# Compressed Indexing with Signature Grammars 

Anders Roy Christiansen and Mikko Berggren Ettienne

The Technical University of Denmark


#### Abstract

The compressed indexing problem is to preprocess a string $S$ of length $n$ into a compressed representation that supports pattern matching queries. That is, given a string $P$ of length $m$ report all occurrences of $P$ in $S$. We present a data structure that supports pattern matching queries in $O\left(m+\operatorname{occ}\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ time using $O(z \lg (n / z))$ space where $z$ is the size of the LZ77 parse of $S$ and $\epsilon>0$ is an arbitrarily small constant, when the alphabet is small or $z=O\left(n^{1-\delta}\right)$ for any constant $\delta>0$. We also present two data structures for the general case; one where the space is increased by $O(z \lg \lg z)$, and one where the query time changes from worst-case to expected. These results improve the previously best known solutions. Notably, this is the first data structure that decides if $P$ occurs in $S$ in $O(m)$ time using $O(z \lg (n / z))$ space. Our results are mainly obtained by a novel combination of a randomized grammar construction algorithm with well known techniques relating pattern matching to 2D-range reporting.


## 1 Introduction

Given a string $S$ and a pattern $P$, the core problem of pattern matching is to report all locations where $P$ occurs in $S$. Pattern matching problems can be divided into two: the algorithmic problem where the text and the pattern are given at the same time, and the data structure problem where one is allowed to preprocess the text (pattern) before a query pattern (text) is given. Many problems within both these categories are well-studied in the history of stringology, and optimal solutions to many variants have been found.

In the last decades, researchers have shown an increasing interest in the compressed version of this problem, where the space used by the index is related to the size of some compressed representation of $S$ instead of the length of $S$. This could be measures such as the size of the LZ77-parse of $S$, the smallest grammar representing $S$, the number of runs in the BWT of $S$, etc. see e.g. [10|3|8|17|16|3]. This problem is highly relevant as the amount of highlyrepetitive data increases rapidly, and thus it is possible to handle greater amounts of data by compressing it. The increase in such data is due to things like DNA sequencing, version control repositories, etc.

In this paper we consider what we call the compressed indexing problem, which is to preprocess a string $S$ of length $n$ into a compressed representation that supports fast pattern matching queries. That is, given a string $P$ of length $m$, report all occ occurrences of substrings in $S$ that match $P$.

Table 1 gives an overview of the results on this problem.

Table 1. Selection of previous results and our new results on compressed indexing. The variables are the text size $n$, the LZ77-parse size $z$, the pattern length $m$, occ is the number of occurrences and $\sigma$ is the size of the alphabet. (The time complexity marked by $\dagger$ is expected whereas all others are worst-case)

| Index | Space | Locate time | $\sigma$ |
| :---: | :---: | :---: | :---: |
| Gagie et al. [9] | $O(z \lg (n / z))$ | $O(m \lg m+$ occ $\lg \lg n)$ | $O(1)$ |
| Nishimoto et al. 17 | $O\left(z \lg n \lg ^{*} n\right)$ | $O(m \lg \lg n \lg \lg z$ $\lg z \lg m \lg n\left(\lg ^{*} n\right)^{2} \quad+$ occ $\lg n$ ) | $n^{O(1)}$ |
| Bille et al. 3] | $O\left(z\left(\lg (n / z)+\lg ^{\epsilon} z\right)\right)$ | $O\left(m+\operatorname{occ}\left(\lg ^{\epsilon} n+\lg \lg n\right)\right)$ | $n^{O(1)}$ |
| Bille et al. 3] | $O(z \lg (n / z) \lg \lg z)$ | $O(m+$ occ $\lg \lg n)$ | $O(1)$ |
| Bille et al. 3] | $O(z \lg (n / z))$ | $\begin{aligned} & O\left(m\left(1+\frac{\lg ^{\epsilon} z}{\lg (n / z)}\right)+\right. \\ & \left.O \operatorname{Occ}\left(\lg ^{\epsilon} n+\lg \lg n\right)\right) \end{aligned}+$ | $O(1)$ |
| Theorem 1 | $O(z \lg (n / z))$ | $O\left(m+\operatorname{occ}\left(\lg ^{\epsilon} z+\lg \lg n\right)\right)$ | $O(1)$ |
| Theorem 2 (1) | $O(z(\lg (n / z)+\lg \lg z))$ | $O\left(m+\operatorname{occ}\left(\lg ^{\epsilon} z+\lg \lg n\right)\right)$ | $n^{O(1)}$ |
| Theorem 2 (2) | $O(z(\lg (n / z))$ | $\mid O\left(m+\operatorname{occ}\left(\lg ^{\epsilon} z+\lg \lg n\right)\right)^{\dagger}$ | $n^{O(1)}$ |

### 1.1 Our Results

In this paper we improve previous solutions that are bounded by the size of the LZ77-parse. For constant-sized alphabets we obtain the following result:

Theorem 1. Given a string $S$ of length $n$ from a constant-sized alphabet with an LZ77 parse of length $z$, we can build a compressed-index supporting pattern matching queries in $O\left(m+\operatorname{occ}\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ time using $O(z \lg (n / z))$ space.

In particular, we are the first to obtain optimal search time using only $O(z \lg (n / z))$ space. For general alphabets we obtain the following:

Theorem 2. Given a string $S$ of length $n$ from an integer alphabet polynomially bounded by $n$ with an LZ77-parse of length $z$, we can build a compressed-index supporting pattern matching queries in:
(1) $O\left(m+\operatorname{occ}\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ time using $O(z(\lg (n / z)+\lg \lg z))$ space.
(2) $O\left(m+\operatorname{occ}\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ expected time using $O(z \lg (n / z))$ space.
(3) $O\left(m+\lg ^{\epsilon} z+\operatorname{occ}\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ time using $O(z \lg (n / z))$ space.

Note $\lg \lg z=O(\lg (n / z))$ when either the alphabet size is $O\left(2^{\lg ^{\epsilon} n}\right)$ or $z=o\left(\frac{n}{\lg ^{\epsilon^{\epsilon} n}}\right)$ where $\epsilon$ and $\epsilon^{\prime}$ are arbitrarily small positive constants. Theorem 1 follows directly from Theorem 2 (1) given these observations. Theorem 2 is a consequence of Lemma 9, 11, 12 and 13

### 1.2 Technical Overview

Our main new contribution is based on a new grammar construction. In 15 Melhorn et al. presented a way to maintain dynamic sequences subject to equality testing using a technique called signatures. They presented two signature
construction techniques. One is randomized and leads to complexities that hold in expectation. The other is based on a deterministic coin-tossing technique of Cole and Vishkin [5] and leads to worst-case running times but incurs an iterated logarithmic overhead compared to the randomized solution. This technique has also resembles the string labeling techniques found e.g. in [19]. To the best of our knowledge, we are the first to consider grammar compression based on the randomized solution from [15]. Despite it being randomized we show how to obtain worst-case query bounds for text indexing using this technique.

The main idea in this grammar construction is that similar substrings will be parsed almost identically. This property also holds true for the deterministic construction technique which has been used to solve dynamic string problems with and without compression, see e.g. [1711]. In [12] Jeż devices a different grammar construction algorithm with similar properties to solve the algorithmic pattern matching problem on grammar compressed strings which has later been used for both static and dynamic string problems, see [20 11]

Our primary solution has an $\lg ^{\epsilon} z$ term in the query time which is problematic for short query patterns. To handle this, we show different solutions for handling short query patterns. These are based on the techniques from LZ77-based indexing combined with extra data structures to speed up the queries.

## 2 Preliminaries

We assume a standard unit-cost RAM model with word size $\Theta(\lg n)$ and that the input is from an integer alphabet $\Sigma=\left\{1,2, \ldots, n^{O(1)}\right\}$. We measure space complexity in terms of machine words unless explicitly stated otherwise. A string $S$ of length $n=|S|$ is a sequence of $n$ symbols $S[1] \ldots S[n]$ drawn from an alphabet $\Sigma$. The sequence $S[i, j]$ is the substring of $S$ given by $S[i] \ldots S[j]$ and strings can be concatenated, i.e. $S=S[1, k] S[k+1, n]$. The empty string is denoted $\epsilon$ and $S[i, i]=S[i]$ while $S[i, j]=\epsilon$ if $j<i, S[i, j]=S[1, j]$ if $i<1$ and $S[i, n]$ if $j>n$. The reverse of $S$ denoted $\operatorname{rev}(s)$ is the string $S[n] S[n-1] \ldots S[1]$. A run in a string $S$ is a substring $S[i, j]$ with identical letters, i.e. $S[k]=S[k+1]$ for $k=i, \ldots, j-1$. Let $S[i, j]$ be a run in $S$ then it is a maximal run if it cannot be extended, i.e. $S[i-1] \neq S[i]$ and $S[j] \neq S[j+1]$. If there are no runs in $S$ we say that $S$ is run-free and it follows that $S[i] \neq S[i+1]$ for $1 \leq i<n$. Denote by $[u]$ the set of integers $\{1,2, \ldots, u\}$.

Let $X \subseteq[u]^{2}$ be a set of points in a 2-dimensional grid. The $2 D$-orthogonal range reporting problem is to compactly represent $Z$ while supporting range reporting queries, that is, given a rectangle $R=\left[a_{1}, b_{1}\right] \times\left[a_{2}, b_{2}\right]$ report all points in the set $R \cap X$. We use the following:

Lemma 1 (Chan et al. [4]). For any set of $n$ points in $[u] \times[u]$ and constant $\epsilon>0$, we can solve 2D-orthogonal range reporting with $O(n \lg n)$ expected preprocessing time using:
$i O(n)$ space and $(1+k) \cdot O\left(\lg ^{\epsilon} n \lg \lg u\right)$ query time
ii $O(n \lg \lg n)$ space and $(1+k) \cdot O(\lg \lg u)$ query time
where $k$ is the number of occurrences inside the rectangle.
A Karp-Rabin fingerprinting function [14 is a randomized hash function for strings. Given a string $S$ of length $n$ and a fingerprinting function $\phi$ we can in $O(n)$ time and space compute and store $O(n)$ fingerprints such that the fingerprint of any substring of $S$ can be computed in constant time. Identical strings have identical fingerprints. The fingerprints of two strings $S$ and $S^{\prime}$ collide when $S \neq S^{\prime}$ and $\phi(S)=\phi\left(S^{\prime}\right)$. A fingerprinting function is collision-free for a set of strings when there are no collisions between the fingerprints of any two strings in the set. We can find collision-free fingerprinting function for a set of strings with total length $n$ in $O(n)$ expected time [18].

Let $D$ be a lexicographically sorted set of $k$ strings. The weak prefix search problem is to compactly represent $D$ while supporting weak prefix queries, that is, given a query string $P$ of length $m$ report the rank of the lexicographically smallest and largest strings in $D$ of which $P$ is a prefix. If no such strings exist, the answer can be arbitrary.

Lemma 2 (Belazzougui et al. [2], appendix H.3). Given a set $D$ of $k$ strings with average length $l$, from an alphabet of size $\sigma$, we can build a data structure using $O(k(\lg l+\lg \lg \sigma))$ bits of space supporting weak prefix search for a pattern $P$ of length $m$ in $O(m \lg \sigma / w+\lg m)$ time where $w$ is the word size.

We will refer to the data structure of Lemma 2 as a $z$-fast trie following the notation from [2]. The $m$ term in the time complexity is due to a linear time preprocessing of the pattern and is not part of the actual search. Therefore it is simple to do weak prefix search for any length $l$ substring of $P$ in $O(\lg l)$ time after preprocessing $P$ once in $O(m)$ time.

The LZ77-parse [21] of a string $S$ of length $n$ is a string $\mathcal{Z}$ of the form $\left(s_{1}, l_{1}, \alpha_{1}\right) \ldots\left(s_{z}, l_{z}, \alpha_{z}\right) \in([n],[n], \Sigma)^{z}$. We define $u_{1}=1, u_{i}=u_{i-1}+l_{i-1}+1$ for $i>1$. For $\mathcal{Z}$ to be a valid parse, we require $l_{1}=0, s_{i}<u_{i}, S\left[u_{i}, u_{i}+l_{i}-1\right]=$ $S\left[s_{i}, s_{i}+l_{i}-1\right]$, and $S\left[u_{i}+l_{i}\right]=\alpha_{i}$ for $i \in[z]$. This guarantees $\mathcal{Z}$ represents $S$ and $S$ is uniquely defined in terms of $\mathcal{Z}$. The substring $S\left[u_{i}, u_{i}+l_{i}\right]$ is called the $i^{t h}$ phrase of the parse and $S\left[s_{i}, s_{i}+l_{i}-1\right]$ is its source. A minimal LZ77-parse of $S$ can be found greedily in $O(n)$ time and stored in $O(z)$ space 21. We call the positions $u_{1}+l_{1}, \ldots, u_{z}+l_{z}$ the borders of $S$.

## 3 Signature Grammars

We consider a hierarchical representation of strings given by Melhorn et al. 15] with some slight modifications. Let $S$ be a run-free string of length $n$ from an integer alphabet $\Sigma$ and let $\pi$ be a uniformly random permutation of $\Sigma$. Define a position $S[i]$ as a local minimum of $S$ if $1<i<n$ and $\pi(S[i])<\pi(S[i-1])$ and $\pi(S[i])<\pi(S[i+1])$. In the block decomposition of $S$, a block starts at position 1 and at every local minimum in $S$ and ends just before the next block begins (the last block ends at position $n$ ). The block decomposition of a string $S$ can be used to construct the signature tree of $S$ denoted $\operatorname{sig}(S)$ which is an ordered labeled tree with several useful properties.

Lemma 3. Let $S$ be a run-free string $S$ of length $n$ from an alphabet $\Sigma$ and let $\pi$ be a uniformly random permutation of $\Sigma$ such that $\pi(c)$ is the rank of the symbol $c \in \Sigma$ in this permutation. Then the expected length between two local minima in the sequence $\pi(S[1]), \pi(S[2]), \ldots, \pi(S[n])$ is at most 3 and the longest gap is $O(\lg n)$ in expectation.

Proof. First we show the expected length between two local minima is at most 3. Look at a position $1 \leq i \leq n$ in the sequence $\pi(S[1]), \pi(S[2]), \ldots, \pi(S[n])$. To determine if $\pi(S[i])$ is a local minimum, we only need to consider the two neighbouring elements $\pi(S[i-1])$ and $\pi(S[i+1])$ thus let us consider the triple $(\pi(S[i-1]), \pi(S[i]), \pi(S[i+1]))$. We need to consider the following cases. First assume $S[i-1] \neq S[i] \neq S[i+1]$. There exist $3!=6$ permutations of a triple with unique elements and in two of these the minimum element is in the middle. Since $\pi$ is a uniformly random permutation of $\Sigma$ all 6 permutations are equally likely, and thus there is $1 / 3$ chance that the element at position $i$ is a local minimum. Now instead assume $S[i-1]=S[i+1] \neq S[i]$ in which case there is $1 / 2$ chance that the middle element is the smallest. Finally, in the case where $i=1$ or $i=n$ there is also $1 / 2$ chance. As $S$ is run-free, these cases cover all possible cases. Thus there is at least $1 / 3$ chance that any position $i$ is a local minimum independently of $S$. Thus the expected number of local minima in the sequence is therefore at least $n / 3$ and the expected distance between any two local minima is at most 3 .

The expected longest distance between two local minima of $O(\lg n)$ was shown in [15].

### 3.1 Signature Grammar Construction

We now give the construction algorithm for the signature tree $\operatorname{sig}(S)$. Consider an ordered forest $F$ of trees. Initially, $F$ consists of $n$ trees where the $i^{\text {th }}$ tree is a single node with label $S[i]$. Let the label of a tree $t$ denoted $l(t)$ be the label of its root node. Let $l(F)$ denote the string that is given by the in-order concatenation of the labels of the trees in $F$. The construction of $\operatorname{sig}(S)$ proceeds as follows:

1. Let $t_{i}, \ldots, t_{j}$ be a maximal subrange of consecutive trees of $F$ with identical labels, i.e. $l\left(t_{i}\right)=\ldots=l\left(t_{j}\right)$. Replace each such subrange in $F$ by a new tree having as root a new node $v$ with children $t_{i}, \ldots, t_{j}$ and a label that identifies the number of children and their label. We call this kind of node a run node. Now $l(F)$ is run-free.
2. Consider the block decomposition of $l(F)$. Let $t_{i}, \ldots, t_{j}$ be consecutive trees in $F$ such that their labels form a block in $l(F)$. Replace all identical blocks $t_{i}, \ldots, t_{j}$ by a new tree having as root a new node with children $t_{i}, \ldots, t_{j}$ and a unique label. We call this kind of node a run-free node.
3. Repeat step 1 and 2 until $F$ contains a single tree, we call this tree $\operatorname{sig}(S)$.

In each iteration the size of $F$ decreases by at least a factor of two and each iteration takes $O(|F|)$ time, thus it can be constructed in $O(n)$ time.

Consider the directed acyclic graph (DAG) of the tree $\operatorname{sig}(S)$ where all identical subtrees are merged. Note we can store run nodes in $O(1)$ space since all out-going edges are pointing to the same node, so we store the number of edges along with a single edge instead of explicitly storing each of them. For run-free nodes we use space proportional to their out-degrees. We call this the signature DAG of $S$ denoted $\operatorname{dag}(S)$. There is a one-to-one correspondence between this DAG and an acyclic run-length grammar producing $S$ where each node corresponds to a production and each leaf to a terminal.

### 3.2 Properties of the Signature Grammar

We now show some properties of $\operatorname{sig}(S)$ and $\operatorname{dag}(S)$ that we will need later. Let $\operatorname{str}(v)$ denote the substring of $S$ given by the labels of the leaves of the subtree of $\operatorname{sig}(S)$ induced by the node $v$ in left to right order.

Lemma 4. Let $v$ be a node in the signature tree for a string $S$ of length $n$. If $v$ has height $h$ then $|\operatorname{str}(v)|$ is at least $2^{h}$ and thus sig(S) (and $\operatorname{dag}(S)$ ) has height $O(\lg n)$.

Proof. This follows directly from the out-degree of all nodes being at least 2 .
Denote by $T(i, j)$ the set of nodes in $\operatorname{sig}(S)$ that are ancestors of the $i^{\text {th }}$ through $j^{\text {th }}$ leaf of $\operatorname{sig}(S)$. These nodes form a sequence of adjacent nodes at every level of $\operatorname{sig}(S)$ and we call them relevant nodes for the substring $S[i, j]$.

Lemma 5. $T(i, j)$ and $T\left(i^{\prime}, j^{\prime}\right)$ have identical nodes except at most the two first and two last nodes on each level whenever $S[i, j]=S\left[i^{\prime}, j^{\prime}\right]$.

Proof. Trivially, the leaves of $T(i, j)$ and $T\left(i^{\prime}, j^{\prime}\right)$ are identical if $S[i, j]=S\left[i^{\prime}, j^{\prime}\right]$. Now we show it is true for nodes on level $l$ assuming it is true for nodes on level $l-1$. We only consider the left part of each level as the argument for the right part is (almost) symmetric. Let $v_{1}, v_{2}, v_{3}, \ldots$ be the nodes on level $l-1$ in $T(i, j)$ and $u_{1}, u_{2}, u_{3}, \ldots$ the nodes on level $l-1$ in $T\left(i^{\prime}, j^{\prime}\right)$ in left to right order. From the assumption, we have $v_{a}, v_{a+1}, \ldots$ are identical with $u_{b}, u_{b+1}, \ldots$ for some $1 \leq a, b \leq 3$. When constructing the $l^{t h}$ level of $\operatorname{sig}(S)$, these nodes are divided into blocks. Let $v_{a+k}$ be the first block that starts after $v_{a}$ then by the block decomposition, the first block after $u_{b}$ starts at $u_{b+k}$. The nodes $v_{1}, \ldots, v_{a+k}$ are spanned by at most two blocks and similarly for $u_{1}, \ldots, u_{b+k}$. These blocks become the first one or two nodes on level $l$ in $T(i, j)$ and $T\left(i^{\prime}, j^{\prime}\right)$ respectively. The block starting at $v_{a+k}$ is identical to the block starting at $u_{b+k}$ and the same holds for the following blocks. These blocks result in identical nodes on level $l$. Thus, if we ignore the at most two first (and last) nodes on level $l$ the remaining nodes are identical.

We call nodes of $T(i, j)$ consistent in respect to $T(i, j)$ if they are guaranteed to be in any other $T\left(i^{\prime}, j^{\prime}\right)$ where $S[i, j]=S\left[i^{\prime}, j^{\prime}\right]$. We denote the remaining nodes of $T(i, j)$ as inconsistent. From the above lemma, it follows at most the left-most and right-most two nodes on each level of $T(i, j)$ can be inconsistent.

Lemma 6. The expected size of the signature $D A G \operatorname{dag}(S)$ is $O(z \lg (n / z))$.
Proof. We first bound the number of unique nodes in $\operatorname{sig}(S)$ in terms of the LZ77-parse of $S$ which has size $z$. Consider the decomposition of $S$ into the $2 z$ substrings $S\left[u_{1}, u_{1}+l_{1}\right], S\left[u_{1}+l_{1}+1\right], \ldots, S\left[u_{z}, u_{z}+l_{z}\right], S\left[u_{z}+l_{z}+1\right]$ given by the phrases and borders of the LZ77-parse of $S$ and the corresponding sets of relevant nodes $R=\left\{T\left(u_{1}, u_{1}+l_{1}\right), T\left(u_{1}+l_{1}+1, u_{1}+l_{1}+1\right), \ldots\right\}$. Clearly, the union of these sets are all the nodes of $\operatorname{sig}(S)$. Since identical nodes are represented only once in $\operatorname{dag}(S)$ we need only count one of their occurrences in $\operatorname{sig}(S)$. We first count the nodes at levels lower than $\lg (n / z)$. A set $T(i, i)$ of nodes relevant to a substring of length one has no more than $O(\lg (n / z))$ such nodes. By Lemma 5 only $O(\lg (n / z))$ of the relevant nodes for a phrase are not guaranteed to also appear in the relevant nodes of its source. Thus we count a total of $O(z \lg (n / z))$ nodes for the $O(z)$ sets of relevant nodes. Consider the leftmost appearance of a node appearing one or more times in $\operatorname{sig}(S)$. By definition, and because every node of $\operatorname{sig}(S)$ is in at least one relevant set, it must already be counted towards one of the sets. Thus there are $O(z \lg (n / z))$ unique vertices in $\operatorname{sig}(S)$ at levels lower than $\lg (n / z)$. Now for the remaining at most $\lg (z)$ levels, there are no more than $O(z)$ nodes because the out-degree of every node is at least two. Thus we have proved that there are $O(z \lg (n / z))$ unique nodes in $\operatorname{sig}(S)$. By Lemma 3 the average block size and thus the expected out-degree of a node is $O(1)$. It follows that the expected number of edges and the expected size of $\operatorname{dag}(S)$ is $O(z \lg (n / z))$.
Lemma 7. A signature grammar of $S$ using $O(z \lg (n / z)$ ) (worst case) space can be constructed in $O(n)$ expected time.

Proof. Construct a signature grammar for $S$ using the signature grammar construction algorithm. If the average out-degree of the run-free nodes in $\operatorname{dag}(S)$ is more than some constant greater than 3 then try again. In expectation it only takes a constant number of retries before this is not the case.

Lemma 8. Given a node $v \in \operatorname{dag}(S)$, the child that produces the character at position $i$ in $\operatorname{str}(v)$ can be found in $O(1)$ time.

Proof. First assume $v$ is a run-free node. If we store $|\operatorname{str}(u)|$ for each child $u$ of $v$ in order, the correct child corresponding to position $i$ can simply be found by iterating over these. However, this may take $O(\log n)$ time since this is the maximum out-degree of a node in $\operatorname{dag}(S)$. This can be improved to $O(\log \log n)$ by doing a binary search, but instead we use a Fusion Tree from [7] that allows us to do this in $O(1)$ time since we have at most $O(\log n)$ elements. This does not increase the space usage. If $v$ is a run node then it is easy to calculate the right child by a single division.

## 4 Long Patterns

In this section we present how to use the signature grammar to construct a compressed index that we will use for patterns of length $\Omega\left(\lg ^{\epsilon} z\right)$ for constant $\epsilon>0$. We obtain the following lemma:

Lemma 9. Given a string $S$ of length $n$ with an LZ77-parse of length $z$ we can build a compressed index supporting pattern matching queries in $O(m+(1+$ occ) $\left.\lg ^{\epsilon} z\right)$ time using $O(z \lg (n / z))$ space for any constant $\epsilon>0$.

### 4.1 Data Structure

Consider a vertex $v$ with children $u_{1}, \ldots u_{k}$ in $\operatorname{dag}(S)$. Let $\operatorname{pre}(v, i)$ denote the prefix of $\operatorname{str}(v)$ given by concatenating the strings represented by the first $i$ children of $v$ and let $\operatorname{suf}(v, i)$ be the suffix of $\operatorname{str}(v)$ given by concatenating the strings represented by the last $k-i$ children of $x$.

The data structure is composed of two z-fast tries (see Lemma 2) $T_{1}$ and $T_{2}$ and a 2D-range reporting data structure $R$.

For every non-leaf node $v \in \operatorname{dag}(S)$ we store the following. Let $k$ be the number of children of $v$ if $v$ is a run-free node otherwise let $k=2$ :

- The reverse of the strings $\operatorname{pre}(v, i)$ for $i \in[k-1]$ in the z-fast trie $T_{1}$.
- The strings $\operatorname{suf}(v, i)$ for $i \in[k-1]$ in the z-fast trie $T_{2}$.
- The points $(a, b)$ where $a$ is the rank of the reverse of $\operatorname{pre}(v, i)$ in $T_{1}$ and $b$ is the rank of $\operatorname{suf}(v, i)$ in $T_{2}$ for $i \in[k-1]$ are stored in $R$. A point stores the vertex $v \in \operatorname{dag}(S)$ and the length of $\operatorname{pre}(v, i)$ as auxiliary information.

There are $O(z \lg (n / z))$ vertices in $\operatorname{dag}(S)$ thus $T_{1}$ and $T_{2}$ take no more than $O(z \lg (n / z))$ words of space using Lemma 2. There $O(z \lg (n / z))$ points in $R$ which takes $O(z \lg (n / z))$ space using Lemma (i) thus the total space in words is $O(z \lg (n / z))$.

### 4.2 Searching

Assume in the following that there are no fingerprint collisions. Compute all the prefix fingerprints of $P \phi(P[1]), \phi(P[1,2]), \ldots, \phi(P[1, m])$. Consider the signature tree $\operatorname{sig}(P)$ for $P$. Let $l_{i}^{k}$ denote the $k$ 'th left-most vertex on level $i$ in $\operatorname{sig}(P)$ and let $j$ be the last level. Let $P_{L}=\left\{\left|\operatorname{str}\left(l_{1}^{1}\right)\right|,\left|\operatorname{str}\left(l_{1}^{1}\right)\right|+\left|\operatorname{str}\left(l_{1}^{2}\right)\right|,\left|\operatorname{str}\left(l_{2}^{1}\right)\right|,\left|\operatorname{str}\left(l_{2}^{1}\right)\right|+\right.$ $\left.\left|\operatorname{str}\left(l_{2}^{2}\right)\right|, \ldots,\left|\operatorname{str}\left(l_{j}^{1}\right)\right|,\left|\operatorname{str}\left(l_{j}^{1}\right)\right|+\left|\operatorname{str}\left(l_{j}^{2}\right)\right|\right\}$. Symmetrically, let $r_{i}^{k}$ denote the $k$ 'th right-most vertex on level $i$ in $\operatorname{sig}(P)$ and let $P_{R}=\left\{m-\left|\operatorname{str}\left(r_{1}^{1}\right)\right|, m-\left|\operatorname{str}\left(r_{1}^{1}\right)\right|-\right.$ $\left|\operatorname{str}\left(r_{1}^{2}\right)\right|, m-\left|\operatorname{str}\left(r_{2}^{1}\right)\right|, m-\left|\operatorname{str}\left(r_{2}^{1}\right)\right|-\left|\operatorname{str}\left(r_{2}^{2}\right)\right|, \ldots, m-\left|\operatorname{str}\left(r_{j}^{1}\right)\right|, m-\left|\operatorname{str}\left(r_{j}^{1}\right)\right|-$ $\left.\left|\operatorname{str}\left(r_{j}^{2}\right)\right|\right\}$. Let $P_{S}=P_{L} \cup P_{R}$.

For $p \in P_{S}$ search for the reverse of $P[1, p]$ in $T_{1}$ and for $P[p+1, m]$ in $T_{2}$ using the precomputed fingerprints. Let $[a, b]$ and $[c, d]$ be the respective ranges returned by the search. Do a range reporting query for the (possibly empty) range $[a, b] \times[c, d]$ in $R$. Each point in the range identifies a node $v$ and a position $i$ such that $P$ occurs at position $i$ in the string $\operatorname{str}(v)$. If $v$ is a run node, there is furthermore an occurrence of $P$ in $\operatorname{str}(v)$ for all positions $i+k \cdot|\operatorname{str}(\operatorname{child}(v))|$ where $k=1, \ldots, j$ and $j \cdot|\operatorname{str}(\operatorname{child}(v))|+m \leq \operatorname{str}(v)$.

To report the actual occurrences of $P$ in $S$ we traverse all ancestors of $v$ in $\operatorname{dag}(S)$; for each occurrence of $P$ in $\operatorname{str}(v)$ found, recursively visit each parent $u$ of $v$ and offset the location of the occurrence to match the location in $\operatorname{str}(u)$
instead of $\operatorname{str}(v)$. When $u$ is the root, report the occurrence. Observe that the time it takes to traverse the ancestors of $v$ is linear in the number of occurrences we find.

We now describe how to handle fingerprint collisions. Given a z-fast trie, Gagie et al. [9] show how to perform $k$ weak prefix queries and identify all false positives using $O(k \lg m+m)$ extra time by employing bookmarked extraction and bookmarked fingerprinting. Because we only compute fingerprints and extract prefixes (suffixes) of the strings represented by vertices in $\operatorname{dag}(S)$ we do not need bookmarking to do this. We refer the reader to [9] for the details. Thus, we modify the search algorithm such that all the searches in $T_{1}$ and $T_{2}$ are carried out first, then we verify the results before progressing to doing range reporting queries only for ranges that were not discarded during verification.

### 4.3 Correctness

For any occurrence $S[l, r]$ of $P$ in $S$ there is a node $v$ in $\operatorname{sig}(S)$ that stabs $S[l, r]$, ie. a suffix of $\operatorname{pre}(v, i)$ equals a prefix $P[1, j]$ and a prefix of $s u f(v, i)$ equals the remaining suffix $P[j+1, m]$ for some $i$ and $j$. Since we put all combinations of $\operatorname{pre}(v, i), \operatorname{suf}(v, i)$ into $T_{1}, T_{2}$ and $R$, we would be guaranteed to find all nodes $v$ that contains $P$ in $\operatorname{str}(v)$ if we searched for all possible split-points $1, \ldots, m-1$ of $P$ i.e. $P[1, i]$ and $P[i+1, m]$ for $i=1, \ldots, m-1$.

We now argue that we do not need to search for all possible split-points of $P$ but only need to consider those in the set $P_{S}$. For a position $i$, we say the node $v$ stabs $i$ if the nearest common ancestor of the $i^{t h}$ and $i+1^{t h}$ leaf of $\operatorname{sig}(S)$ denoted $N C A\left(l_{i}, l_{i+1}\right)$ is $v$.

Look at any occurrence $S[l, r]$ of $P$. Consider $T_{S}=T(l, r)$ and $T_{P}=\operatorname{sig}(P)$. Look at a possible split-point $i \in[1, m-1]$ and the node $v$ that stabs position $i$ in $T_{P}$. Let $u_{l}$ and $u_{r}$ be adjacent children of $v$ such that the rightmost leaf descendant of $u_{l}$ is the $i^{t h}$ leaf and the leftmost leaf descendant of $u_{r}$ is the $i+1^{\text {th }}$ leaf. We now look at two cases for $v$ and argue it is irrelevant to consider position $i$ as split-point for $P$ in these cases:

1. Case $v$ is consistent (in respect to $T_{P}$ ). In this case it is guaranteed that the node that stabs $l+i$ in $T_{S}$ is identical to $v$. Since $v$ is a descendant of the root of $T_{P}$ (as the root of $T_{P}$ is inconsistent) $\operatorname{str}(v)$ cannot contain $P$ and thus it is irrelevant to consider $i$ as a split-point.
2. Case $v$ is inconsistent and $u_{l}$ and $u_{r}$ are both consistent (in respect to $T_{P}$ ). In this case $u_{l}$ and $u_{r}$ have identical corresponding nodes $u_{l}^{\prime}$ and $u_{r}^{\prime}$ in $T_{S}$. Because $u_{l}$ and $u_{r}$ are children of the same node it follows that $u_{l}^{\prime}$ and $u_{r}^{\prime}$ must also both be children of some node $v^{\prime}$ that stabs $l+i$ in $T_{S}$ (however $v$ and $v^{\prime}$ may not be identical since $v$ is inconsistent). Consider the node $u_{l l}^{\prime}$ to the left of $u_{l}^{\prime}$ (or symmetrically for the right side if $v$ is an inconsistent node in the right side of $T_{P}$ ). If $\operatorname{str}\left(v^{\prime}\right)$ contains $P$ then $u_{l l}^{\prime}$ is also a child of $v^{\prime}$ (otherwise $u_{l}$ would be inconsistent). So it suffices to check the split-point $i-\left|u_{l}\right|$. Surely $i-\left|u_{l}\right|$ stabs an inconsistent node in $T_{P}$, so either we consider that position relevant, or the same argument applies again and a split-point further to the left is eventually considered relevant.

Thus only split-points where $v$ and at least one of $u_{l}$ or $u_{r}$ are inconsistent are relevant. These positions are a subset of the position in $P_{S}$, and thus we try all relevant split-points.

### 4.4 Complexity

A query on $T_{1}$ and $T_{2}$ takes $O(\lg m)$ time by Lemma 2 while a query on $R$ takes $O\left(\lg ^{\epsilon} z\right)$ time using Lemma 1 (i) (excluding reporting). We do $O(\lg m)$ queries as the size of $P_{S}$ is $O(\lg m)$. Verification of the $O(\lg m)$ strings we search for takes total time $O\left(\lg ^{2} m+m\right)=O(m)$. Constructing the signature DAG for $P$ takes $O(m)$ time, thus total time without reporting is $O(m+$ $\left.\lg m \lg ^{\epsilon} z\right)=O\left(m+\lg ^{\epsilon^{\prime}} z\right.$ ) for any $\epsilon^{\prime}>\epsilon$. This holds because if $m \leq \lg ^{2 \epsilon} z$ then $\lg m \lg ^{\epsilon} z \leq \lg \lg ^{2 \epsilon} z \lg ^{\epsilon} z=O\left(\lg ^{\epsilon^{\prime}} z\right)$, otherwise $m>\lg ^{2 \epsilon} z \Leftrightarrow \sqrt{m}>\lg ^{\epsilon} z$ and then $\lg m \lg ^{\epsilon} z=O(\lg m \sqrt{m})=O(m)$. For every query on $R$ we may find multiple points each corresponding to an occurrence of $P$. It takes $O\left(\lg ^{\epsilon} z\right)$ time to report each point thus the total time becomes $O\left(m+(1+\right.$ occ $\left.) \lg ^{\epsilon^{\prime}} z\right)$.

## 5 Short Patterns

Our solution for short patterns uses properties of the LZ77-parse of S. A primary substring of $S$ is a substring that contains one or more borders of $S$, all other substrings are called secondary. A primary substring that matches a query pattern $P$ is a primary occurrence of $P$ while a secondary substring that matches $P$ is a secondary occurrence of $P$. In a seminal paper on LZ77 based indexing [13] Kärkkäinen and Ukkonen use some observations by Farach and Thorup [6] to show how all secondary occurrences of a query pattern $P$ can be found given a list of the primary occurrences of $P$ through a reduction to orthogonal range reporting. Employing the range reporting result given in Lemma (ii), all secondary occurrences can be reported as stated in the following lemma:

Lemma 10 (Kärkkäinen and Ukkonen [13]). Given the LZ77-parse of a string $S$ there exists a data structure that uses $O(z \lg \lg z)$ space that can report all secondary occurrences of a pattern $P$ given the list of primary occurrences of $P$ in $S$ in $O(o c c \lg \lg n)$ time.

We now describe a data structure that can report all primary occurrences of a pattern $P$ of length at most $k$ in $O(m+o c c)$ time using $O(z k)$ space.

Lemma 11. Given a string $S$ of length $n$ and a positive integer $k \leq n$ we can build a compressed index supporting pattern matching queries for patterns of length $m$ in $O(m+$ occ $\lg \lg n)$ time using $O(z k+z \lg \lg z)$ space that works for $m \leq k$.

Proof. Consider the set $C$ of $z$ substrings of $S$ that are defined by $S\left[u_{i}-k, u_{i}+\right.$ $k-1]$ for $i \in[z]$, ie. the substrings of length $2 k$ surrounding the borders of the LZ77-parse. The total length of these strings is $\Theta(z k)$. Construct the generalized
suffix tree $T$ over the set of strings $C$. This takes $\Theta(z k)$ words of space. To ensure no occurrence is reported more than once, if multiple suffixes in this generalized suffix tree correspond to substrings of $S$ that starts on the same position in $S$, only include the longest of these. This happens when the distance between two borders is less than $2 k$.

To find the primary occurrences of $P$ of length $m$, simply find all occurrences of $P$ in $T$. These occurrences are a super set of the primary occurrences of $P$ in $S$, since $T$ contains all substrings starting/ending at most $k$ positions from a border. It is easy to filter out all occurrences that are not primary, simply by calculating if they cross a border or not. This takes $O(m+$ occ) time (where occ includes secondary occurrences). Combined with Lemma 10 this gives Lemma 11 .

## 6 Semi-Short Patterns

In this section, we show how to handle patterns of length between $\lg \lg z$ and $\lg ^{\epsilon} z$. It is based on the same reduction to 2D-range reporting as used for long patterns. However, the positions in $S$ that are inserted in the range reporting structure is now based on the LZ77-parse of $S$ instead. Furthermore we use Lemma 1 (ii) which gives faster range reporting but uses super-linear space, which is fine because we instead put fewer points into the structure. We get the following lemma:

Lemma 12. Given a string $S$ of length $n$ we solve the compressed indexing problem for a pattern $P$ of length $m$ with $\lg \lg z \leq m \leq \lg ^{\epsilon} z$ for any positive constant $\epsilon<\frac{1}{2}$ in $O\left(m+\right.$ occ $\left.\left(\lg \lg n+\lg ^{\epsilon} z\right)\right)$ time using $O(z(\lg \lg z+\log (n / z)))$ space.

### 6.1 Data Structure

As in the previous section for short patterns, we only need to worry about primary occurrences of $P$ in $S$. Let $B$ be the set of all substrings of length at most $\lg ^{\epsilon} z$ that cross a border in $S$. The split positions of such a string are the offsets of the leftmost borders in its occurrences. All primary occurrences of $P$ in $S$ are in this set. The size of this set is $|B|=O\left(z \lg ^{2 \epsilon} z\right)$. The data structure is composed by the following:

- A dictionary $H$ mapping each string in $B$ to its split positions.
- A z-fast trie $T_{1}$ on the reverse of the strings $T\left[u_{i}, l_{i}\right]$ for $i \in[z]$.
- A z-fast trie $T_{2}$ on the strings $T\left[u_{i}, n\right]$ for $i \in[z]$.
- A range reporting data structure $R$ with a point $(c, d)$ for every pair of strings $C_{i}=T\left[u_{i}, l_{i}\right], D_{i}=T\left[u_{i+1}, n\right]$ for $i \in[z]$ where $D_{z}=\epsilon$ and $c$ is the lexicographical rank of the reverse of $C_{i}$ in the set $\left\{C_{1}, \ldots, C_{z}\right\}$ and $d$ is the lexicographical rank of $D_{i}$ in the set $\left\{D_{1}, \ldots D_{z}\right\}$. We store the border $u_{i}$ along with the point $(c, d)$.
- The data structure described in Lemma 10 to report secondary occurrences.
- The signature grammar for $S$.

Each entry in $H$ requires $\lg ^{\lg }{ }^{\epsilon} z=O(\lg \lg z)$ bits to store since a split position can be at most $\lg ^{\epsilon} z$. Thus the dictionary can be stored in $O(|B| \cdot \lg \lg z)=$ $O\left(z \lg ^{2 \epsilon} z \lg \lg z\right)$ bits which for $\epsilon<\frac{1}{2}$ is $O(z)$ words. The tries $T_{1}$ and $T_{2}$ take $O(z)$ space while $R$ takes $O(z \lg \lg z)$ space. The signature grammar takes $O(z \log (n / z))$. Thus the total space is $O(z(\lg \lg z+\log (n / z)))$.

### 6.2 Searching

Assume a lookup for $P$ in $H$ does not give false-positives. Given a pattern $P$ compute all prefix fingerprints of $P$. Next do a lookup in $H$. If there is no match then $P$ does not occur in $S$. Otherwise, we do the following for each of the splitpoints $s$ stored in $H$. First split $P$ into a left part $P_{l}=P[0, s-1]$ and a right part $P_{r}=P[s, m]$. Then search for the reverse of $P_{l}$ in $T_{1}$ and for $P_{r}$ in $T_{2}$ using the corresponding fingerprints. The search induces a (possibly empty) range for which we do a range reporting query in $R$. Each occurrence in $R$ corresponds to a primary occurrence of $P$ in $S$, so report these. Finally use Lemma 10 to report all secondary occurrences.

Unfortunately, we cannot guarantee a lookup for $P$ in $H$ does not give a false positive. Instead, we pause the reporting step when the first possible occurrence of $P$ has been found. At this point, we verify the substring $P$ matches the found occurrence in $S$. We know this occurrence is around an LZ-border in $S$ such that $P_{l}$ is to the left of the border and $P_{r}$ is to the right of the border. Thus we can efficiently verify that $P$ actually occurs at this position using the grammar.

### 6.3 Analysis

Computing the prefix fingerprints of $P$ takes $O(m)$ time. First, we analyze the running time in the case $P$ actually exists in $S$. The lookup in $H$ takes $O(1)$ time using perfect hashing. For each split-point we do two z-fast trie lookups in time $O(\lg m)=O(\lg \lg z)$. Since each different split-point corresponds to at least one unique occurrence, this takes at most $O$ (occ $\lg \lg z)$ time in total. Similarly each lookup and occurrence in the 2D-range reporting structure takes $\lg \lg z$ time, which is therefore also bounded by $O(\operatorname{occ} \lg \lg z)$ time. Finally, we verified one of the found occurrence against $P$ in $O(m)$ time. So the total time is $O(m+$ occ $\lg \lg z)$ in this case.

In the case $P$ does not exists, either the lookup in $H$ tells us that, and we spend $O(1)$ time, or the lookup in $H$ is a false-positive. In the latter case, we perform exactly two z-fast trie lookups and one range reporting query. These all take time $O(\lg \lg z)$. Since $m \geq \lg \lg z$ this is $O(m)$ time. Again, we verified the found occurrence against $P$ in $O(m)$ time. The total time in this case is therefore $O(m)$.

Note we ensure our fingerprint function is collision free for all substrings in $B$ during the preprocessing thus there can only be collisions if $P$ does not occur in $S$ when $m \leq \lg ^{\epsilon} z$.

## 7 Randomized Solution

In this section we present a very simple way to turn the $O\left(m+(1+\right.$ occ $\left.) \lg ^{\epsilon} z\right)$ worst-case time of Lemma 9 into $O\left(m+\right.$ occ $\left.^{\lg }{ }^{\epsilon} z\right)$ expected time. First observe, this is already true if the pattern we search for occurs at least once or if $m \geq \lg ^{\epsilon} z$.

As in the semi-short patterns section, we consider the set $B$ of substrings of $S$ of length at most $\lg ^{\epsilon} z$ that crosses a border. Create a dictionary $H$ with $z \lg ^{3 \epsilon} z$ entries and insert all the strings from $B$. This means only a $\lg ^{\epsilon} z$ fraction of the entries are used, and thus if we lookup a string $s$ (where $|s| \leq \lg ^{\epsilon} z$ ) that is not in $H$ there is only a $\frac{1}{\lg ^{\epsilon} z}$ chance of getting a false-positive.

Now to answer a query, we first check if $m \leq \lg ^{\epsilon} z$ in which case we look it up in $H$. If it does not exist, report that. If it does exist in $H$ or if $m>\lg ^{\epsilon} z$ use the solution from Lemma 9 to answer the query.

In the case $P$ does not exist, we spend either $O(m)$ time if $H$ reports no, or $O\left(m+\lg ^{\epsilon} z\right)$ time if $H$ reports a false-positive. Since there is only $\frac{1}{\lg ^{\epsilon} z}$ chance of getting a false positive, the expected time in this case is $O(m)$. In all other cases, the running time is $O\left(m+\right.$ occ $\left.\lg ^{\epsilon} z\right)$ in worst-case, so the total expected running time is $O\left(m+{ }^{\circ} c \lg ^{\epsilon} z\right)$. The space usage of $H$ is $O\left(z \lg ^{3 \epsilon} z\right)$ bits since we only need to store one bit for each entry. This is $O(z)$ words for $\epsilon \leq 1 / 3$. To sum up, we get the following lemma:

Lemma 13. Given a signature grammar for a text $S$ of length $n$ with an LZ77parse of length $z$ we can build a compressed index supporting pattern matching queries in $O\left(m+\mathrm{occ}^{\lg }{ }^{\epsilon} z\right)$ expected time using $O(z \lg (n / z))$ space for any constant $0<\epsilon \leq 1 / 3$.

## References

1. Stephen Alstrup, Gerth Stølting Brodal, and Theis Rauhe. Pattern matching in dynamic texts. In Proceedings of the 11th Annual Symposium on Discrete Algorithms. Citeseer, 2000.
2. Djamal Belazzougui, Paolo Boldi, Rasmus Pagh, and Sebastiano Vigna. Fast prefix search in little space, with applications. In Proc. 18th ESA, pages 427-438, 2010.
3. Philip Bille, Mikko Berggren Ettienne, Inge Li Gørtz, and Hjalte Wedel Vildhøj. Time-space trade-offs for Lempel-Ziv compressed indexing. In 28th Annual Symposium on Combinatorial Pattern Matching. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2017.
4. Timothy M. Chan, Kasper Green Larsen, and Mihai Patrascu. Orthogonal range searching on the ram, revisited. In Proc. 27th SOCG, pages 1-10, 2011.
5. Richard Cole and Uzi Vishkin. Deterministic coin tossing with applications to optimal parallel list ranking. Information and Control, 70(1):32-53, 1986.
6. M. Farach and M. Thorup. String Matching in Lempel-Ziv Compressed Strings. Algorithmica, 20(4):388-404, 1998.
7. M. L. Fredman and D. E. Willard. Blasting through the information theoretic barrier with fusion trees. In Proceedings of the Twenty-second Annual ACM Symposium on Theory of Computing, STOC '90, pages 1-7, New York, NY, USA, 1990. ACM.
8. Travis Gagie, Paweł Gawrychowski, Juha Kärkkäinen, Yakov Nekrich, and Simon Puglisi. A faster grammar-based self-index. Language and Automata Theory and Applications, pages 240-251, 2012.
9. Travis Gagie, Paweł Gawrychowski, Juha Kärkkäinen, Yakov Nekrich, and Simon J Puglisi. LZ77-based self-indexing with faster pattern matching. In Latin American Symposium on Theoretical Informatics, pages 731-742. Springer, 2014.
10. Travis Gagie, Gonzalo Navarro, and Nicola Prezza. Optimal-time text indexing in bwt-runs bounded space. arXiv preprint arXiv:1705.10382, 2017.
11. Paweł Gawrychowski, Adam Karczmarz, Tomasz Kociumaka, Jakub Łacki, and Piotr Sankowski. Optimal dynamic strings. arXiv preprint arXiv:1511.02612, 2015.
12. Artur Jeż. Faster fully compressed pattern matching by recompression. ACM Transactions on Algorithms (TALG), 11(3):20, 2015.
13. Juha Kärkkäinen and Esko Ukkonen. Lempel-Ziv parsing and sublinear-size index structures for string matching. Proceedings of the 3rd South American Workshop on String Processing (WSP'96), 26(Teollisuuskatu 23):141-155, 1996.
14. Richard M. Karp and Michael O. Rabin. Efficient randomized pattern-matching algorithms. IBM J. Res. Dev., 31(2):249-260, 1987.
15. Kurt Mehlhorn, Rajamani Sundar, and Christian Uhrig. Maintaining dynamic sequences under equality tests in polylogarithmic time. Algorithmica, 17(2):183198, 1997.
16. Gonzalo Navarro and Veli Mäkinen. Compressed full-text indexes. ACM Computing Surveys (CSUR), 39(1):2, 2007.
17. Takaaki Nishimoto, I Tomohiro, Shunsuke Inenaga, Hideo Bannai, and Masayuki Takeda. Dynamic index, LZ factorization, and LCE queries in compressed space. arXiv preprint arXiv:1504.06954, 2015.
18. Benny Porat and Ely Porat. Exact and approximate pattern matching in the streaming model. In Proc. 50th FOCS, pages 315-323, 2009.
19. S. C. Sahinalp and U. Vishkin. Efficient approximate and dynamic matching of patterns using a labeling paradigm. In Proceedings of 37th Conference on Foundations of Computer Science, Oct 1996.
20. I Tomohiro. Longest common extension with recompression. 2017.
21. Jacob Ziv and Abraham Lempel. A Universal Algorithm for Sequential Data Compression. IEEE Transactions on Information Theory, (3), 1977.
