

A textbook solution for dynamic strings

Zsuzsanna Lipták ^{a,*}, Francesco Masillo ^b, Gonzalo Navarro ^c

^a Department of Informatics, University of Verona, Italy

^b Department of Computer Science, Technical University of Dortmund, Germany

^c Center for Biotechnology and Bioengineering (CeBiB), Department of Computer Science, University of Chile, Chile

ARTICLE INFO

Prof. Seth Pettie

Keywords:

Dynamic strings

Binary search trees

Splay trees

Dynamic data structures

ABSTRACT

We consider the problem of maintaining a collection of strings while efficiently supporting splits and concatenations on them, as well as comparing two substrings, and computing the longest common prefix between two suffixes. This problem can be solved in optimal time $\mathcal{O}(\log N)$ whp for the updates and $\mathcal{O}(1)$ worst-case time for the queries, where N is the total collection size [Gawrychowski et al., SODA 2018]. We present here a much simpler solution based on a forest of enhanced splay trees (FeST), where both the updates and the substring comparison take $\mathcal{O}(\log n)$ amortized time, n being the sum of the lengths of the strings involved in the operation. The length ℓ of the longest common prefix is computed in $\mathcal{O}(\log n + \log^2 \ell)$ amortized time. Our query results are correct whp. Our simpler solution enables other more general updates in $\mathcal{O}(\log n)$ amortized time, such as reversing a substring and/or mapping its symbols. We can also make FeST use compact space, and extend it to regard substrings as circular or as their omega extension. A C++ implementation of our FeST data structure is available at <https://github.com/fmasillo/FeST>.

1. Introduction

Consider the problem in which we have to maintain a collection of *dynamic strings*, that is, strings we want to modify over time. The modifications may be edit operations such as insertion, deletion, or substitution of a single character; inserting or deleting an entire substring (possibly creating a new string from the deleted substring); adding a fresh string to the collection; etc. In terms of queries, we may want to retrieve a symbol or substring of a dynamic string, determine whether two substrings from anywhere in the collection are equal, or even determine the longest prefix shared by two suffixes in the collection (LCP). The collection must be maintained in such a way that both updates and queries have little cost.

This setup is known in general as the *dynamic strings* problem. A partial and fairly straightforward solution are the so-called ropes, or cords [1]. These are binary trees¹ where the leaves store short substrings, whose left-to-right concatenation forms the string. Ropes were introduced for the Cedar programming language to speed up handling very long strings; a C implementation (termed cords) was also given in the same paper [1]. As the motivating application of ropes/cords was that of implementing a text editor, they support edit operations and extraction/insertion of substrings to enable fast typing and cut&paste, as well as retrieving substrings, but do not support queries like substring equality or LCPs. The trees must be periodically rebalanced to maintain logarithmic times. Recently, a modified version of ropes was implemented for the Ruby language as a basic data type [2]. This variant supports the same updates but does not give theoretical guarantees.

* Corresponding author.

E-mail address: zsuzsanna.liptak@univr.it (Zs. Lipták).

¹ The authors [1] actually state that they are DAGs and referring to them as binary trees is just a simplification. The reason is that the nodes can have more than one parent, so subtrees may be shared.

Table 1

Summary of existing solutions and the solution proposed in this paper. Here, N is the total length of all the strings in the collection, m is the number of updates so far, n is the length of the string(s) affected by the operation, and ℓ is the length of the LCP.

	update time	string eq. time	LCP time
Sundar & Tarjan 1990 [3]	$\mathcal{O}(\sqrt{N \log m} + \log m)$	$\mathcal{O}(\log m)$	—
Mehlhorn et al. 1997 [4]	$\mathcal{O}(\log^2 N)$ randomized, or $\mathcal{O}(\log N (\log m \log^* m + \log N))$ deterministic	$\mathcal{O}(1)$	—
Alstrup et al. 2000 [5]	$\mathcal{O}(\log N \log^* N)$ whp	$\mathcal{O}(\log N)$	$\mathcal{O}(\log N)$
Gawrychowski et al. 2018 [7]	$\mathcal{O}(\log N)$ whp	$\mathcal{O}(1)$	$\mathcal{O}(1)$
Nishimoto et al. 2020 [9]	$\mathcal{O}((n + \log N \log^* N) \frac{(\log \log N)^2}{\log \log \log N})$	$\mathcal{O}(\log n + \log \ell \log^* N)$	$\mathcal{O}(\log n + \log \ell \log^* N)$
Cardinal and Iacono 2021 [10]	$\mathcal{O}(\log N)$ exp. amortized	$\mathcal{O}(\log N)$ exp. amortized	$\mathcal{O}(\log N)$ exp. amortized
this work	$\mathcal{O}(\log n)$ amortized	$\mathcal{O}(\log n)$ amortized (correct whp)	$\mathcal{O}(\log n + \log^2 \ell)$ am. (correct whp)

The first solution we know of that enables equality tests, by Sundar and Tarjan [3], supports splitting and concatenating whole sequences, and whole-string equality in constant time, with updates taking $\mathcal{O}(\sqrt{N \log m} + \log m)$ amortized time, where N is the total length of all the strings in the collection and m is the number of updates so far. It is easy to see that these three primitives encompass all the operations and queries above, except for LCP (substring retrieval is often implicit). The update complexity was soon improved by Mehlhorn et al. [4] to $\mathcal{O}(\log^2 N)$ expected time with a randomized data structure, and $\mathcal{O}(\log N (\log m \log^* m + \log N))$ worst-case time with a deterministic one. The deterministic time complexity was later improved by Alstrup et al. [5] to $\mathcal{O}(\log N \log^* N)$ (which holds with high probability, whp), also computing LCPs in $\mathcal{O}(\log N)$ worst-case time. Recently, Gawrychowski et al. [6,7] obtained $\mathcal{O}(\log N)$ update time whp, retaining constant time to compare substrings, and also decreasing the LCP time to constant, among many other results. They also showed that the problem is essentially closed because just updates and substring equality require $\Omega(\log N)$ time even if allowing amortization. Nishimoto et al. [8,9] showed how to compute LCPs in worst-case time $\mathcal{O}(\log N + \log \ell \log^* N)$, where ℓ is the LCP length, while inserting/deleting substrings of length ℓ in worst-case time $\mathcal{O}((\ell + \log N \log^* N) \frac{(\log \log N)^2}{\log \log \log N})$. Table 1 summarizes these solutions.

All these results build on the idea of parsing a string hierarchically by consistently cutting it into blocks, giving unique names to the blocks, and passing the sequence of names to the next level of parsing. The string is then represented by a parse tree of logarithmic height, whose root consists of a single name, which can be compared to the name at the root of another substring to determine string equality. While there is a general consensus on the fact that those solutions are overly complicated, Gawrychowski et al. [7] mention that

“We note that it is very simple to achieve $\mathcal{O}(\log n)$ update time [...], if we allow the equality queries to give an incorrect result with polynomially small probability. We represent every string by a balanced search tree with characters in the leaves and every node storing a fingerprint of the sequence represented by its descendant leaves. However, it is not clear how to make the answers always correct in this approach [...]. Furthermore, it seems that both computing the longest common prefix of two strings of length n and comparing them lexicographically requires $\Omega(\log^2 n)$ time in this approach.”

This suggestion, indeed, connects to the original idea of ropes [1]. Cardinal and Iacono [10] built on the suggestion to develop a kind of tree dubbed “Data Dependent Tree (DDT)”, which enables updates and LCP computation in $\mathcal{O}(\log N)$ expected amortized time, yet with no errors. DDTs eliminate the chance of errors by ensuring that the fingerprints have no collisions—they simply rebuild all DDTs for all strings in the collection, using a new hash function, when this low-probability event occurs—and reduce the LCP complexity to $\mathcal{O}(\log N)$ by ensuring that subtrees representing the same string have the same shape (so one can descend in the subtrees of both strings synchronously).

In this paper we build on the same suggestion [7], but explore the use of another kind of tree—an enhanced splay tree—which yields a beautifully simple yet powerful data structure for maintaining dynamic string collections. We obtain logarithmic amortized update times for most operations (our cost to compute LCPs lies between logarithmic and squared-logarithmic, see later) and our queries return correct answers whp. The ease of implementation of splay trees makes our solution attractive to be included in a textbook for undergraduate students.

An important consequence of using simpler data structures is that our space usage is $\mathcal{O}(N)$, whereas the solutions based on parsings require in addition $\mathcal{O}(\log N)$ space per update performed, as each one adds a new path to the parse tree. Since the previous parse tree is still available, those structures are *persistent*: one can access any previous version. Our solution is not persistent in principle, but we can make it persistent using $\mathcal{O}(\log n)$ extra space per update or query made so far (we cannot make direct use of the techniques of Driscoll et al. [11]). These add only $\mathcal{O}(1)$ amortized time to the operations.

It would not be hard to obtain *worst-case* times instead of amortized ones, by choosing AVL, α -balanced, or other trees that guarantee logarithmic height. One can indeed find the use of such binary trees for representing strings in the literature [12–14]. Our solution using splay trees has the key advantage of being very simple and easy to understand. The basic operations of splitting

and concatenating strings, using worst-case balanced trees, imply attaching and detaching many subtrees, plus careful rebalancing, which is a nightmare to explain and implement.² Knuth, for example, considered them too complicated to include in his book [16, p. 473] “Deletion, concatenation, etc. It is possible to do many other things to balanced trees and maintain the balance, but the algorithms are sufficiently lengthy that the details are beyond the scope of this book.” Instead, he says [16, p. 478] “A much simpler self-adjusting data structure called a splay tree was developed subsequently [...] Splay trees, like the other kinds of balanced trees already mentioned, support the operations of concatenation and splitting as well as insertion and deletion, and in a particularly simple way.”

1.1. Other related work

A related but distinct line of work aims at maintaining a data structure such that the solution to some particular problem on one or two strings can be efficiently updated when these strings undergo an edit operation (deletion, insertion, or substitution). Examples are longest common factor of two strings [17,18], optimal alignment of two strings [19], approximating the edit distance [20], longest palindromic substring [21], longest square [22], or longest Lyndon factor [23] of one string. The setup can be what is referred to as partially dynamic, when the original string or strings are returned to their state before the edit, or fully dynamic, when the edit operations are reflected on the original string or strings. Clifford et al. [24] give lower bounds on various problems of this kind when a single substitution is applied.

This setup, also referred to as *dynamic strings*, differs from ours in several ways: (a) we are not only interested in solving one specific problem on strings; (b) we have an entire collection of strings, and will want to ask queries on any one or any pair of these; and (c) we allow various kinds of update operations.

Locally consistent parsings to maintain dynamic strings have been used to support more complex problems, such as simulating suffix arrays [25,26].

1.2. Our contribution

We use a splay tree [27], enhanced with additional information, to represent each string in the collection, where all the nodes contain string symbols and Karp-Rabin-like fingerprints [28,29] of the symbols in their subtree. We refer to our data structure as a *forest of enhanced splay trees*, or FeST. As we will see, we can create new strings in $\mathcal{O}(n)$ time, extract substrings of length ℓ in $\mathcal{O}(\ell + \log n)$ time, perform updates and (correctly whp) compare substrings in $\mathcal{O}(\log n)$ time, where n is the length of the strings involved—as opposed to the total length N of all the strings—and the times are amortized (the linear terms are also worst-case). Further, we can compute LCPs correctly whp in amortized time $\mathcal{O}(\log n + \log^2 \ell)$, where ℓ is the length of the returned LCP.

While our LCP time is $\mathcal{O}(\log^2 n)$ for long enough ℓ , LCPs are usually much shorter than the suffixes. For example, in considerably general probabilistic models [30], the maximum LCP value between *any* distinct suffixes of two strings of length n is almost surely $\mathcal{O}(\log n)$, in which case our algorithm runs in $\mathcal{O}(\log n)$ amortized time.

The versatility of our FeST data structure allows us to easily support other kinds of operations, such as reversing or complementing substrings, or both. We can thus implement the reverse complementation of a substring in a DNA or RNA sequence, whereby the substring is reversed and each character is replaced by its Watson-Crick complement. Substring reversal alone is used in classic problems on genome rearrangements where genomes are represented as sequences of genes, and have to be sorted by reversals (see, e.g., [31–36], to cite just a few). Note that chromosomes can be viewed either as permutations or as strings, when gene duplication is taken into account, see Fertin et al. [37]; our FeST data structure accommodates both. We can also implement signed reversals [38,39], another model of evolutionary operation used in genome rearrangements. In general, we can combine reversals with any involution on the alphabet, of which signed or Watson-Crick complementation are only examples. In order to support these operations in $\mathcal{O}(\log n)$ amortized time, we only need to add new constant-space annotations, further enhancing our splay trees while retaining the running times for the other operations. The obvious solution of maintaining modified copies of the strings (e.g., reversed, complemented, etc.) is less attractive in practice due to the extra space and time needed to store and update all the copies.

1.3. Operations supported

We maintain a collection of strings of total length N in $\mathcal{O}(N)$ space, and support the following operations, where we distinguish the basic string data type from dynamic strings (all times are amortized).

- `make-string(w)` creates a dynamic string s from a basic string w , in $\mathcal{O}(|s|)$ time.
- `access(s, i)` returns the symbol $s[i]$ in $\mathcal{O}(\log |s|)$ time.
- `retrieve(s, i, j)` returns the basic string $w[1..j - i + 1] = s[i..j]$, in $\mathcal{O}(|w| + \log |s|)$ time.
- `substitute(s, i, c)`, `insert(s, i, c)`, and `delete(s, i)` perform the basic edit operations on s : substituting $s[i]$ by character c , inserting c at $s[i]$, and deleting $s[i]$, respectively, all in $\mathcal{O}(\log |s|)$ time. For appending c at the end of s one can use `insert($s, |s| + 1, c$)`.
- `introduce(s_1, i, s_2)` inserts s_2 at position i of s_1 (for $1 \leq i \leq |s_1| + 1$), converting s_1 to $s_1[1..i - 1] \cdot s_2 \cdot s_1[i..|s_1|]$ and destroying s_2 , in $\mathcal{O}(\log |s_1 s_2|)$ time.
- `extract(s, i, j)` creates dynamic string $s' = s[i..j]$, removing it from s , in $\mathcal{O}(\log |s|)$ time.

² As an example, an efficient implementation [15] of Rytter’s AVL grammar [12] has over 10,000 lines of C++ code considering only their “basic” variant.

- $\text{equal}(s_1, i_1, s_2, i_2, \ell)$ determines the equality of substrings $s_1[i_1..i_1 + \ell - 1]$ and $s_2[i_2..i_2 + \ell - 1]$ in $\mathcal{O}(\log |s_1 s_2|)$ time, correctly whp.
- $\text{lcp}(s_1, i_1, s_2, i_2)$ computes the length ℓ of the longest common prefix between suffixes $s_1[i_1..]$ and $s_2[i_2..]$, in $\mathcal{O}(\log |s_1 s_2| + \log^2 \ell)$ time, correctly whp, and also tells which suffix is lexicographically smaller.
- $\text{reverse}(s, i, j)$ reverses the substring $s[i..j]$ of s , in $\mathcal{O}(\log |s|)$ time.
- $\text{map}(s, i, j)$ applies a fixed involution (a symbol mapping that is its own inverse) to all the symbols of $s[i..j]$, in $\mathcal{O}(\log |s|)$ time.

We have not chosen a minimal set of primitives. For example, our operations `introduce` and `extract` are easily reducible to the more basic `concat` and `split` used in the literature (the converse is also true, $\text{concat}(s_1, s_2) = \text{introduce}(s_1, |s_1| + 1, s_2)$ and $\text{split}(s, i) = \text{extract}(s, i, |s|)$). Further, with those operations it suffices to support the queries `retrieve` and `equal`, as well as the transformations `reverse` and `map`, on whole strings only. The edit operations can also be simulated with `make-string`, `concat`, and `split`, while `access` is a particular case of `retrieve`. We prefer the richer interface because with splay trees we can give solutions that are less cumbersome than building on the most basic primitives, even if equivalent in time complexity.

Thanks to the simplicity of our solution, our publicly available C++ implementation (<https://github.com/fmasillo/FeST>) takes less than 1000 lines of code and includes all the operations presented on linear strings. It does not depend on external libraries and uses only common C++ standard library functions. The code itself closely follows the explanations reported in this paper, in this way allowing the user to better understand the data structure, as well as enabling them to possibly extend its functionalities by adding custom applications.

1.4. Extensions

FeST uses $\mathcal{O}(N)$ space, that is, words of memory. This can be excessive when the alphabet Σ of the strings is small; think for example of the ASCII alphabet (which needs one byte per string symbol) or DNA (where two bits suffice). We show how FeST can be adapted to use essentially $\mathcal{O}(N \log |\Sigma|)$ bits of space, which is truly proportional to the size of the string data, while retaining all the complexities described above. Once again, there exist solutions for dynamic strings using $N \log_2 |\Sigma| + o(N \log |\Sigma|)$ bits, and even compressed space, that implement most of the operations we support with worst-case time guarantees [40], but these are overly complicated.

Our data structure also enables easy implementation of other features, such as handling circular strings. This is an important and emerging topic [41–45], as many current sequence collections, in particular in computational biology, consist of circular rather than linear strings. Recent data structures built for circular strings [46,47], based on the extended Burrows-Wheeler Transform (eBWT) [48], avoid the detour via the linearization and handle the circular input strings directly. Finally, FeST also allows queries on the omega extensions of strings, that is, on the infinite concatenation $s^\omega = s \cdot s \cdot s \dots$. These occur, for example, in the context of the eBWT, which is based on the so-called omega-order.

1.5. Overview of paper

The rest of the paper is organized as follows. In the next section, we give the necessary technical background (Section 2). In Section 3, we introduce our FeST data structure and standard operations such as creating a new dynamic string, introducing or extracting substrings, or testing substring equality. The following section treats more complex operations such as LCP or substring reversals (Section 4). Then, in Section 5 we show how to reduce the space of our FeST data structure so that it uses compact space. In Section 6 we extend our operations to circular strings and omega extensions of linear strings. This section also includes some open problems on circular strings. We close with some conclusions in Section 7.

This is an extended version of the conference paper [49]. Apart from expanded explanations, we have added new results on compact representation, circular strings, omega extensions, and a reference implementation.

2. Basic concepts

2.1. Strings

We use array-based notation for strings, indexing from 1, so a string s is a finite sequence over a finite ordered alphabet Σ , written $s = s[1..n] = s[1]s[2] \dots s[n]$, for some $n \geq 0$. We assume that the alphabet Σ is integer. The length of s is denoted $|s|$, and ϵ denotes the *empty string*, the unique string of length 0. For $1 \leq i, j \leq |s|$, we write $s[i..j] = s[i]s[i+1] \dots s[j]$ for the *substring* from i to j , where $s[i..j] = \epsilon$ if $i > j$. We write *prefixes* as $s[..i] = s[1..i]$ and *suffixes* as $s[i..] = s[i..|s|]$. Given two strings s, t , their concatenation is written $s \cdot t$ or simply st , and s^k denotes the k -fold concatenation of s , with $s^0 = \epsilon$. A substring (prefix, suffix) of s is called *proper* if it does not equal s .

The *longest common prefix* (LCP) of two strings s and t is defined as the longest string u that is both a prefix of s and t , and $\text{lcp}(s, t) = |u|$ as its length. One can define the lexicographic order based on the lcp: $s <_{\text{lex}} t$ if s is a proper prefix of t , or otherwise if $s[\ell + 1] < t[\ell + 1]$, where $\ell = \text{lcp}(s, t)$.




Fig. 1. Scheme of the zig-zig subroutine applied on node x .




Fig. 2. Scheme of the zig-zag subroutine applied on node x .

2.2. Splay trees

The *splay tree* [27] is a binary search tree that guarantees that a sequence of insertions, deletions, and node accesses costs $\mathcal{O}(\log n)$ amortized time per operation on a tree of n nodes that starts initially empty. In addition, splay trees support splitting and joining trees, both in $\mathcal{O}(\log n)$ amortized time, where n is the total number of nodes involved in the operation.

In the context of binary trees, an edge rotation is a transformation that involves two nodes connected by a single parent-child edge, and flips that edge so that the parent becomes the child or vice versa, while preserving the tree order of all keys. In order to keep the logarithmic amortized time for every operation, splay trees apply a function called *splay* after every query or update on the tree. This operation moves a tree node x to the root by a sequence of so-called “rotations” named zig, zig-zig, zig-zag, and their symmetric versions, each composed of a constant number of edge rotations.

Formally, given a splay tree \mathcal{T} , the $\text{splay}(x)$ function returns a modified splay tree \mathcal{T}' , where the node x is the root of \mathcal{T}' . The tree \mathcal{T}' is obtained deterministically by applying a sequence of rotations, as follows. Let $x(A, B)$ denote a tree rooted at x with left and right subtrees A and B , then the rotation zig-zig converts $z(y(x(A, B), C), D)$ into $x(A, y(B, z(C, D)))$, while the rotation zig-zag converts $z(y(A, x(B, C)), D)$ into $x(y(A, B), z(C, D))$. Whether zig-zig or zig-zag (or their symmetric variants zag-zig or zag-zag) is applied to x depends on its relative position w.r.t. its grandparent. Figs. 1 and 2 illustrate these two rotations. Note that both of these operations are composed by two edge rotations. The rotations are repeated on x as long as possible. Finally, operation zig, which is only applied if x is a child of the root, converts $y(x(A, B), C)$ into $x(A, y(B, C))$ (operation zag is symmetric). Because only edge rotations are involved, splaying preserves the in-order values of the affected nodes.

Every access or update on the tree is thus followed by a *splay* on the deepest reached node. In particular, after finding a node x in a downward traversal, we do $\text{splay}(x)$ to make x the tree root. The goal is that the costs of all the operations are proportional to the cost of all the related *splay* operations performed, so we can focus on analyzing only the splays. Many of the splay tree properties can be derived from a general “access lemma” [27, Lem. 1].

Lemma 1 (Access Lemma [27]). *Let us assign any positive weight $w(x)$ to the nodes x of a splay tree T , and define $sw(x)$ as the sum of the weights of all the nodes in the subtree rooted at x . Then, the amortized time of $\text{splay}(x)$ is $\mathcal{O}(\log(W/sw(x))) \subseteq \mathcal{O}(\log(W/w(x)))$, where $W = \sum_{x \in T} w(x)$.*

The result is obtained by defining $r(x) = \log sw(x)$ (all our logarithms are in base 2) and $\Phi(T) = \sum_{x \in T} r(x)$ as the potential function for the splay tree T . If we choose $w(x) = 1$ for all x , then $W = n$ on a splay tree of n nodes, and thus we obtain $\mathcal{O}(\log n)$ amortized

cost for each operation. By choosing other functions $w(x)$, one can prove splay tree properties like static optimality, the static finger property, and the working set property [27].

The update operations supported by splay trees include inserting new nodes, deleting nodes, joining two trees (where all the nodes in the second tree go to the right of the nodes in the first tree), and splitting a tree into two at some node (where all the nodes to its right become a second tree). The times of those operations are ruled by the “balance theorem with updates” [27, Thm. 6].

Lemma 2 (Balance Theorem with Updates [27]). *Any sequence of access, insert, delete, join and split operations on a collection of initially empty splay trees has an amortized cost of $\mathcal{O}(\log n)$ per operation, where n is the size of the tree(s) where the operation is carried out.*

This theorem is proved with the potential function that assigns $w(x) = 1$ to every node x . Note the theorem considers a forest of splay trees, whose potential function is the sum of the functions $\Phi(T)$ over the trees T in the forest. For details, see the original paper [27].

2.3. Karp-Rabin fingerprinting

Our queries will be correct “with high probability” (whp), meaning a probability of at least $1 - 1/N^c$ for an arbitrarily large constant c , where N is the total size of the collection. This will come from the use of a variant of the original Karp-Rabin fingerprint [28] (cf. [29]) defined as follows. Let $[1..a]$ be the alphabet of our strings and $p \geq a$ a prime number. We choose a random base b uniformly from $[1..p-1]$. The fingerprint κ of string $s[1..n]$ is defined as $\kappa(s) = \left(\sum_{i=0}^{n-1} s[n-i] \cdot b^i \right) \bmod p$. We say that two strings $s \neq s'$ of the same length n collide through κ if $\kappa(s) = \kappa(s')$, that is, $\kappa(s'') = 0$ where $s'' = s - s'$ is the string defined by $s''[i] = (s[i] - s'[i]) \bmod p$. Since $\kappa(s'')$ is a polynomial, in the variable b , of degree at most $n-1$ over the field \mathbb{Z}_p , it has at most $n-1$ roots. The probability of a collision between two strings of length n is then bounded by $(n-1)/(p-1)$ because b is uniformly chosen in $[1..p-1]$. By choosing $p \in \Theta(N^{c+1})$ for any desired constant c , we obtain that κ is collision-free on any $s \neq s'$ whp. We will actually choose $p \in \Theta(N^{c+2})$ because some of our operations perform $\mathcal{O}(\text{polylog } N)$ string comparisons, not just one. Since N varies over time, we can use instead a fixed upper bound, like the total amount of main memory. We use the RAM machine model where logical and arithmetic operations on $\Theta(\log N)$ -bit machine words take constant time.

Two fingerprints $\kappa(s)$ and $\kappa(s')$ can then be composed in constant time to form $\kappa(s' \cdot s) = (\kappa(s') \cdot b^{|s|} + \kappa(s)) \bmod p$. To avoid the $\mathcal{O}(\log |s|)$ time for modular exponentiation, we will maintain the value $b^{|s|} \bmod p$ together with $\kappa(s)$. The corresponding value for $s' \cdot s$ is $(b^{|s'|} \cdot b^{|s|}) \bmod p$, so we can maintain those powers in constant time upon concatenations.

3. Our data structure and standard operations

In this section we describe our data structure called FeST (Forest of enhanced Splay Trees), composed of a collection of (enhanced) splay trees, and then show it implements classical operations on dynamic strings. We will discuss less standard operations in the next section.

3.1. The data structure

We will use a FeST for maintaining the collection of strings, one splay tree per string. A dynamic string $s[1..n]$ is encoded in a splay tree with n nodes such that $s[k]$ is stored in the node x with in-order k (the in-order of a node is the position in which it is listed if we recursively traverse first the left subtree, then the node, and finally the right subtree). We will say that node x represents the substring $s[i..j]$, where $[i..j]$ is the range of the in-orders of all the nodes in the subtree rooted at x . Let T be the splay tree representing string s , then for $1 \leq i \leq |s|$, we call $\text{node}(i)$ the node with in-order i , and for a node x of T , we call $\text{pos}(x)$ the in-order of node x . The root of T is denoted $\text{root}(T)$; we sometimes identify $\text{root}(T)$ with T .

For the amortized analysis of our FeST, our potential function Φ will be the sum of the potential functions $\Phi(T)$ over all the splay trees T representing our string collection. The collection starts initially empty, with $\Phi = 0$. New strings are added to the collection with `make-string`; then edited with `substitute`, `insert`, and `delete`, and redistributed with `introduce` and `extract`.

3.1.1. Information stored at nodes

A node x of the splay tree representing $s[i..j]$ will contain pointers to its left and right children, called $x.\text{left}$ and $x.\text{right}$, its symbol $x.\text{char} = s[\text{pos}(x)]$, its subtree size $x.\text{size} = j - i + 1$, its fingerprint $x.\text{fp} = \kappa(s[i..j])$, and the value $x.\text{power} = b^{j-i+1} \bmod p$. These fields are recomputed in constant time whenever a node x acquires new children $x.\text{left}$ and/or $x.\text{right}$ (e.g., during the splay rotations) with the following formulas: (1) $x.\text{size} = x.\text{left}.\text{size} + 1 + x.\text{right}.\text{size}$, (2) $x.\text{fp} = ((x.\text{left}.\text{fp} \cdot b + x.\text{char}) \cdot x.\text{right}.\text{power} + x.\text{right}.\text{fp}) \bmod p$, and (3) $x.\text{power} = (x.\text{left}.\text{power} \cdot b \cdot x.\text{right}.\text{power}) \bmod p$, as explained in Section 2.3. For the formula to be complete when the left and/or right child is *null*, we assume $\text{null}.\text{size} = 0$, $\text{null}.\text{fp} = 0$, and $\text{null}.\text{power} = 1$. We will later incorporate other fields.

Subtree sizes allow us identify $\text{node}(i)$ given i , in the splay tree T representing string s , in $\mathcal{O}(\log |s|)$ amortized time. This means we can answer `access(s, i)` in $\mathcal{O}(\log |s|)$ amortized time, since $s[i] = \text{node}(i).\text{char}$. Finding $\text{node}(i)$ is done in the usual way, with the recursive function `find(i) = find(T, i)` that returns the i th smallest element in the subtree rooted at the given node. More precisely, `find(x, i) = x` if $i = x.\text{left}.\text{size} + 1$, `find(x, i) = find(x.left, i)` if $i < x.\text{left}.\text{size} + 1$, and `find(x, i) = find(x.right, i - (x.left.size + 1))` if $i > x.\text{left}.\text{size} + 1$. To obtain logarithmic amortized time, `find` splays the node it returns, thus $\text{pos}(\text{root}(T)) = i$ holds after calling `find(T, i)`.

3.1.2. Isolating substrings

We will make use of another primitive we call `isolate(i, j)`, for $1 \leq i, j \leq |s|$ and $i \leq j + 1$, on a tree T representing string s . This operation rearranges T in such a way that $s[i..j]$ becomes represented by one subtree, and returns this subtree's root y .

If $i = 1$ and $j = |s|$, then $y = \text{root}(T)$ and we are done. If $i = 1$ and $j < |s|$, then we find (and splay) `node(j + 1)` using `find(j + 1)`; this will move `node(j + 1)` to the root, and $s[i..j]$ will be represented by the left subtree of the root, so $y = \text{root}(T).\text{left}$. Similarly, if $1 < i$ and $j = |s|$, then we perform `find(i - 1)`, so `node(i - 1)` is splayed to the root and $s[i..j]$ is represented by the right subtree of the root, thus $y = \text{root}(T).\text{right}$.

Finally, if $1 < i, j < |s|$, then splaying first `node(j + 1)` and then `node(i - 1)` will typically result in `node(i - 1)` being the root and `node(j + 1)` its right child, thus the left subtree of `node(j + 1)` contains $s[i..j]$, that is, $y = \text{root}(T).\text{right}.\text{left}$. The only exception arises if the last splay operation on `node(i - 1)` is a zig-zig, as in this case `node(j + 1)` would become a grandchild, not a child, of the root. Therefore, in this case, we modify the last splay operation: if `node(i - 1)` is a grandchild of the root and a zig-zig must be applied, we perform instead two consecutive zig operations on `node(i - 1)` in a bottom-up manner, that is, we first rotate the edge between `node(i - 1)` and its parent, and then the edge between `node(i - 1)` and its new parent (former grandparent), see Fig. 3.

We now consider the effect of the modified zig-zig operation on the potential. In the proof of Lemma 1 [27, Lem. 1], Sleator and Tarjan show that the zig-zig and the zig-zag cases contribute $3(r'(x) - r(x))$ to the amortized cost, where $r'(x)$ is the new value of $r(x)$ after the operation. The sum then telescopes to $3(r(t) - r(x)) = 3 \log(sw(t)/sw(x))$ along an upward path towards a root node t . The zig rotation, instead, contributes $1 + r'(x) - r(x)$, where the 1 would be problematic if it was not applied only once in the path. Our new zig-zig may, at most one time in the path, cost like two zig's, $2 + 2(r'(x) - r(x))$, which raises the cost bound of the whole splay operation from $1 + 3 \log(sw(t)/sw(x))$ to $2 + 3 \log(sw(t)/sw(x))$. This retains the amortized complexity, that is, the amortized time for `isolate` is $\mathcal{O}(\log |s|)$.

3.2. Creating a new dynamic string

Given a basic string $w[1..n]$, operation `make-string(w)` creates a new dynamic string $s[1..n]$ with the same content as w , which is added to the FeST. While this can be accomplished in $\mathcal{O}(n \log n)$ amortized time via successive `ins` operations on an initially empty string, we describe a “bulk-loading” technique that achieves linear worst-case (and amortized) time.

The idea is to create, in $\mathcal{O}(n)$ time, a perfectly balanced splay tree using the standard recursive procedure. As we show in the next lemma, this shape of the tree adds only $\mathcal{O}(n)$ to the potential function, and therefore the amortized time of this procedure is also $\mathcal{O}(n)$.

Lemma 3. *The potential $\Phi(T)$ of a perfectly balanced splay tree T with n nodes is at most $2n + \mathcal{O}(\log^2 n) \subseteq \mathcal{O}(n)$.*

Proof. Let d be the depth of the deepest leaves in a perfectly balanced binary tree, and call $l = d - d' + 1$ the level of any node of depth d' . It is easy to see that there are at most $1 + n/2^l$ subtrees of level l . Those subtrees have at most $2^l - 1$ nodes. Separating the sum $\Phi(T) = \sum_{x \in T} r(x)$ by levels l and using the bound $sw(x) < 2^l$ if x is of level l , we get $\Phi(T) < \sum_{l=1}^{\log n} \left(1 + \frac{n}{2^l}\right) \log 2^l = 2n + \mathcal{O}(\log^2 n)$. \square

Once the tree is created and the fields `x.char` are assigned in in-order, we perform a post-order traversal to compute the other fields. This is done in constant time per node using the formulas given in Section 3.1.

3.3. Retrieving a substring

Given a string s in the FeST and two indices $1 \leq i \leq j \leq |s|$, operation `retrieve(s, i, j)` extracts the substring $s[i..j]$ and returns it as a basic string. The special case $i = j$ is given by `access(s, i)`, which finds `node(i)`, splays it, and returns $(T).\text{char}$, recall Section 3.1. If $i < j$, we perform $y = \text{isolate}(i, j)$ and then we return $s[i..j]$ with an in-order traversal of the subtree rooted at y . Overall, the operation `retrieve(s, i, j)` takes $\mathcal{O}(\log |s|)$ amortized time for `isolate`, and then $\mathcal{O}(j - i + 1)$ worst case time for the traversal of the subtree.

3.4. Edit operations


Let s be a string in the FeST, i an index of s , and c a character. The simplest edit operation, `substitute(s, i, c)` writes c at $s[i]$, that is, s becomes $s' = s[..i - 1] \cdot c \cdot s[i + 1..]$. It is implemented by doing `find(i)` in the splay tree T of s , in $\mathcal{O}(\log |s|)$ amortized time. After the operation, `node(i)` is the root, so we set $\text{root}(T).\text{char} = c$ and recompute (only) its fingerprint as explained in Section 3.1.

Now consider operation `insert(s, i, c)`, which converts s into $s' = s[..i - 1] \cdot c \cdot s[i..]$. This corresponds to the standard insertion of a node in the splay tree, at in-order position i . We first use `find(i)` in order to make $x = \text{node}(i)$ the tree root, and then create a new root node y , with $y.\text{left} = x.\text{left}$ and $y.\text{right} = x$. We then set $x.\text{left} = \text{null}$ and recompute the other fields of x as shown in Section 3.1. Finally, we set $y.\text{char} = c$ and also compute its other fields. By Lemma 2, the amortized cost for an insertion is $\mathcal{O}(\log |s|)$.


Finally, the operation `delete(s, i)` converts s into $s' = s[..i - 1] \cdot s[i + 1..]$. This corresponds to standard deletion in the splay tree: we first do `find(i)` in the tree T of s , so that $x = \text{node}(i)$ becomes the root, and then join the splay trees of $x.\text{left}$ and $x.\text{right}$, isolating the root node x and freeing it. The joined tree now represents s' ; the amortized cost is $\mathcal{O}(\log |s|)$.

3.5. Introducing and extracting substrings

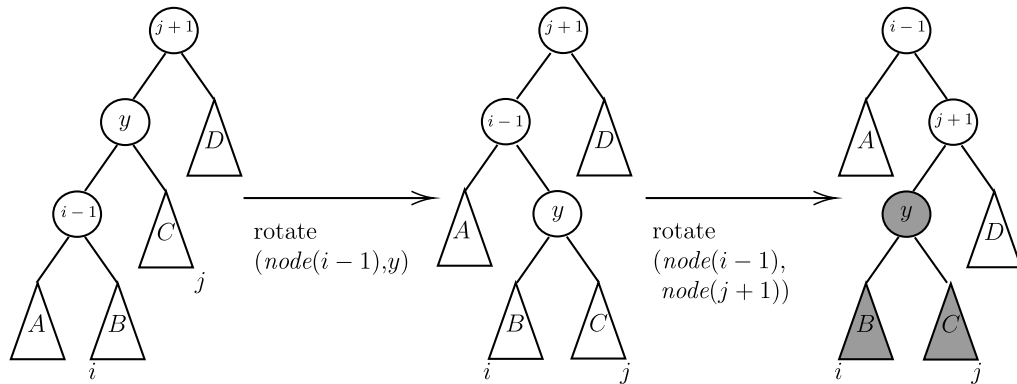
Given two strings s_1 and s_2 represented by trees T_1 and T_2 in the FeST, and an insertion position i in s_1 , operation `introduce(s1, i, s2)` generates a new string $s = s_1[..i - 1] \cdot s_2 \cdot s_1[i..]$ (the original strings are not anymore available). We implement this operation by first



(a) General sequence of operations for $\text{isolate}(i, j)$.



(b) Case of zig-zag as the last splaying operation for $\text{isolate}(i, j)$.



(c) Case of the modified zig-zig as the last splaying operation for $\text{isolate}(i, j)$.

Fig. 3. Scheme of the $\text{isolate}(i, j)$ operation applied on a splay tree. Subfigures 3(a) and 3(b) show two cases of the last splay operation of $\text{isolate}(i, j)$, yielding a single (shaded) subtree that represents the substring $s[i..j]$.

doing $y = \text{isolate}(i, i - 1)$ on the tree T_1 . Note that in this case y will be a *null* node, whose in-order position is between $i - 1$ and i . We then replace this null node by (the root of) the tree T_2 . As shown in Section 3.1, the node y that we replace has at most two ancestors in T_1 , say x_1 (the root) and x_2 . We must then recompute the fields of x_2 and then of x_1 .

Apart from the $\mathcal{O}(\log |s_1|)$ amortized time for isolate , the other operations take constant time. We must consider the change in the potential introduced by connecting T_2 to T_1 . In the potential Φ , the summands $\log sw(x_1)$ and $\log sw(x_2)$ will increase to $\log(sw(x_1) + |s_2|)$ and $\log(sw(x_2) + |s_2|)$, thus the increase is $\mathcal{O}(\log |s_2|)$. The total amortized time is thus $\mathcal{O}(\log |s_1| + \log |s_2|) = \mathcal{O}(\log |s_1 s_2|)$.

Let s be a string represented by tree T in the FeST and $i \leq j$ indices in s . Function $\text{extract}(s, i, j)$ removes $s[i..j]$ from s and creates a new dynamic string s' from it. This can be carried out by first doing $y = \text{isolate}(i, j)$ on T , then detaching y from its parent in T to make it the root of the tree that will represent s' , and finally recomputing the fields of the (former) ancestors x_2 and x_1 of y . The change in potential is negative, as $\log sw(x_1)$ and $\log sw(x_2)$ decrease by up to $\mathcal{O}(\log(j - i + 1))$. The total amortized time is then $\mathcal{O}(\log |s|)$.

3.6. Substring equality

Let $s_1[i_1..i_1 + \ell - 1]$ and $s_2[i_2..i_2 + \ell - 1]$ be two substrings, where possibly $s_1 = s_2$. Per Section 2.3, we can compute equal whp by comparing $\kappa(s_1[i_1..i_1 + \ell - 1])$ and $\kappa(s_2[i_2..i_2 + \ell - 1])$. We compute $y_1 = \text{isolate}(i_1, i_1 + \ell - 1)$ on the tree of s_1 and $y_2 = \text{isolate}(i_2, i_2 + \ell - 1)$ on the tree of s_2 . Once node y_1 represents $s_1[i_1..i_1 + \ell - 1]$ and y_2 represents $s_2[i_2..i_2 + \ell - 1]$, we compare $y_1.\text{fp} = \kappa(s_1[i_1..i_1 + \ell - 1])$ with $y_2.\text{fp} = \kappa(s_2[i_2..i_2 + \ell - 1])$.

The splay operations take $\mathcal{O}(\log |s_1 s_2|)$ amortized time, while the comparison of the fingerprints takes constant time and returns the correct answer whp. Note this is a one-sided error; if the method answers negatively, the strings are distinct.

3.7. String uniqueness

Our string collections allow the existence of repeated dynamic strings. While debatable in a static scenario, this makes all the sense in the dynamic case, since those strings that are equal at some moment may undergo different updates later. If for some reason we wanted to ensure uniqueness (whp), we could maintain a balanced tree with the root fingerprints $\text{root}(T).\text{fp}$ of all the trees T of the FeST. Every string creation or update would then check the tree for uniqueness, increasing the cost by $\mathcal{O}(\log M)$, where $M \leq N$ is the number of strings in the collection.

4. Extended operations

In this section we consider less standard operations of dynamic strings, including the computation of LCPs and others we have not seen addressed before.

4.1. Longest common prefixes

Operation $\text{lcp}(s_1, i_1, s_2, i_2)$ computes $\text{lcp}(s_1[i_1..], s_2[i_2..])$ correctly whp, by exponentially searching for the maximum value ℓ such that $s_1[i_1..i_1 + \ell - 1] = s_2[i_2..i_2 + \ell - 1]$. The exponential search requires $\mathcal{O}(\log \ell)$ equality tests, which are done using equal operations. The amortized cost of this basic solution is then $\mathcal{O}(\log |s_1 s_2| \log \ell)$; we now improve it.

We note that all the accesses the exponential search performs in s_1 and s_2 are at distance $\mathcal{O}(\ell)$ from $s_1[i_1]$ and $s_2[i_2]$. We could then use the dynamic finger property [50] to show, with some care, that the amortized time is $\mathcal{O}(\log |s_1 s_2| + \log^2 \ell)$. This property, however, uses a different mechanism of potential functions where trees cannot be joined or split.³ We then use an alternative approach. The main idea is that, if we could bound ℓ beforehand, we could isolate those areas so that the accesses inside them would cost $\mathcal{O}(\log \ell)$ and then we could reach the desired amortized time. Bounding ℓ in less than $\mathcal{O}(\log \ell)$ accesses (i.e., $\mathcal{O}(\log |s_1 s_2| \log \ell)$ time) is challenging, however. Assuming for now that $s_1 \neq s_2$ (we later handle the case $s_1 = s_2$), our plan is as follows (see Fig. 4):

1. Find a (crude) upper bound $\ell' \geq \ell$.
2. Extract substrings $s'_1 = s_1[i_1..i_1 + \ell' - 1]$ and $s'_2 = s_2[i_2..i_2 + \ell' - 1]$.
3. Run the basic exponential search for ℓ between $s'_1[1..]$ and $s'_2[1..]$.
4. Reinsert substrings s'_1 and s'_2 into s_1 and s_2 .

Steps 2 and 4 are carried out in $\mathcal{O}(\log |s_1 s_2|)$ amortized time using the operations `extract` and `introduce`, respectively. Step 3 will still require $\mathcal{O}(\log \ell)$ substring comparisons, but since they will be carried out on the shorter substrings s'_1 and s'_2 , they will take $\mathcal{O}(\log \ell \log \ell')$ amortized time. The main challenge is to balance the cost to find ℓ' in Step 1 with the quality of the approximation of ℓ' so that $\log \ell'$ is not much larger than $\log \ell$.

Consider the following strategy for Step 1. Let $n = |s_1 s_2|$ and $n' = \min(|s_1| - i_1 + 1, |s_2| - i_2 + 1)$. We first check a few border cases that we handle in $\mathcal{O}(\log n)$ amortized time: if $s_1[i_1..i_1 + n' - 1] = s_2[i_2..i_2 + n' - 1]$ we finish with the answer $\ell = n'$, or else if $s_1[i_1..i_1 + 1] \neq s_2[i_2..i_2 + 1]$ we finish with the answer $\ell = 0$ or $\ell = 1$. Otherwise, we define the sequence $\ell_0 = 2$ and $\ell_i = \min(n', \ell_{i-1}^2)$ and try out the values ℓ_i for $i = 1, 2, \dots$, until we obtain $s_1[i_1..i_1 + \ell_i - 1] \neq s_2[i_2..i_2 + \ell_i - 1]$. This implies that $\ell_{i-1} \leq \ell < \ell_i$, so we can use $\ell' = \ell_i \leq \ell^2$. This yields $\mathcal{O}(\log \ell \log \ell') = \mathcal{O}(\log^2 \ell)$ amortized time for Step 3. On the other hand, since $\ell \geq \ell_{i-1} = 2^{2^{i-1}}$, it holds $i \leq 1 + \log \log \ell$. Since each of the i values is tried out in $\mathcal{O}(\log n)$ time with `equal`, the amortized cost of Step 1 is $\mathcal{O}(\log n \log \log \ell)$ and the total cost to compute `lcp` is $\mathcal{O}(\log n \log \log \ell + \log^2 \ell)$. In particular, this is $\mathcal{O}(\log^2 \ell)$ when ℓ is large enough, $\log \ell = \Omega(\sqrt{\log n \log \log n})$.

4.1.1. Hitting twice

To obtain our desired time $\mathcal{O}(\log n + \log^2 \ell)$ for every value of $\log \ell$, we will apply our general strategy twice. First, we will set $\ell'' = 2^{\log^{2/3} n}$ and determine whether $s_1[i_1..i_1 + \ell'' - 1] = s_2[i_2..i_2 + \ell'' - 1]$. If they are equal, then $\log \ell = \Omega(\log^{2/3} n)$ and we can apply the strategy of the previous paragraph verbatim, obtaining amortized time $\mathcal{O}(\log^2 \ell)$. If they are not equal, then we know that $\ell'' > \ell$, so we extract $s''_1 = s_1[i_1..i_1 + \ell'' - 1]$ and $s''_2 = s_2[i_2..i_2 + \ell'' - 1]$ to complete the search for ℓ' inside those (note we are still in Step 1). We use the same sequence ℓ_i of the previous paragraph, with the only difference that the accesses are done on trees of size ℓ''

³ The static finger property cannot be used either, because we need new fingers every time an LCP is computed. Extending the “unified theorem” [27, Thm. 5] to m fingers (to support m LCP operations in the sequence) introduces an $\mathcal{O}(\log m)$ additive amortized time in the operations, since now $W = \Theta(m)$.

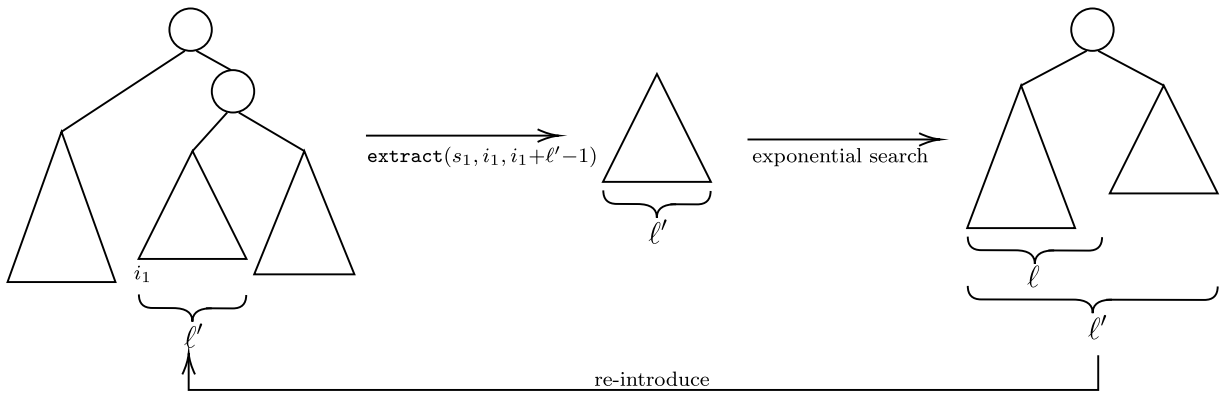


Fig. 4. Scheme of operations for 1cp shown on one of the two strings.

and not n ; therefore each step costs $\mathcal{O}(\log \ell'') = \mathcal{O}(\log^{2/3} n)$ instead of $\mathcal{O}(\log n)$. After finally finding ℓ' , we introduce back s_1'' and s_2'' into s_1 and s_2 . Step 1 then completes in amortized time $\mathcal{O}(\log n + \log^{2/3} n \log \log \ell) = \mathcal{O}(\log n)$. Having found $\ell' \leq \ell^2$, we proceed with Step 2 onwards as above, taking $\mathcal{O}(\log^2 \ell)$ additional time.

4.1.2. When the strings are the same

In the case $s_1 = s_2$, assume w.l.o.g. $i_1 < i_2$. We can still carry out Step 1 and, if $i_1 + \ell' \leq i_2$, proceed with the plan in the same way, extracting s_1' and s_2' from the same string and later reintroducing them. In case $i_1 + \ell' > i_2$, however, both substrings overlap. In this case we extract just one substring, $s' = s_1[i_1..i_2 + \ell' - 1]$, which is of length at most $2\ell'$, and run the basic exponential search between $s'[1..]$ and $s'[i_2 - i_1 + 1..]$ still in amortized time $\mathcal{O}(\log \ell \log \ell')$. We finally reintroduce s' in s_1 . The same is done if we need to extract s_1'' and s_2'' : if both come from the same string and $i_1 + \ell'' > i_2$, then we extract just one single string $s'' = s[i_1..i_2 + \ell'' - 1]$ and obtain the same asymptotic times.

4.1.3. Lexicographic comparisons

Once we know that (whp) the LCP of the suffixes is of length ℓ , we can determine which is smaller by accessing (using `access`) the symbols at positions $s_1[i_1 + \ell]$ and $s_2[i_2 + \ell]$ and comparing them, in $\mathcal{O}(\log |s_1 s_2|)$ additional amortized time.

4.2. Substring reversals

Operation `reverse(s, i, j)` changes s to $s[..i - 1]s[j]s[j - 1] \dots s[i + 1]s[i]s[j + 1..]$. Reflecting it directly in our current structure requires $\Omega(j - i + 1)$ time, which is potentially $\Omega(|s|)$. Our strategy, instead, is to just “mark” the subtrees where the reversal should be carried out, and de-amortize its cost across future operations, materializing it progressively as we traverse the marked subtrees. To this end, we extend our FeST data structure with a new Boolean field $x.rev$ in each node x , which indicates that its whole subtree should be regarded as reversed, that is, its descending nodes should be read right-to-left, but that this update has not yet been carried out. This field is set to `false` on newly created nodes. We also add a field $x.fprev$, so that if x represents $s[i..j]$, then $x.fprev = \kappa(s[j]s[j - 1] \dots s[i + 1]s[i])$ is the fingerprint of the reversed string. When $x.rev$ is `true`, the fields of x (including $x.fp$ and $x.fprev$) still do not reflect the reversal.

The fields $x.fprev$ must be maintained in the same way as the fields $x.fp$. Concretely, upon every update where the children of node x change, we not only recompute $x.fp$ as shown in Section 3.1, but also $x.fprev = ((x.right.fprev \cdot b + x.char) \cdot x.left.power + x.left.fprev) \bmod p$.

In order to apply the described reversal to a substring $s[i..j]$, we first compute $y = \text{isolate}(i, j)$ on the tree of s , and then toggle the Boolean value $y.rev = \neg y.rev$ (note that, if y had already an unprocessed reversal, this is undone without ever materializing it). The operation `reverse` then takes $\mathcal{O}(\log |s|)$ amortized time, dominated by the cost of `isolate(i, j)`. We must, however, handle potentially reversed nodes.

4.2.1. Fixing marked nodes

Every time we access a tree node, if it is marked as reversed, we *fix* it, after which it can be treated as a regular node because its fields will already reflect the reversal of its represented string (though some descendant nodes may still need fixing).

Fixing a node involves exchanging its left and right children, toggling their reverse marks, and updating the node fingerprint. More precisely, we define the primitive `fix(x)` as follows: if $x.rev$ is `true`, then (i) set $x.rev = \text{false}$, $x.left.rev = \neg x.left.rev$, $x.right.rev = \neg x.right.rev$, (ii) swap $x.left$ with $x.right$, and (iii) swap $x.fp$ with $x.fprev$. See Fig. 5 for an example. It is easy to see that `fix` maintains the invariants about the meaning of the reverse fields.

Because all the operations in splay trees, including the `splay`, are done along paths that are first traversed downwards from the root, it suffices that we run `fix(x)` on every node x we find as we descend from the root (for example, on every node x where we




Fig. 5. Scheme of the `fix` operation on node x .

perform `find(x, i)`, before taking any other action on the node. This ensures that all the accesses and structural changes to the splay tree are performed over fixed nodes, and therefore no algorithm needs further changes. For example, when we perform `splay(x)`, all the ancestors of x are already fixed. As another example, if we run `equal` as in Section 3.6, the nodes y_1 and y_2 will already be fixed by the time we read their fingerprint fields. As a third example, if we run `retrieve(s, i, j)` as in Section 3.3 and the subtree of y has reversed nodes inside, we will progressively fix all those nodes as we traverse the subtree, therefore correctly retrieving $s[i..j]$ within $\mathcal{O}(j - i + 1)$ time.

Note that `fix` takes constant time per node and does not change the potential function Φ , so no time complexities change due to our adjustments. The new fields also enable other queries, for example to decide whether a string is a palindrome.

4.3. Involutions

We support the operation `map(s, i, j)` analogously to substring reversals, that is, isolating $s[i..j]$ in a node $y = \text{isolate}(i, j)$ and then marking that the substring covered by node y is mapped using a new Boolean field $y.\text{map}$, which is set to `true`. This will indicate that every symbol $s[k]$, for $i \leq k \leq j$, must be interpreted as $f(s[k])$, but that the change has not yet been materialized. Similarly to `reverse`, this information will be propagated downwards as we descend into a subtree, otherwise it is maintained in the subtree's root only. The operation will then take $\mathcal{O}(\log |s|)$ amortized time.

To manage the mapping and deamortize its linear cost across subsequent operations, we will also store fields $x.\text{mfp} = \kappa(f(s[i])f(s[i+1]) \dots f(s[j]))$ and $x.\text{mfprev} = \kappa(f(s[j])f(s[j-1]) \dots f(s[i]))$, which maintain the fingerprint of the mapped string, and its reverse, represented by x . Those are maintained analogously as the previous fingerprints: (1) $x.\text{mfp} = ((x.\text{left.mfp} \cdot b + f(x.\text{char})) \cdot x.\text{right.power} + x.\text{right.mfp}) \bmod p$, and (2) $x.\text{mfprev} = ((x.\text{right.mfprev} \cdot b + f(x.\text{char})) \cdot x.\text{left.power} + x.\text{left.mfprev}) \bmod p$.

As for string reversals, every time we access a tree node, if it is marked as mapped, we unmark it and toggle the mapped mark of its children, before proceeding with any other action. Precisely, we define the primitive `fixm(x)` as follows: if $x.\text{map}$ is `true`, then (i) set $x.\text{map} = \text{false}$, $x.\text{left.map} = \neg x.\text{left.map}$, $x.\text{right.map} = \neg x.\text{right.map}$, (ii) set $x.\text{char} = f(x.\text{char})$, and (iii) swap $x.\text{fp}$ with $x.\text{mfp}$, and $x.\text{fprev}$ with $x.\text{mfprev}$. We note that, in addition, the `fix` operation defined in Section 4.2 must also exchange $x.\text{mfp}$ with $x.\text{mfprev}$ if we also support involutions. Note how, as for reversals, two applications of f cancel each other, which is correct because f is an involution. Operation `fixm` is applied in the same way as `fix` along tree traversals.

4.3.1. Reverse complementation

By combining string reversals and involutions, we can for example support the application of *reverse complementation* of substrings in DNA sequences, where a substring $s[i..j]$ is reversed and in addition its symbols are replaced by their Watson-Crick complement, applying the involution $f(A) = T$, $f(T) = A$, $f(C) = G$, and $f(G) = C$. In case we *only* want to perform reverse complementation (and not reversals and involutions independently), we can simplify our fields and maintain only a Boolean field $x.\text{rc}$ and the fingerprint $x.\text{mfprev}$ in addition to $x.\text{fp}$. Fixing a node consists of: if $x.\text{rc}$ is `true`, then (i) set $x.\text{rc} = \text{false}$, $x.\text{left.rc} = \neg x.\text{left.rc}$, $x.\text{right.rc} = \neg x.\text{right.rc}$, (ii) set $x.\text{char} = f(x.\text{char})$, (iii) swap $x.\text{left}$ with $x.\text{right}$, (iv) swap $x.\text{fp}$ with $x.\text{mfprev}$.

5. Using compact space

A disadvantage of our scheme is that it uses $\mathcal{O}(N)$ words of space to represent a set of strings of total length N . This is excessive when the alphabet size $|\Sigma|$ of the strings is small compared to the number of bits, w , of the computer word: the M strings can be represented using $Mw + N \lceil \log_2 |\Sigma| \rceil$ bits, while our representation is using $\mathcal{O}(Nw)$ bits. In this section we offer the same functionality we have described using essentially $\mathcal{O}(Mw + N \log |\Sigma|)$ bits of space. We prioritize simplicity in our presentation, retaining the desired time complexities but ignoring possible programming speedups.

Precisely, on a RAM machine of w bits, we will be able to represent M strings of total length N over alphabet Σ , using $\mathcal{O}(M + (N \log |\Sigma|)/w + 2^{\alpha w})$ space, for any fixed constant $0 < \alpha < 1$, so that we will retain the times of all the operations obtained with the basic approach. We will need a preprocessing time and space of $\mathcal{O}(2^{\alpha w})$ for building some “universal” tables that do not depend on the string collection. That preprocessing time and space is, in practice, insignificant; for example, with $\alpha = 1/2$ and under the usual assumption that $w = \Theta(\log N)$, it is just $\mathcal{O}(\sqrt{N})$. The key idea to achieve compact space is to regard s as a shorter string on an enlarged alphabet, exploiting the fact that the space and time of FeST is insensitive to the alphabet size as long as symbols fit in a computer word.

5.1. Working on packed strings

We define a variant of FeST that regards s as a sequence of short strings we call “chunks”, whose length is in $\Theta(w/\log |\Sigma|)$. Concretely, for a constant $0 < \alpha < 1$, the chunks will range from length $l_{\min} = \lfloor (\alpha/2)w/\lceil \log_2 |\Sigma| \rceil \rfloor$ to $l_{\max} = \lfloor \alpha w/\lceil \log_2 |\Sigma| \rceil \rfloor$ (if $l_{\min} = 0$, then we already obtain compact space with the basic FeST using no chunking). Only the first and last chunks of the string can be shorter than l_{\min} .

Given a partition of s into chunks, the “packed” version \tilde{s} of s is a string where each chunk is seen as a single symbol over alphabet $\tilde{\Sigma} = \cup_{l=l_{\min}}^{l_{\max}} \Sigma^l$, which is of size $\leq 2^{\alpha w}$. We represent them as numbers of $l_{\max} \cdot \lceil \log_2 |\Sigma| \rceil \leq \alpha w$ bits, where each symbol of $\tilde{\Sigma}$ is represented in $\lceil \log_2 |\Sigma| \rceil$ bits. Symbols of $\tilde{\Sigma}$ still fit in one computer word, so they can be manipulated in constant time; in particular we can extract any symbol of $\tilde{\Sigma}$ from any chunk in $\mathcal{O}(1)$ time.

If we have M strings s of total length N , the length of all the M packed strings \tilde{s} is $\mathcal{O}(M + N \log |\Sigma|/w)$, and the total space used by FeST on those packed strings is $\mathcal{O}(Mw + N \log |\Sigma|)$ bits. That is, we achieve compact space.

We can create the tree \tilde{T} that represents string \tilde{s} in time $\mathcal{O}(|s|)$, by first building \tilde{s} as a sequence of chunks of s of length, say, $\lfloor (l_{\min} + l_{\max})/2 \rfloor$ (with the last one being possibly shorter), and then creating \tilde{T} on \tilde{s} .

In the rest of the section, we show how we simulate all the operations on strings s using the tree \tilde{T} that represents their packed version \tilde{s} . To do so, we add a new field to the FeST nodes x , called $x.\text{len}$, which stores the sum of the lengths of the chunks that are stored in the subtree rooted at x (instead, $x.\text{size}$ counts the number of chunks). It is easy to maintain $x.\text{len}$ upon rotations and attachment/detachment of subtrees by storing an array L of $2^{\alpha w}$ entries, such that $L[x.\text{char}]$ is the length of the chunk $x.\text{char}$: when the children of x change we recompute $x.\text{len} = x.\text{left}.\text{len} + L[x.\text{char}] + x.\text{right}.\text{len}$. This must also be updated upon changes to $x.\text{char}$ (as in operations `substitute` on \tilde{s} , for example).

Because chunks have variable lengths, the same strings can be represented by different sequences of chunks. We cannot then rely on an arbitrary fingerprint κ built on $\tilde{\Sigma}$, but must instead ensure that $\kappa(\tilde{c}) = \kappa(c_1 \cdots c_l)$ when $\tilde{c} = c_1 \cdots c_l$ (i.e., \tilde{c} must be regarded as a string of symbols in $\tilde{\Sigma}$). We will have a table K of $2^{\alpha w}$ cells with $K[\tilde{c}] = \kappa(\tilde{c})$ precomputed. Further, we will precompute a small table B of l_{\max} entries, so that $B[l] = b^l \bmod p$ and redefine $x.\text{power} = b^{x.\text{len}} \bmod p$. When the children of x change, we then recompute $x.\text{fp} = ((x.\text{left}.\text{fp} \cdot B[L[x.\text{char}]] + K[x.\text{char}]) \cdot x.\text{right}.\text{power} + x.\text{right}.\text{fp}) \bmod p$ and $x.\text{power} = (x.\text{left}.\text{power} \cdot B[L[x.\text{char}]] \cdot x.\text{right}.\text{power}) \bmod p$.

Tables like L and K (and several others of $2^{\alpha w}$ entries we will introduce in the sequel) are easily built in time $\mathcal{O}(2^{\alpha w})$: each entry representing a chunk of l symbols is built in constant time from another representing $l - 1$ symbols.

5.2. Implementing access and retrieve

With the field $x.\text{len}$, we easily implement the primitive `cfind`(\tilde{T}, i), which returns the inorder of the node x of the tree \tilde{T} of \tilde{s} whose chunk $x.\text{char}$ contains $s[i]$, the chunk $x.\text{char}$, and the offset of i within that chunk: starting from $x = \text{root}(\tilde{T})$, if $i \leq x.\text{left}.\text{len}$, then `cfind`(x, i) = `cfind`($x.\text{left}, i$); if $i > x.\text{left}.\text{len} + L[x.\text{char}]$, then `cfind`(x, i) = `cfind`($x.\text{right}, i - x.\text{left}.\text{len} - L[x.\text{char}]$); otherwise `cfind`(x, i) = ($x.\text{left}.\text{size} + 1, x.\text{char}, i - x.\text{left}.\text{len}$). We then apply `splay`(x) on the final node x reached along the recursion.

We implement `access`(s, i) by first computing $(i', \tilde{c}, o) = \text{cfind}(\tilde{T}, i)$ and then extracting extracting the o th symbol of \tilde{c} . The cost is asymptotically the same as that of performing `access`(\tilde{s}, i').

To implement `retrieve`(s, i, j), we first find $(i', \tilde{c}_1, o_1) = \text{cfind}(\tilde{T}, i)$ and $(j', \tilde{c}_2, o_2) = \text{cfind}(\tilde{T}, j)$. If $j' - i' \leq 1$, we have sufficient information to complete the operation in $\mathcal{O}(j - i + 1)$ additional time, by extracting the desired symbols from \tilde{c}_1 and \tilde{c}_2 . Otherwise, we must also extract the intermediate chunks, from the $(i' + 1)$ th to the $(j' - 1)$ th, with `retrieve`($\tilde{s}, i' + 1, j' - 1$).

5.3. Implementing edit operations

Let us now consider the edit operations. The easiest to implement is `substitute`(s, i, c), for which we first compute $(i', \tilde{c}, o) = \text{cfind}(\tilde{T}, i)$. We then replace, in constant time, the o th position of the chunk \tilde{c} by the new symbol c , obtaining the new chunk \tilde{c}' . Finally, we issue `substitute`($\tilde{s}, i', \tilde{c}'$).

The procedure for the operations `insert`(s, i, c) and `delete`(s, i) differs only in how we modify \tilde{c} to obtain \tilde{c}' , requiring constant-time bit shifting and masking on computer words. There is a problem, however, if the resulting chunk \tilde{c}' exceeds l_{\max} symbols upon `insert`, or falls below l_{\min} symbols upon `delete`. In the first case, we split \tilde{c}' into \tilde{c}'_1 and \tilde{c}'_2 , holding $l_{\min} + 1$ and $l_{\max} - l_{\min}$ symbols, and issue `substitute`($\tilde{s}, i', \tilde{c}'_1$) and `insert`($\tilde{s}, i' + 1, \tilde{c}'_2$). In the second case, we obtain the next chunk⁴, $\tilde{c}'' = \text{access}(\tilde{s}, i' + 1)$, and distribute the symbols of \tilde{c}' and \tilde{c}'' into two new chunks \tilde{c}'_1 and \tilde{c}'_2 of permitted lengths, ending with `substitute`($\tilde{s}, i', \tilde{c}'_1$) and `substitute`($\tilde{s}, i' + 1, \tilde{c}'_2$). If this is not possible because the length of \tilde{c}'' is already l_{\min} , we concatenate \tilde{c}' and \tilde{c}'' into \tilde{c}''' , which is below the maximum length l_{\max} , and issue instead `substitute`($\tilde{s}, i', \tilde{c}'''$) and `delete`($\tilde{s}, i' + 1$).

5.4. Implementing introduce and extract

To implement these operations, and several that follow, we must implement the equivalent of `isolate` on \tilde{T} . Concretely, primitive `cisolate`(\tilde{T}, i, j) rearranges \tilde{T} so that $s[i..j]$ (note this is the original string) becomes represented by one subtree of \tilde{T} , and returns that subtree’s root y .

The basic plan to implement `cisolate`(\tilde{T}, i, j) is as follows:

⁴ If i' is the last position of \tilde{s} , we can allow \tilde{c}' be shorter than l_{\min} .

1. We compute $(i', \tilde{c}_1, o_1) = \text{cfind}(\tilde{T}, i)$. If $o_1 > 1$ (i.e., \tilde{c}_1 does not start at $s[i]$) we split, as done for overflows in `insert`, \tilde{c}_1 into \tilde{c}' (holding the first $o_1 - 1$ symbols of \tilde{c}_1) and \tilde{c}'' (holding the rest); then we do `substitute`($\tilde{T}, i', \tilde{c}'$) and `insert`($\tilde{T}, i' + 1, \tilde{c}''$), and increment i' .
2. We compute $(j', \tilde{c}_2, o_2) = \text{cfind}(\tilde{T}, j)$. If $o_2 < L[\tilde{c}_2]$ (i.e., \tilde{c}_2 does not end at $s[j]$) we split \tilde{c}_2 into \tilde{c}' (holding the first o_2 symbols of \tilde{c}_2) and \tilde{c}'' (with the rest), and then `substitute`($\tilde{T}, j', \tilde{c}'$) and `insert`($\tilde{T}, j' + 1, \tilde{c}''$).
3. Now that $s[i..j]$ spans a sequence of whole chunks, $\tilde{s}[i'..j']$, we execute $y = \text{isolate}(\tilde{T}, i', j')$ and return y .

This basic plan needs some fixes in case the resulting chunks \tilde{c}' and \tilde{c}'' are too short. In points 1 and 2 above, if \tilde{c}' is too short and not the first symbol in \tilde{s} , we merge it with the preceding chunk (and re-split them if needed) as done for `delete`; if \tilde{c}'' is too short and not the last symbol in \tilde{s} , we merge it with the following chunk (and re-split them if needed). An exception occurs if, in point 2 above, \tilde{c}' is too short but it is the unique chunk spanning $s[i..j]$; in this case we let it be short (this is illegal in \tilde{T} , but it will be fixed soon in the operations that use `isolate`). Note that values i' and j' may have to be updated upon chunk merges not followed by re-splits.

We now easily implement `introduce`(s_1, i, s_2), where each \tilde{s}_i is represented by \tilde{T}_i : we call $y = \text{isolate}(\tilde{T}_1, i, i - 1)$ and replace node y (which contains an empty chunk) by the root of \tilde{T}_2 . If the first or last chunks of \tilde{T}_2 were too short (which was allowed), we must fix this by merging them with their new preceding or following chunk, respectively, and possibly re-splitting.

For `extract`(s, i, j) with \tilde{s} represented by \tilde{T} , we call $y = \text{isolate}(\tilde{T}, i, j)$, detach y from \tilde{T} , and return it as the root of the tree representing $s[i..j]$.

5.5. Implementing `equal` and `lcp`

To implement `equal`(s_1, i_1, s_2, i_2, ℓ), we distinguish two cases. If $\ell \leq l_{\max}$, we materialize chunks corresponding to $s_1[i_1..i_1 + \ell - 1]$ and to $s_2[i_2..i_2 + \ell - 1]$, doing as follows for each one: we obtain \tilde{c}_1 and \tilde{c}_2 with `cfind`, as done for `retrieve`, and use bit shifting and masking to produce, in constant additional time, the chunk from \tilde{c}_1 and \tilde{c}_2 (it might also be that $\tilde{c}_1 = \tilde{c}_2$ and the desired chunk is inside it). We then compare the two materialized chunks, without error.

If, instead, $\ell > l_{\max}$, we resort to fingerprinting and can err with low probability. We use $y_1 = \text{isolate}(\tilde{T}_1, i_1, i_1 + \ell - 1)$ and $y_2 = \text{isolate}(\tilde{T}_2, i_2, i_2 + \ell - 1)$, and compare $y_1.\text{fp}$ with $y_2.\text{fp}$.

The computation of $\ell = \text{lcp}(s_1, i_1, s_2, i_2)$ is built on top of `equal`, and as such can be implemented verbatim on the trees \tilde{T}_1 and \tilde{T}_2 that represent \tilde{s}_1 and \tilde{s}_2 , in time $O(\log n + \log^2 \ell)$.

5.6. Implementing reversals and involutions

To implement `reverse`(s, i, j), we compute $y = \text{isolate}(\tilde{T}, i, j)$ and set the flag $y.\text{rev}$. Those flags are fixed upon traversals, just as in the original solution, including when we apply the new operation `cfind` on \tilde{T} . A difference with the original `fix` operation is that, this time, `fix`(x) must also reverse the chunk $x.\text{char}$. To do this in constant time, we maintain a table R with 2^{aw} entries, so that $R[x.\text{char}]$ is the reversed string of $x.\text{char}$. We also maintain the fields $x.\text{fprev}$ analogously to $x.\text{fp}$, using that $\kappa(c_1 \dots c_\ell) = K[R[c_1 \dots c_\ell]]$.

For implementing `map`(s, i, j), we will have another precomputed table F of 2^{aw} entries, so that $F[\tilde{c}] = f(c_1) \dots f(c_\ell)$ for every chunk $\tilde{c} = c_1 \dots c_\ell$. To apply f over $s[i..j]$, we compute $y = \text{isolate}(\tilde{T}, i, j)$ and set the flag $y.\text{map}$. The operation `fixm`(x) used to fix the mapping works as before, except that we set $x.\text{char} = F[x.\text{char}]$ instead of $x.\text{char} = f(x.\text{char})$. We again maintain the fields $x.\text{mfp}$ using that $\kappa(f(c_1) \dots f(c_\ell)) = K[F[c_1 \dots c_\ell]]$.

An exception in both cases arises when $j - i + 1 < l_{\min}$, as in this case the node y obtained from `isolate` is too short. In such a case, instead of the above procedures, we directly update $y.\text{char}$ to $R[y.\text{char}]$ or $F[y.\text{char}]$, and then apply the same procedure used by `equal` when a chunk becomes too short.

6. Circular strings and omega extensions

In this section, we explain how to extend our data structure to circular strings and to omega extensions of strings (i.e., to the infinite string that results from concatenating a finite string an infinite number of times).

A new routine, called `rotate`, will be fundamental in handling circular strings; it allows us linearize the circular string starting at any of its indices. By carefully using this primitive, along with a slight modification for the computation of fingerprints, we can support every operation that we presented on linear strings with the same time bounds. By supporting operations on circular strings, we can also handle the omega extension of strings, and will see that here, too, we are able to meet the same time bounds on every operation as on linear strings.

We start by introducing further terminology needed in this section.

6.1. Additional definitions

A string s is *periodic* with period r if $s[i + r] = s[i]$ for all $1 \leq i \leq |s| - r$. Two strings s, t are *conjugates* if there exist strings u, v , possibly empty, such that $s = uv$ and $t = vu$. Conjugacy is an equivalence; the equivalence classes $[s]$ are also called *circular strings*, and any $t \in [s]$ is called a *linearization* of this circular string. Abusing notation, any linear string s can be viewed as a circular string, in which case it is taken as a representative of its conjugacy class. A *substring* of a circular string s is any prefix of any $t \in [s]$, or, equivalently, a string of the form $s[i..j]$ for $1 \leq i, j \leq |s|$ (a linear substring), or $s[i..|s|..j]$, where $j < i$. A *necklace* is a string s with the property that $s \leq_{\text{lex}} t$ for all $t \in [s]$. Every conjugacy class contains exactly one necklace.

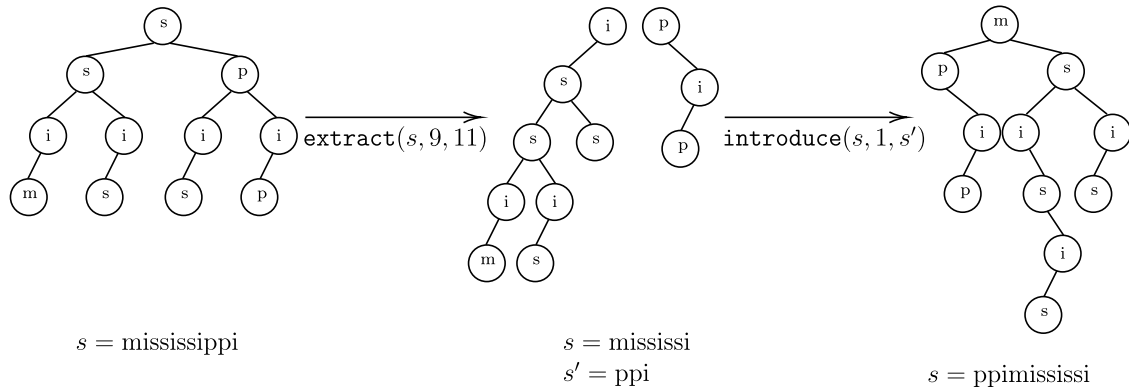


Fig. 6. Cycle-rotation operation: $\text{rotate}(s, 9)$ moves $s[9..]$ to the left of $s[..8]$. After the rotation the string becomes $s[9..]s[..8]$.

When the dynamic strings in our collection are to be interpreted as circular strings, we need to adjust some of our operations. Our model is that we will maintain a canonical representative \hat{s} of the class of rotations of s . All the indices of the operations refer to positions in \hat{s} . Internally, we may store in the FeST another representative s of the class, not necessarily \hat{s} .

6.2. Circular strings

Our general approach to handle operations on \hat{s} regarding it as circular is to rotate it conveniently before accessing it. The splay tree T of \hat{s} will then maintain some (string) rotation $s = \hat{s}[r..] \hat{s}[..r-1]$ of \hat{s} , and we will maintain a field $\text{start}(\hat{s}) = r$ so that we can map any index $\hat{s}[i]$ referred to in update or query operations to $s[(|s| + i - \text{start}(\hat{s})) \bmod |s| + 1]$.

When we want to change the rotation of \hat{s} to another index r' , so that we now store $s' = \hat{s}[r'..] \hat{s}[..r'-1]$, we make use of a new operation $\text{rotate}(s, i)$, which rotates s so that its splay tree represents $s[i..]s[..i-1]$. This is implemented as $\text{extract}(s, i, |s|)$, which splits the substring $s' = s[i..|s|]$ from s producing two distinct trees, followed by $\text{introduce}(s, 1, s')$. See Fig. 6 for an example. We then move from rotation r to r' in $\mathcal{O}(\log |s|)$ amortized time by doing $\text{rotate}(s, r' - r + 1)$ if $r' > r$, or $\text{rotate}(s, |s| + r' - r + 1)$ if $r' < r$. We then set $\text{start}(\hat{s}) = r'$.

Operation $s = \text{make-string}(w)$ stays as before, in the understanding that $\hat{s} = w$ will be seen as the canonical representation of the class, so we set $\text{start}(\hat{s}) = 1$; this can be changed later with a string rotation if desired. All the operations that address a single position $\hat{s}[i]$, like access and the edit operations, are implemented verbatim by just shifting the index i using $\text{start}(\hat{s})$ as explained. Instead, the operations retrieve , extract , equal , reverse , and map , which act on a range $\hat{s}[i..j]$, may give trouble when $i > j$, as in this case the substring is $\hat{s}[i..] \hat{s}[..j]$ by circularity. In this case, those operations will be preceded by a change of rotation from the current one, $r = \text{start}(\hat{s})$, to $r' = 1$, using rotate as explained. This guard will get rid of those cases. Note that, in the case of equal , we may need to rotate both s_1 and s_2 , independently, to compute each of the two signatures.

The two remaining operations deserve some consideration. Operation $\text{introduce}(s_1, i, s_2)$ could be implemented verbatim (with the shifting of i), but in this case it would introduce in $\hat{s}_1[i]$ the current rotation of s_2 , instead of \hat{s}_2 as one would expect. Therefore, we precede the operation by a change of rotation in s_2 to $r' = 1$, which makes the splay tree store \hat{s}_2 with $\text{start}(\hat{s}_2) = 1$.

Finally, in operation $\text{lcp}(s_1, i_1, s_2, i_2)$ we do not know for how long the LCP will extend, so we precede it by changes of rotations in both s_1 and s_2 that make them start at position 1 of \hat{s}_1 and \hat{s}_2 . In case $s_1 = s_2$, however, this trick cannot be used. One simple solution is to rotate the string every time we call equal during Step 1; recall Section 4.1. This will be needed as long as the accesses are done on s_1 and s_2 ; as soon as we extract the substrings of length ℓ'' (and, later, ℓ' for Step 3), we work only on the extracted strings. While the complexity is preserved, rotating the string every time can be too cumbersome. We can use an alternative way to compute signatures of circular substrings, $\kappa(s[i..]s[..j])$: we compute as in Section 3.6 $\sigma = \kappa(s[i..])$ and $\tau = \kappa(s[..j])$, as well as $b^j \bmod p$, which comes for free with the computation of τ ; then $\kappa(s[i..]s[..j]) = (\sigma \cdot b^j + \tau) \bmod p$.

Overall, we maintain for all the operations the same asymptotic running times given in the Introduction when the strings are interpreted as circular.

6.2.1. Signed reversals on circular strings

By combining reversals and involutions, we can support signed reversals on circular strings, too. Note that the original paper in which reversals were introduced [31] used circular chromosomes. In this application, circular chromosomes can be viewed as circular strings over an integer alphabet, where each gene in the chromosome has its own gene identifier taken from Σ , for example chromosome = 162534 and $\Sigma = \{1, 2, 3, 4, 5, 6\}$. To perform a signed reversal, we do this in the same way as for linear strings, namely by doubling the alphabet Σ of gene identifiers such that each gene i has a negated version $-i$, and using the involution $f(i) = -i$ (and $f(-i) = i$).

6.3. Omega extensions

Circular dynamic strings allow us to implement operations that act on the omega extensions of the underlying strings. Recall that for a (linear) string s , the infinite string s^ω is defined as the infinite concatenation $s^\omega = s \cdot s \cdot s \dots$. These are, for example, used in the definition of the *extended Burrows-Wheeler Transform* (eBWT) of Mantaci et al. [48], where the underlying string order is based on omega extensions. In this case, comparisons of substrings may need to be made whose length exceeds the shorter of the two strings s_1 and s_2 . We therefore introduce a generalization of circular substrings as follows: t is called an *omega-substring* of s if $t = s[i..]s^k s[..j]$ for some $j < i - 1$ and $k \geq 0$. Note that the suffix $s[i..]$ and the prefix $s[..j]$ may also be empty. Thus, t is an omega-substring of s if and only if $t = v^k v[..j]$ for some $k \geq 1$ and some conjugate v of s .

An important tool in this section will be the famous Fine and Wilf Lemma [51], which states that if a string w has two periods r, q and $|w| \geq r + q - \gcd(r, q)$, then w is also periodic with period $\gcd(r, q)$ (a string s is called periodic with period r if $s[i + r] = s[i]$ for all $1 \leq i \leq |s| - r$). The following is a known corollary, a different formulation of which was proven, e.g., in [48]; we reprove it here for completeness.

Lemma 4. *Let u, v be two strings. If $\text{lcp}(u^\omega, v^\omega) \geq |u| + |v| - \gcd(|u|, |v|)$, then $u^\omega = v^\omega$.*

Proof. Let $\ell = \text{lcp}(u^\omega, v^\omega) \geq |u| + |v| - \gcd(|u|, |v|)$. Then the string $t = s_1^\omega[..\ell]$ is periodic both with period $|u|$ and with period $|v|$, and thus, by the Fine and Wilf lemma, it is also periodic with period $\gcd(|u|, |v|)$. Since $\gcd(|u|, |v|) \leq |u|, |v|$, this implies that both u and v are powers of the same string x , of length $\gcd(|u|, |v|)$ and therefore, $u^\omega = x^\omega = v^\omega$. \square

We further observe that the fingerprint of strings of the form u^k can be computed from the fingerprint of string u . More precisely, let u be a string, $\pi = \kappa(u)$ its fingerprint, and $k \geq 1$. Then, calling $d = b^{|u|} \bmod p$ (which we also obtain in the field y .power when computing $\kappa(u)$), it holds

$$\begin{aligned} \kappa(u^k) &= (\pi \cdot d^{k-1} + \pi \cdot d^{k-2} + \dots + \pi \cdot d + \pi) \bmod p \\ &= (\pi \cdot (d^{k-1} + d^{k-2} + \dots + 1)) \bmod p, \end{aligned} \tag{1}$$

where $\text{geomsum}(d, k - 1) = (d^{k-1} + d^{k-2} + \dots + 1) \bmod p$ can be computed in $\mathcal{O}(\log k)$ time using the identity $d^{2k+1} + d^{2k} + \dots + 1 = (d + 1) \cdot ((d^2)^k + (d^2)^{k-1} + \dots + 1)$, as follows⁵ (all modulo p):

$$\begin{aligned} \text{geomsum}(d, 0) &= 1 \\ \text{geomsum}(d, 2k + 1) &= (d + 1) \cdot \text{geomsum}(d^2, k) \\ \text{geomsum}(d, 2k) &= d \cdot \text{geomsum}(d, 2k - 1) + 1 \end{aligned} \tag{2}$$

6.3.1. Extended substring equality

We devise at least two ways in which our equal query can be extended to omega extensions. First, consider the query $\text{equal}_\omega(s_1, i_1, s_2, i_2, \ell) = \text{equal}(s_1^\omega, i_1, s_2^\omega, i_2, \ell)$, that is, the normal substring equality interpreted on the omega extensions of s_1 and s_2 . We let $v_1 = \text{rotate}(s_1, i_1)$ and $v_2 = \text{rotate}(s_2, i_2)$. Then we have $s_1^\omega[i_1..i_1 + \ell - 1] = v_1^{k_1} v_1[..j_1]$, where $k_1 = \lfloor \ell / |s_1| \rfloor$ and $j_1 = \ell \bmod |s_1|$. If $k_1 = 0$, we simply compute $\kappa_1 = \kappa(s_1^\omega[i_1..i_1 + \ell - 1]) = \kappa(v_1[..j_1])$. Otherwise, we compute $\kappa_1 = \kappa(s_1^\omega[i_1..i_1 + \ell - 1])$ by applying Eq. (1) as follows:

$$\kappa_1 = (\kappa(v_1) \cdot (d^{k_1-1} + \dots + 1) \cdot b^{j_1} + \kappa(v_1[..j_1])) \bmod p. \tag{3}$$

There are various components to compute in this formula apart from the fingerprints themselves. First, note that $d = b^{|s_1|} \bmod p = b^{|v_1|} \bmod p = \text{root}(T_1)$.power for the tree T_1 of s_1 (or v_1), so we have it in constant time. Second, $b^{j_1} \bmod p$ is the field y .power after we compute $\kappa(v_1[..j_1])$ via $y = \text{isolate}(v_1, 1, j_1)$ after completion of $\text{rotate}(s_1, i_1)$, thus we also have it in constant time. Third, $d^{k_1-1} + \dots + 1 = \text{geomsum}(d, k_1 - 1)$ is computed with Eq. (2) in time $\mathcal{O}(\log k_1) \subseteq \mathcal{O}(\log \ell)$.

By Lemma 4 we can define $\ell_\omega = |s_1| + |s_2|$ and, if $\ell \geq \ell_\omega$, run the equal_ω query with ℓ_ω instead of ℓ . The lemma shows that $s_1[i_1..i_1 + \ell - 1] = s_2[i_2..i_2 + \ell - 1]$ iff $s_1[i_1..i_1 + \ell_\omega - 1] = s_2[i_2..i_2 + \ell_\omega - 1]$. This limits ℓ to $|s_1| + |s_2|$ in our query and therefore the cost $\mathcal{O}(\log \ell)$ is in $\mathcal{O}(\log |s_1 s_2|)$.

We compute κ_2 analogously, and return *true* if and only if $\kappa_1 = \kappa_2$, after undoing the rotations to get back the original strings s_1 and s_2 . The total amortized time for operation equal_ω is then $\mathcal{O}(\log |s_1 s_2|)$. Note that our results still hold whp because we are deciding on fingerprints of strings of length $\mathcal{O}(N)$, not $\mathcal{O}(\ell)$ (which is in principle unbounded).

A second extension of equal is $\text{equal}_\omega^\omega(s_1, i_1, \ell_1, s_2, i_2, \ell_2)$, interpreted as $(s_1^\omega[i_1..i_1 + \ell_1 - 1])^\omega = (s_2^\omega[i_2..i_2 + \ell_2 - 1])^\omega$, that is, the omega extension of $s_1^\omega[i_1..i_1 + \ell_1 - 1]$ is equal to the omega extension of $s_2^\omega[i_2..i_2 + \ell_2 - 1]$. By Lemma 4, this is equivalent to $(s_1^\omega[i_1..i_1 + \ell_1 - 1])^{\ell_2} = (s_2^\omega[i_2..i_2 + \ell_2 - 1])^{\ell_1}$. So we first compute $\kappa_1 = \kappa(s_1^\omega[i_1..i_1 + \ell_1 - 1])$ and $\kappa_2 = \kappa(s_2^\omega[i_2..i_2 + \ell_2 - 1])$ as above, compute $d_1 = b^{\ell_1} \bmod p$ and $d_2 = b^{\ell_2} \bmod p$, and then return whether $(\kappa_1 \cdot (d_1^{\ell_2-1} + \dots + 1)) \bmod p = (\kappa_2 \cdot (d_2^{\ell_1-1} + \dots + 1)) \bmod p$. Operation $\text{equal}_\omega^\omega$ is then also computed in amortized time $\mathcal{O}(\log |s_1 s_2|)$.

⁵ This technique seems to be folklore. Note that the better known formula $\text{geomsum}(d, k) = ((d^{k+1} - 1) \cdot (d - 1)^{-1}) \bmod p$ requires computing multiplicative inverses, which takes $\mathcal{O}(\log N)$ time using the extended Euclid's algorithm, or $\mathcal{O}(\log \log N)$ with faster algorithms [52]; those terms would not be absorbed by others in our cost formula.

6.3.2. Extended longest common prefix

We are also able to extend LCPs to omega extensions: the operation $\text{lcp}_\omega(s_1, i_1, s_2, i_2)$ computes, for the corresponding rotations $v_1 = \text{rotate}(s_1, i_1)$ and $v_2 = \text{rotate}(s_2, i_2)$, the longest common prefix length $\text{lcp}(v_1^\omega, v_2^\omega)$, as well as the lexicographic order of v_1^ω and v_2^ω . That this can be done efficiently follows again from Lemma 4. We first compare their omega-substrings of length $\ell_\omega = |s_1| + |s_2|$. If $\text{equal}_\omega(s_1, i_1, s_2, i_2, \ell_\omega)$ answers *true*, then it follows that $\text{lcp}(s_1, i_1, s_2, i_2)$ is ∞ . Otherwise, we run a close variant of the algorithm described in Section 4.1; note that ℓ_ω can be considerably larger than one of s_1 or s_2 . For Step 1, we define $n' = n = |s_1 s_2|$; the other formulas do not change. We run the equal_ω computations on s_1 and s_2 using Eq. (3) to compute the fingerprints. We extract the substrings of length ℓ' in Step 3 (analogously, ℓ'' in Step 1) using the `extract` for circular strings, but do so only if $\ell' \leq |s_1|$ (resp., $\ell' \leq |s_2|$); otherwise we keep accessing the original string using Eq. (3). The total amortized time to compute LCPs on omega extensions is thus $\mathcal{O}(\log |s_1 s_2|)$.

6.4. Future work on circular strings

One feature that we would like to add to our data structure is allowing identification of conjugates. The rationale behind this is that a circular string can be represented by any of its linearizations, so these should all be regarded as equivalent. Furthermore, when the collection contains several conjugates of the same string, then this may be just an artifact caused by the data acquisition process.

This could be solved by replacing each circular string with its necklace representative, that is, the unique conjugate that is lexicographically minimal in the conjugacy class, before applying `make-string`; this representative is computable in linear time in the string length [51]. However, updates can change the lexicographic relationship of the rotations, and thus the necklace representative of the conjugacy class. Recomputing the necklace rotation of s after each update would add worst-case $\mathcal{O}(|s|)$ time to our running times, which is not acceptable. Computing the necklace rotation after an edit operation, or more in general, after any one of our update operations, is an interesting research question, which to the best of our knowledge has not yet been addressed.

7. Conclusion

We presented a new data structure, a forest of enhanced splay trees (FeST), to handle collections of dynamic strings. Our solution is much simpler than those offering the best theoretical results, while still having logarithmic amortized times for most update and query operations. We answer queries correctly whp, and updates are always correct.

To build our data structure, we employ an approach that differs from theoretical solutions: we use a splay tree for representing each string, enhancing it with additional annotations. The use of binary trees to represent dynamic strings is not new, but exploiting the simplicity of splay trees for attaching and detaching subtrees is. As our FeST is easy to understand, explain, and implement, we believe that it offers the opportunity of wide usability and can become a textbook implementation of dynamic strings. Further, we have found nontrivial—yet perfectly implementable—solutions to relevant queries, like computing the length ℓ of the longest common prefix of two suffixes in time $\mathcal{O}(\log n + \log^2 \ell)$ instead of the trivial $\mathcal{O}(\log^2 n)$. The simplicity of our solution enables new features, like the possibility of reversing a substring, or reverse-complementing it, to be easily implemented in logarithmic amortized time. Our data structure can be implemented within compact space—that is, proportional to the size of the represented data—and can be extended to handle circular strings and omega-extensions of strings—, features that competing solutions have not explored.

We have implemented our FeST data structure and all operations on linear strings presented in Sections 3 and 4. Our C++ implementation is freely available at <https://github.com/fmasillo/FeST>. The complete code consists of fewer than 1000 lines, evidencing the simplicity of our solution.

We believe that the FeST data structure is sufficiently simple but at the same time powerful enough to be an excellent candidate for a data structure course. It is simple enough that it can be given as a coding project for advanced undergraduate students. At the same time, it allows introducing several non-trivial concepts such as dynamicity, amortized analysis, tree balancing operations, and hashing. It is also highly efficient, getting close to theoretically optimal solutions, and can thus be used in practical settings, by programmers who wish to implement data structures for dynamic text collections.

CRedit authorship contribution statement

Zsuzsanna Lipták: Writing - review & editing, Writing - original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization; **Francesco Masillo:** Writing - review & editing, Writing - original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization; **Gonzalo Navarro:** Writing - review & editing, Writing - original draft, Visualization, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Funding acquisition, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

We thank the reviewers for their constructive comments, in particular for the suggestion to achieve compact space. Zs.L. partially funded by the Italian Ministry of University and Research (MUR) PRIN Project PINC, Pangenome INformatiCs: from Theory to Applications (Grant No. 2022YRB97K), and by the INdAM - GNCS Project CUP_E53C24001950001. G.N. was supported in part by Basal Funds FB0001 and AFB240001, as well as Fondecyt Grant 1–230755, ANID, Santiago, Chile.

References

- [1] H. Boehm, R.R. Atkinson, M.F. Plass, Ropes: an alternative to strings, *Softw. Prac. Exper.* 25 (12) (1995) 1315–1330. <https://doi.org/10.1002/spe.4380251203>
- [2] K. Menard, C. Seaton, B. Daloz, Specializing ropes for Ruby, in: E. Tilevich, H. Mössenböck (Eds.), *Proc. 15th International Conference on Managed Languages & Runtimes (ManLang)*, ACM, 2018, pp. 10:1–10:7.
- [3] R. Sundar, R.E. Tarjan, Unique Binary-Search-Tree representations and equality testing of sets and sequences, *SIAM J. Comput.* 23 (1) (1994) 24–44.
- [4] K. Mehlhorn, R. Sundar, C. Uhrig, Maintaining dynamic sequences under equality tests in polylogarithmic time, *Algorithmica* 17 (2) (1997) 183–198. <https://doi.org/10.1007/BF02522825>
- [5] S. Alstrup, G.S. Brodal, T. Rauhe, Pattern matching in dynamic texts, in: D.B. Shmoys (Ed.), *Proc. 11th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, ACM/SIAM, 2000, pp. 819–828. <http://dl.acm.org/citation.cfm?id=338219.338645>.
- [6] P. Gawrychowski, A. Karczmarz, T. Kociumaka, J. Lacki, P. Sankowski, Optimal dynamic strings, *CoRR* (2015). 1511.02612.
- [7] P. Gawrychowski, A. Karczmarz, T. Kociumaka, J. Lacki, P. Sankowski, Optimal dynamic strings, in: *Proc. 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, SIAM, 2018, pp. 1509–1528.
- [8] T. Nishimoto, T. I. S. Inenaga, H. Bannai, M. Takeda, Fully dynamic data structure for LCE queries in compressed space, in: P. Faliszewski, A. Muscholl, R. Niedermeier (Eds.), *Proc. 41st International Symposium on Mathematical Foundations of Computer Science (MFCS)*, 58 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2016, pp. 72:1–72:14. <https://doi.org/10.4230/LIPIcs.MFCS.2016.72>
- [9] T. Nishimoto, I. Tomohiro, S. Inenaga, H. Bannai, M. Takeda, Dynamic index and LZ factorization in compressed space, *Discrete Appl. Math.* 274 (2020) 116–129.
- [10] J. Cardinal, J. Iacono, Modular subset sum, dynamic strings, and zero-sum sets, in: *Proc. 4th Symposium on Simplicity in Algorithms (SOSA)*, SIAM, 2021, pp. 45–56.
- [11] J.R. Driscoll, N. Sarnak, D.D. Sleator, R.E. Tarjan, Making data structures persistent, in: J. Hartmanis (Ed.), *Proc. 18th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, 1986, pp. 109–121.
- [12] W. Rytter, Application of Lempel-Ziv factorization to the approximation of grammar-based compression, *Theor. Comput. Sci.* 302 (1–3) (2003) 211–222.
- [13] M. Charikar, E. Lehman, D. Liu, R. Panigrahy, M. Prabhakaran, A. Sahai, A. Shelat, The smallest grammar problem, *IEEE Trans. Inf. Theory* 51 (7) (2005) 2554–2576.
- [14] P. Gawrychowski, Pattern matching in Lempel-Ziv compressed strings: fast, simple, and deterministic, in: C. Demetrescu, M.M. Halldórsson (Eds.), *Proc. 19th Annual European Symposium on Algorithms (ESA)*, 6942 of *Lecture Notes in Computer Science*, Springer, 2011, pp. 421–432. https://doi.org/10.1007/978-3-642-23719-5_36
- [15] D. Kempa, B. Langmead, Fast and space-efficient construction of AVL grammars from the LZ77 parsing, in: *Proc. 29th Annual European Symposium on Algorithms (ESA)*, 204, 2021, pp. 56:1–56:14.
- [16] D.E. Knuth, *The Art of Computer Programming, volume 3: Sorting and Searching*, Addison-Wesley, Reading, MA, USA, 2nd edition, Reading, MA, USA, 1998.
- [17] A. Amir, P. Charalampopoulos, C.S. Iliopoulos, S.P. Pissis, J. Radoszewski, Longest common factor after one edit operation, in: *Proc. 24th International Symposium on String Processing and Information Retrieval (SPIRE)*, 10508 of *Lecture Notes in Computer Science*, Springer, 2017, pp. 14–26.
- [18] A. Amir, P. Charalampopoulos, S.P. Pissis, J. Radoszewski, Dynamic and internal longest common substring, *Algorithmica* 82 (12) (2020) 3707–3743.
- [19] P. Charalampopoulos, T. Kociumaka, S. Mozes, Dynamic string alignment, in: *Proc. 31st Annual Symposium on Combinatorial Pattern Matching (CPM)*, 161 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 9:1–9:13.
- [20] T. Kociumaka, A. Mukherjee, B. Saha, Approximating edit distance in the fully dynamic model, in: *Proc. 64th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE, 2023, pp. 1628–1638.
- [21] M. Funakoshi, Y. Nakashima, S. Inenaga, H. Bannai, M. Takeda, Longest substring palindrome after edit, in: *Proc. 29th Annual Symposium on Combinatorial Pattern Matching (CPM)*, 105 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 12:1–12:14.
- [22] A. Amir, I. Boneh, P. Charalampopoulos, E. Kondratovsky, Repetition detection in a dynamic string, in: M.A. Bender, O. Svensson, G. Herman (Eds.), *Proc. 27th Annual European Symposium on Algorithms (ESA)*, 144 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019, pp. 5:1–5:18.
- [23] Y. Urabe, Y. Nakashima, S. Inenaga, H. Bannai, M. Takeda, Longest Lyndon substring after edit, in: G. Navarro, D. Sankoff, B. Zhu (Eds.), *Proc. 29th Annual Symposium on Combinatorial Pattern Matching (CPM)*, 105 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 19:1–19:10.
- [24] R. Clifford, A. Grønlund, K.G. Larsen, T. Starikovskaya, Upper and lower bounds for dynamic data structures on strings, in: *Proc. 35th Symposium on Theoretical Aspects of Computer Science (STACS)*, 96 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 22:1–22:14.
- [25] D. Kempa, T. Kociumaka, Dynamic suffix array with polylogarithmic queries and updates, in: *Proc. 54th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, ACM, 2022, pp. 1657–1670.
- [26] D. Kempa, T. Kociumaka, Dynamic suffix array with polylogarithmic queries and updates, *CoRR* (2022). 2201.01285
- [27] D.D. Sleator, R.E. Tarjan, Self-adjusting binary search trees, *J. ACM* 32 (3) (1985) 652–686.
- [28] R.M. Karp, M.O. Rabin, Efficient randomized pattern-matching algorithms, *IBM J. Res. Dev.* 31 (2) (1987) 249–260.
- [29] G. Navarro, N. Prezza, Universal compressed text indexing, *Theor. Comput. Sci.* 762 (2019) 41–50.
- [30] W. Szpankowski, A generalized suffix tree and its (un)expected asymptotic behaviors, *SIAM J. Comput.* 22 (6) (1993) 1176–1198.
- [31] G.A. Watterson, W.J. Ewens, T.E. Hall, A. Morgan, The chromosome inversion problem, *J. Theor. Biol.* 99 (1982) 1–7.
- [32] V. Bafna, P.A. Pevzner, Genome rearrangements and sorting by reversals, in: *Proc. 34th Annual Symposium on Foundations of Computer Science (FOCS)*, IEEE Computer Society, 1993, pp. 148–157.
- [33] A. Caprara, Sorting by reversals is difficult, in: *Proc. 1st Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, ACM, 1997, pp. 75–83.
- [34] A. Caprara, R. Rizzi, Improved approximation for breakpoint graph decomposition and sorting by reversals, *J. Comb. Optim.* 6 (2) (2002) 157–182.
- [35] A.R. Oliveira, U. Dias, Z. Dias, On the sorting by reversals and transpositions problem, *J. Universal Comput. Sci.* 23 (9) (2017) 868–906.
- [36] G. Cerbai, L.S. Ferrari, Permutation patterns in genome rearrangement problems: The reversal model, *Discrete Appl. Math.* 279 (2020) 34–48.
- [37] G. Fertin, A. Labarre, I. Rusu, E. Tannier, S. Vialette, *Combinatorics of Genome Rearrangements*, Computational Molecular Biology Series, MIT Press, 2009.
- [38] S. Hannenhalli, P.A. Pevzner, Transforming cabbage into turnip: Polynomial algorithm for sorting signed permutations by reversals, in: *Proc. 27th Annual ACM Symposium on Theory of Computing (STOC)*, ACM, 1995, pp. 178–189.
- [39] Y. Han, Improving the efficiency of sorting by reversals, in: *Proc. International Conference on Bioinformatics & Computational Biology (BIOCOMP)*, 2006, pp. 406–409.
- [40] G. Navarro, K. Sadakane, Fully-functional static and dynamic succinct trees, *ACM Trans. Algorithms* 10 (3) (2014) 16:1–16:39.
- [41] L.A.K. Ayad, S.P. Pissis, MARS: Improving multiple circular sequence alignment using refined sequences, *BMC Genom.* 18 (1) (2017) 1–10.
- [42] P. Charalampopoulos, T. Kociumaka, J. Radoszewski, S.P. Pissis, W. Rytter, T. Walen, W. Zuba, Approximate circular pattern matching, in: *Proc. 30th Annual European Symposium on Algorithms (ESA)*, 244 of *LIPIcs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022, pp. 35:1–35:19.

- [43] R. Grossi, C.S. Iliopoulos, J. Jansson, Z. Lim, W. Sung, W. Zuba, Finding the cyclic covers of a string, in: Proc. 17th International Conference and Workshops on Algorithms and Computation (WALCOM), 13973 of *Lecture Notes in Computer Science*, Springer, 2023, pp. 139–150.
- [44] R. Grossi, C.S. Iliopoulos, R. Mercas, N. Pisanti, S.P. Pissis, A. Retha, F. Vayani, Circular sequence comparison: algorithms and applications, *Algorithms Mol. Biol.* 11 (1) (2016) 1–14.
- [45] C.S. Iliopoulos, T. Kociumaka, J. Radoszewski, W. Rytter, T. Walen, W. Zuba, Linear-time computation of cyclic roots and cyclic covers of a string, in: Proc. 34th Annual Symposium on Combinatorial Pattern Matching (CPM), 259 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023, pp. 15:1–15:15.
- [46] C. Boucher, D. Cenzato, Zs. Lipták, M. Rossi, M. Sciortino, Computing the original eBWT faster, simpler, and with less memory, in: T. Lecroq, H. Touzet (Eds.), Proc. 28th International Symposium on String Processing and Information Retrieval (SPIRE), 12944 of *Lecture Notes in Computer Science*, Springer, 2021, pp. 129–142.
- [47] C. Boucher, D. Cenzato, Zs. Lipták, M. Rossi, M. Sciortino, r -indexing the eBWT, *Inform. Comput.* 298 (2024) 105155. <https://doi.org/10.1016/j.ic.2024.105155>
- [48] S. Mantaci, A. Restivo, G. Rosone, M. Sciortino, An extension of the Burrows-Wheeler Transform, *Theor. Comput. Sci.* 387 (3) (2007) 298–312.
- [49] Zs. Lipták, F. Masillo, G. Navarro, A textbook solution for dynamic strings, in: T.M. Chan, J. Fischer, J. Iacono, G. Herman (Eds.), Proc. 32nd Annual European Symposium on Algorithms (ESA), 308 of *LIPICs*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2024, pp. 86:1–86:16. <https://doi.org/10.4230/LIPICs.ESA.2024.86>
- [50] R. Cole, On the dynamic finger conjecture for splay trees. Part II: The proof, *SIAM J. Comput.* 30 (1) (2000) 44–85. <https://doi.org/10.1137/S009753979732699X>
- [51] M. Lothaire, *Applied Combinatorics on Words*, Cambridge University Press, 2005.
- [52] D. Stehlé, P. Zimmermann, A binary recursive gcd algorithm, in: D.A. Buell (Ed.), Proc. 6th International Symposium on Algorithmic Number Theory (ANTS), 3076 of *Lecture Notes in Computer Science*, Springer, 2004, pp. 411–425. https://doi.org/10.1007/978-3-540-24847-7_31