setminus Q$ that we had due to $Q = \emptyset$. However, in general the sets $[i \dots i + \tau] \setminus Q$ can be of arbitrary size between 0 and $\tau + 1$. To deal with this issue, we define the following set (assuming $\tau \geq 2$).

$$B = \left\{ i \in [1 \dots n - \tau + 1] \setminus Q : \text{per}(T[i \dots i + \tau - 1]) \leq \frac{1}{3}\tau \text{ or } \text{per}(T[i + 1 \dots i + \tau]) \leq \frac{1}{3}\tau \right\}$$

(In the special case of $\tau = 1$, we set $B = \emptyset$.) Intuitively, B forms a boundary which separates Q from its complement, as formalized in the fact below. However, it also contains some additional fragments included to make sure that B consists of full classes $P \in \mathcal{P}$.

FACT 8.7 (♣). *If $[l \dots r] \cap Q \neq \emptyset$ and $[l \dots r] \not\subseteq Q$ for two positions $l, r \in [1 \dots n - \tau + 1]$, then $[l \dots r] \cap B \neq \emptyset$.*

We conclude the analysis with a linear-time construction of Q and B , which also reveals an upper bound on $|B|$.

LEMMA 8.8 (♣). *Given a text T and a positive integer τ , the sets Q and B can be constructed in $O(n)$ time. Moreover, $|B| \leq \frac{6n}{\tau}$.*

8.2.2 Randomized Construction. The set B lets us adopt the results of Section 8.1 to arbitrary texts. As indicated in a probabilistic argument, the key trick is to assign the smallest identifiers to classes in B .

FACT 8.9. *There is an identifier function id such that Construction 8.1 yields a synchronizing set of size at most $\frac{18n}{\tau}$.*

PROOF. As in Fact 8.4, we take a random bijection $\pi : \mathcal{P} \rightarrow [0 \dots |\mathcal{P}|)$ and set $\text{id}(j) = \pi(P)$ if $j \in P$. However, this time we draw π uniformly at random only among bijections such that if $P \subseteq B$ and $P' \cap B = \emptyset$ for classes $P, P' \in \mathcal{P}$, then $\pi(P) < \pi(P')$. (Note that each class in \mathcal{P} is either contained in B or is disjoint with this set.)

Consider a position i . Observe that if $[i \dots i + \tau] \cap B \neq \emptyset$, then $i \in S$ holds only if $i \in B$ or $i + \tau \in B$. Hence, the number of such positions $i \in S$ is at most $2|B| \leq \frac{12n}{\tau}$. Otherwise, Fact 8.7 implies that $[i \dots i + \tau] \subseteq Q$ or $[i \dots i + \tau] \cap Q = \emptyset$. In the former case, we are guaranteed that $i \notin S$ by Construction 8.1. On the other hand,

$\mathbb{P}[i \in S] \leq \frac{\epsilon}{\tau}$ holds in the latter case as in the proof of Fact 8.4 since $[i..i+\tau] \setminus Q = [i..i+\tau]$ is of size $\tau+1$. By linearity of expectation, the expected number of such positions $i \in S$ is up to $\frac{6n}{\tau}$.

We conclude that $\mathbb{E}[|S|] \leq \frac{12n}{\tau} + \frac{6n}{\tau} = \frac{18n}{\tau}$. In particular, $|S| \leq \frac{18n}{\tau}$ holds for some identifier function id. \square

8.2.3 Deterministic Construction. Our adaptation of the deterministic construction uses Fact 8.7 and Lemma 8.8 in a similar way. As in the proof of Proposition 8.5, we gradually construct id handling one partition class $P \in \mathcal{P}$ at a time. We start with classes contained in B (in an arbitrary order), then we process classes contained in Q (still in an arbitrary order). In the final third phase, we choose the subsequent classes disjoint with $B \cup Q$ according to their scores.

PROPOSITION 8.10 (♣). *Given a text $T \in [0.. \sigma]^n$ for $\sigma = n^{O(1)}$ and a positive integer $\tau \leq \frac{1}{2}n$, in $O(n)$ time one can construct a τ -synchronizing set of size at most $\frac{30n}{\tau}$.*

8.2.4 Efficient Implementation for Small τ . We conclude by noting that the procedure of Proposition 8.10 can be implemented in $O(\frac{n}{\tau})$ time for $\tau < \epsilon \log_{\sigma} n$ just as we implemented the procedure of Proposition 8.5 to prove Proposition 8.6. The only observation needed to make this seamless adaptation is that we can check in $O(\tau^{O(1)})$ time whether a given position belongs to Q or B. In particular, we have sufficient time to perform these two checks for every position contained in a representative block or its context.

THEOREM 8.11. *For every constant $\epsilon < \frac{1}{5}$, given the packed representation of a text $T \in [0.. \sigma]^n$ and a positive integer $\tau \leq \epsilon \log_{\sigma} n$, one can construct in $O(\frac{n}{\tau})$ time a τ -synchronizing set of size $O(\frac{n}{\tau})$.*

REFERENCES

- Mai Alzamel, Maxime Crochemore, Costas S. Iliopoulos, Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, Juliusz Straszynski, Tomasz Waleń, and Wiktor Zuba. 2019. Quasi-linear-time algorithm for longest common circular factor. In *Proc. CPM*. arXiv:1901.11305 To appear.
- Alberto Apostolico. 1985. The myriad virtues of subword trees. In *Combinatorial Algorithms on Words*. 85–96. https://doi.org/10.1007/978-3-642-82456-2_6
- Maxim Babenko, Paweł Gawrychowski, Tomasz Kociumaka, and Tatiana Starikovskaya. 2015. Wavelet trees meet suffix trees. In *Proc. SODA*. 572–591. <https://doi.org/10.1137/1.9781611973730.39>
- Djamal Belazzougui. 2014. Linear time construction of compressed text indices in compact space. In *Proc. STOC*. 148–193. <https://doi.org/10.1145/2591796.2591885>
- Oren Ben-Kiki, Philip Bille, Dany Breslauer, Leszek Gąsieniec, Roberto Grossi, and Oren Weimann. 2014. Towards optimal packed string matching. *Theor. Comput. Sci.* 525 (2014), 111–129. <https://doi.org/10.1016/j.tcs.2013.06.013>
- Michael A. Bender, Martin Farach-Colton, Giridhar Pemmasani, Steven Skiena, and Pavel Sumazin. 2005. Lowest common ancestors in trees and directed acyclic graphs. *J. Algorithms* 57, 2 (2005), 75–94. <https://doi.org/10.1016/j.jalgor.2005.08.001>
- Or Birenzweig, Shay Golan, and Ely Porat. 2018. Locally consistent parsing for text indexing in small space. arXiv:1812.00359
- Michael Burrows and David J. Wheeler. 1994. *A block-sorting lossless data compression algorithm*. Technical Report 124. Digital Equipment Corporation.
- Timothy M. Chan and Mihai Pătraşcu. 2010. Counting inversions, offline orthogonal range counting, and related problems. In *Proc. SODA*. 161–173. <https://doi.org/10.1137/1.9781611973075.15>
- Yu-Feng Chien, Wing-Kai Hon, Rahul Shah, Sharma V. Thankachan, and Jeffrey Scott Vitter. 2015. Geometric BWT: compressed text indexing via sparse suffixes and range searching. *Algorithmica* 71, 2 (2015), 258–278. <https://doi.org/10.1007/s00453-013-9792-1>
- Maxime Crochemore, Christophe Hancart, and Thierry Lecroq. 2007. *Algorithms on strings*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/cbo9780511546853>
- Martin Farach-Colton, Paolo Ferragina, and S. Muthukrishnan. 2000. On the sorting-complexity of suffix tree construction. *J. ACM* 47, 6 (2000), 987–1011. <https://doi.org/10.1145/355541.355547>
- Paolo Ferragina and Giovanni Manzini. 2005. Indexing compressed text. *J. ACM* 52, 4 (2005), 552–581. <https://doi.org/10.1145/1082036.1082039>
- Nathan J. Fine and Herbert S. Wilf. 1965. Uniqueness theorems for periodic functions. *Proc. Am. Math. Soc.* 16, 1 (1965), 109–114. <https://doi.org/10.2307/2034009>
- Roberto Grossi, Ankur Gupta, and Jeffrey Scott Vitter. 2003. High-order entropy-compressed text indexes. In *Proc. SODA*. 841–850. <http://dl.acm.org/citation.cfm?id=644108.644250>
- Roberto Grossi and Jeffrey Scott Vitter. 2005. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. *SIAM J. Comput.* 35, 2 (2005), 378–407. <https://doi.org/10.1137/s0097539702402354>
- Dan Gusfield. 1997. *Algorithms on strings, trees, and sequences: Computer science and computational biology*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/cbo9780511574931>
- Torben Hagerup. 1998. Sorting and searching on the word RAM. In *Proc. STACS*, Vol. 1373. 366–398. <https://doi.org/10.1007/BFb0028575>
- Yijie Han. 2004. Deterministic sorting in $O(n \log \log n)$ time and linear space. *J. Algorithms* 50, 1 (2004), 96–105. <https://doi.org/10.1016/j.jalgor.2003.09.001>
- Dov Harel and Robert Endre Tarjan. 1984. Fast algorithms for finding nearest common ancestors. *SIAM J. Comput.* 13, 2 (1984), 338–355. <https://doi.org/10.1137/0213024>
- Wing-Kai Hon, Kunihiko Sadakane, and Wing-Kin Sung. 2009. Breaking a time-and-space barrier in constructing full-text indices. *SIAM J. Comput.* 38, 6 (2009), 2162–2178. <https://doi.org/10.1137/070685373>
- Guy Jacobson. 1989. Space-efficient static trees and graphs. In *Proc. FOCS*. 549–554. <https://doi.org/10.1109/SFCS.1989.63533>
- Artur Jež. 2016. Recompression: A simple and powerful technique for word equations. *J. ACM* 63, 1, Article 4 (2016). <https://doi.org/10.1145/2743014>
- Juha Kärkkäinen, Peter Sanders, and Stefan Burkhardt. 2006. Linear word suffix array construction. *J. ACM* 53, 6 (2006), 918–936. <https://doi.org/10.1145/1217856.1217858>
- Dominik Kempa. 2019. Optimal construction of compressed indexes for highly repetitive texts. In *Proc. SODA*. 1344–1357. <https://doi.org/10.1137/1.9781611975482.82>
- Dominik Kempa and Tomasz Kociumaka. 2019. String synchronizing sets: Sublinear-time BWT construction and optimal LCE queries. arXiv:1904.04228
- Tomasz Kociumaka. 2018. *Efficient data structures for internal queries in texts*. PhD thesis. University of Warsaw, Warsaw, Poland. <https://www.mimuw.edu.pl/~kociumaka/files/phd.pdf>
- Tomasz Kociumaka, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. 2015. Internal pattern matching queries in a text and applications. In *Proc. SODA*. 532–551. <https://doi.org/10.1137/1.9781611973730.36>
- Gad M. Landau and Uzi Vishkin. 1988. Fast string matching with k differences. *J. Comput. Syst. Sci.* 37, 1 (1988), 63–78. [https://doi.org/10.1016/0022-0000\(88\)90045-1](https://doi.org/10.1016/0022-0000(88)90045-1)
- Ben Langmead, Cole Trapnell, Mihai Pop, and Steven L Salzberg. 2009. Ultrafast and memory-efficient alignment of short DNA sequences to the human genome. *Genome Biol.* 10, 3, Article R25 (2009). <https://doi.org/10.1186/gb-2009-10-3-r25>
- Heng Li and Richard Durbin. 2009. Fast and accurate short read alignment with Burrows-Wheeler transform. *Bioinformatics* 25, 14 (2009), 1754–1760. <https://doi.org/10.1093/bioinformatics/btp324>
- Veli Mäkinen, Djamal Belazzougui, Fabio Cunial, and Alexandru I. Tomescu. 2015. *Genome-scale algorithm design: Biological sequence analysis in the era of high-throughput sequencing*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/cbo9781139940023>
- Udi Manber and Eugene W. Myers. 1993. Suffix arrays: A new method for on-line string searches. *SIAM J. Comput.* 22, 5 (1993), 935–948. <https://doi.org/10.1137/0222058>
- J. Ian Munro, Gonzalo Navarro, and Yakov Nekrich. 2017. Space-efficient construction of compressed indexes in deterministic linear time. In *Proc. SODA*. 408–424. <https://doi.org/10.1137/1.9781611974782.26>
- J. Ian Munro, Gonzalo Navarro, and Yakov Nekrich. 2017. Text indexing and searching in sublinear time. arXiv:1712.07431
- J. Ian Munro, Yakov Nekrich, and Jeffrey Scott Vitter. 2016. Fast construction of wavelet trees. *Theor. Comput. Sci.* 638 (2016), 91–97. <https://doi.org/10.1016/j.tcs.2015.11.011>
- Gonzalo Navarro. 2014. Wavelet trees for all. *J. Discrete Algorithms* 25 (2014), 2–20. <https://doi.org/10.1016/j.jda.2013.07.004>
- Gonzalo Navarro. 2016. *Compact data structures: A practical approach*. Cambridge University Press, Cambridge, UK. <https://doi.org/10.1017/cbo9781316588284>
- Enno Ohlebusch. 2013. *Bioinformatics algorithms: Sequence analysis, genome rearrangements, and phylogenetic reconstruction*. Oldenbusch Verlag, Ulm, Germany.
- Süleyman Cenk Sahinalp and Uzi Vishkin. 1994. On a parallel-algorithms method for string matching problems. In *Proc. CIAC*. 22–32. https://doi.org/10.1007/3-540-57811-0_3
- Yuka Tanimura, Takaaki Nishimoto, Hideo Bannai, Shunsuke Inenaga, and Masayuki Takeda. 2017. Small-space LCE data structure with constant-time queries. In *Proc. MFCS*. <https://doi.org/10.4230/LIPIcs.MFCS.2017.10>
- Peter Weiner. 1973. Linear pattern matching algorithms. In *Proc. SWAT (FOCS)*. 1–11. <https://doi.org/10.1109/SWAT.1973.13>