

Compressed Multiple Pattern Matching

Dmitry Kosolobov

University of Helsinki, Helsinki, Finland
dkosolobov@mail.ru

Nikita Sivukhin

Ural Federal University, Ekaterinburg, Russia
sivukhin.nikita@yandex.ru

Abstract

Given d strings over the alphabet $\{0, 1, \dots, \sigma-1\}$, the classical Aho–Corasick data structure allows us to find all occ occurrences of the strings in any text T in $O(|T| + occ)$ time using $O(m \log m)$ bits of space, where m is the number of edges in the trie containing the strings. Fix any constant $\varepsilon \in (0, 2)$. We describe a compressed solution for the problem that, provided $\sigma \leq m^\delta$ for a constant $\delta < 1$, works in $O(|T|^\frac{1}{\varepsilon} \log \frac{1}{\varepsilon} + occ)$ time, which is $O(|T| + occ)$ since ε is constant, and occupies $mH_k + 1.443m + \varepsilon m + O(d \log \frac{m}{d})$ bits of space, for all $0 \leq k \leq \max\{0, \alpha \log_\sigma m - 2\}$ simultaneously, where $\alpha \in (0, 1)$ is an arbitrary constant and H_k is the k th-order empirical entropy of the trie. Hence, we reduce the $3.443m$ term in the space bounds of previously best succinct solutions to $(1.443 + \varepsilon)m$, thus solving an open problem posed by Belazzougui. Further, we notice that $L = \log \binom{\sigma(m+1)}{m} - O(\log(\sigma m))$ is a worst-case space lower bound for any solution of the problem and, for $d = o(m)$ and constant ε , our approach allows to achieve $L + \varepsilon m$ bits of space, which gives an evidence that, for $d = o(m)$, the space of our data structure is theoretically optimal up to the εm additive term and it is hardly possible to eliminate the term $1.443m$. In addition, we refine the space analysis of previous works by proposing a more appropriate definition for H_k . We also simplify the construction for practice adapting the fixed block compression boosting technique, then implement our data structure, and conduct a number of experiments showing that it is comparable to the state of the art in terms of time and is superior in space.

2012 ACM Subject Classification Theory of computation → Pattern matching

Keywords and phrases multiple pattern matching, compressed space, Aho–Corasick automaton

Digital Object Identifier 10.4230/LIPIcs.CPM.2019.13

Related Version <https://arxiv.org/abs/1811.01248>

Funding Supported by the Russian Science Foundation (RSF), project 18-71-00002.

1 Introduction

Searching for multiple patterns in text is a fundamental stringology problem that has numerous applications, including bioinformatics [20], search engines [36], intrusion detection systems [27, 32], shortest superstring approximation [2], and others. The classical solution for the multiple pattern matching is the Aho–Corasick data structure [1], which, however, does not always fulfil space requirements of many such applications due to the rapid growth of the amounts of data in modern systems. To address this issue, several space-efficient multiple pattern matching data structures were developed in the last decade. In this paper we improve the space consumption in a state-of-the-art solution for the problem, simplify the compression method used in this solution by adapting the known fixed block compression boosting technique, and give an evidence that the achieved space is, in a sense, close to optimal; in addition, we refine the theoretical space analysis of a previous best result, and implement our construction and conduct a number of experiments showing that it is comparable to the existing practical data structures in terms of time and is superior in space. Before discussing our contribution in details, let us briefly survey known results in this topic.



© Dmitry Kosolobov and Nikita Sivukhin;

licensed under Creative Commons License CC-BY

30th Annual Symposium on Combinatorial Pattern Matching (CPM 2019).

Editors: Nadia Pisanti and Solon P. Pissis; Article No. 13; pp. 13:1–13:14

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The Aho–Corasick solution builds a trie of all patterns augmented with additional structures overall occupying $O(m \log m)$ bits (hereafter, \log denote logarithm with base 2), where m is the number of edges in the trie, and allows us to find all *occ* occurrences of the patterns in any text T in $O(|T| + \text{occ})$ time. The $O(m \log m)$ -bit space overhead imposed by this solution might be unacceptably high if one is processing large sets of patterns; for this case, several succinct and compressed data structures were developed in the last decade [3, 5, 11, 21, 22, 25, 35]. In this paper we are especially interested in the two closely related results from [3] and [21], which provide currently the best time and space bounds for the multiple pattern matching problem. In [3] Belazzougui designed a compact representation of the Aho–Corasick scheme that works in the same $O(|T| + \text{occ})$ time but stores the trie with all additional structures in only $m \log \sigma + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits, where σ is the alphabet size and d is the number of patterns; in addition, he showed that the space can be improved to $mH_0 + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits with no slowdown provided $\sigma \leq m^\delta$ for a constant $\delta < 1$, where H_0 is the zeroth-order empirical entropy of the trie. In [21] Hon et al. further lowered the space to $mH_k + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits (again, assuming $\sigma \leq m^\delta$) by simply applying the compression boosting technique [12], where H_k is the k th-order empirical entropy of the trie (see clarifications below) and k is any fixed integer such that $0 \leq k \leq \alpha \log_\sigma m - 1$, for arbitrary constant $\alpha \in (0, 1)$. This topic is rich with other related results, which, for instance, support dynamic modifications of the patterns [5, 10, 11], try to process T in real-time [25], allow randomization [7, 18], consider hardware implementations [9], etc. In this paper we focus on the basic functionality as in [3] and [21].

Belazzougui posed the following open problem [3]: can we reduce the constant 3.443 in the space of his (and Hon et al.’s) result without any significant slowdown? We solve this problem affirmatively describing a data structure that, provided $\sigma \leq m^\delta$ for a constant $\delta < 1$, occupies $mH_k + 1.443m + \varepsilon m + O(d \log \frac{m}{d})$ bits of space, for an arbitrarily chosen constant $\varepsilon \in (0, 2)$, and answers pattern matching queries on any text T in $O(|T|\varepsilon^{-1} \log \varepsilon^{-1} + \text{occ})$ time, which is $O(|T| + \text{occ})$ since ε is constant. Then, we notice that $L = \log \binom{\sigma(m+1)}{m} - O(\log(\sigma m))$ is a worst-case space lower bound for any multiple pattern matching data structure and, for $d = o(m)$ and constant ε , our construction allows to achieve $L + \varepsilon m$ bits of space; observe that, for $\sigma = \omega(1)$, we have $L = m \log \sigma + m \log e + o(m) \approx m \log \sigma + 1.443m + o(m)$ (see [3]), which gives an evidence that, for $d = o(m)$, our space bound is optimal up to the εm additive term and it is hardly possible to remove the term $1.443m$. In addition, we argue that the definition of H_k borrowed by Hon et al. [21] from [13], denoted H_k^* in our paper, is not satisfactory: in particular, H_k^* can be greater than $\log \sigma$, which contradicts the idea of the empirical entropy (H_k^* was devised in [13] for a slightly different problem); we propose a more appropriate definition for H_k , which is not worse since $H_k \leq H_k^*$, and refine the analysis of Hon et al. showing that their data structure indeed occupies $mH_k + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits even according to our definition of H_k . Further, our solution, unlike [21], does not require to fix k and the space bound holds for all k such that $0 \leq k \leq \max\{0, \alpha \log_\sigma m - 2\}$ simultaneously; this is achieved by adapting the fixed block compression boosting technique [17, 24] to our construction, which is also better for practice than the compression boosting used in [21]. Finally, we implement our data structure and conduct a number of experiments showing that it is comparable to the state of the art in terms of time and is superior in space.

The paper is organized as follows. In the following section we introduce some basic notions and define the k th-order empirical entropy of tries. In Section 3 we survey the solution of Belazzougui [3]. In Section 4 we discuss compression boosting techniques and investigate the flaws of the space analysis of Hon et al. [21]. Section 5 describes our main data structure and considers space optimality. Appendix contains implementation details and experiments.

2 Preliminaries

Throughout the paper, we mainly consider strings drawn from the alphabet $\{0, 1, \dots, \sigma-1\}$ of size σ (not necessarily constant). For a string $s = c_0c_1 \cdots c_{n-1}$, denote by $|s|$ its length n . The *reverse* $c_{n-1} \cdots c_1c_0$ of s is denoted s^r . We write $s[i]$ for the letter c_i of s and $s[i..j]$ for the *substring* $c_i c_{i+1} \cdots c_j$, assuming $s[i..j]$ is the empty string if $i > j$. We say that a string p *occurs* in s at position i if $s[i..i+|p|-1] = p$. A string p is called a *prefix* (resp., *suffix*) of s if p occurs in s at position 0 (resp., $|s| - |p|$). For integer segments, we use the following notation: $[i..j] = \{i, i+1, \dots, j\}$, $(i..j) = [i..j] \setminus \{i\}$, $[i..j) = [i..j] \setminus \{j\}$. The set of all strings of lengths k over an alphabet A is denoted A^k ; we use this notation for the set $[0..\sigma]^k$. For letter c , the only string of the set $\{c\}^k$ is denoted c^k .

The trie containing a set of strings S is the minimal in the number of vertices rooted tree with edges labeled by letters so that each $s \in S$ can be spelled out on the path from the root to a vertex. For vertex v , denote the string written on the path from the root to v by $\text{str}(v)$.

The *zeroth-order empirical entropy* (see [26, 29]) of a string t of length n is defined as $H_0(t) = \sum_{c \in [0..\sigma]} \frac{n_c}{n} \log \frac{n}{n_c}$, where n_c is the number of letters c in t and $\frac{n_c}{n} \log \frac{n}{n_c} = 0$ whenever $n_c = 0$. For a string w of length k , let t_w be a string formed by concatenating all letters immediately following occurrences of w in the string $\$^k t$, where $\$ = \sigma$ is a new special letter introduced for technical convenience; e.g., $t_{ab} = aac$ and $t_{\$a} = b$ for $t = abababc$. The *kth-order empirical entropy* of t is defined as $H_k(t) = \sum_{w \in [0..\sigma]^k} \frac{|t_w|}{n} H_0(t_w)$ (see [26, 28, 29]). It is well known that $\log \sigma \geq H_0 \geq H_1 \geq \dots$ and H_k makes sense as a measure of string compression only for $k < \log_\sigma n$ (see [15] for a deep discussion). For the sake of completeness, let us show that $H_k \geq H_{k+1}$; this proof then can be easily adapted for the empirical entropy of tries below. Curiously, to our knowledge, all sources refer to this simple and intuitive fact as “obvious” but do not give a proof; even the original paper [29], the survey [28], and the book [30].

► **Lemma 1** (see [24, Lemma 3]). *For any strings t_1, t_2, \dots, t_ℓ and the string $t = t_1 t_2 \cdots t_\ell$, we have $|t|H_0(t) \geq \sum_{i=1}^\ell |t_i|H_0(t_i)$.*

Since, without loss of generality, one can assume that $t_w = t_{a_0 w} t_{a_1 w} \cdots t_{a_\sigma w}$, where a_0, \dots, a_σ are all letters of $[0..\sigma]$, Lemma 1 directly implies the inequality $|t|H_k(t) = \sum_{w \in [0..\sigma]^k} |t_w|H_0(t_w) \geq \sum_{w \in [0..\sigma]^{k+1}} |t_w|H_0(t_w) = |t|H_{k+1}(t)$ and, hence, $H_k \geq H_{k+1}$.

By analogy, one can define the empirical entropy for tries (see also [13]). Let \mathcal{T} be a trie with n edges over the alphabet $[0..\sigma]$. For a string w of length k , denote by \mathcal{T}_w a string formed by concatenating in an arbitrary order the letters on the edges (u, v) of \mathcal{T} (here u is the parent of v) such that w is a suffix of $\$^k \text{str}(u)$; e.g., $\mathcal{T}_{\$k}$ consists of all letters labeling the edges incident to the root. Then, the *kth-order empirical entropy* of \mathcal{T} is defined as $H_k(\mathcal{T}) = \sum_{w \in [0..\sigma]^k} \frac{|\mathcal{T}_w|}{n} H_0(\mathcal{T}_w)$. (Note that $\sum_{w \in [0..\sigma]^k} |\mathcal{T}_w| = n$.) Analogously to the case of strings, one can show that $\log \sigma \geq H_0(\mathcal{T}) \geq H_1(\mathcal{T}) \geq \dots$ and $H_k(\mathcal{T})$ makes sense as a measure of compression only for $k < \log_\sigma n$. For the definition of the *kth-order empirical entropy* of tries as given by Hon et al. [21], see Section 4.

3 Basic Algorithm

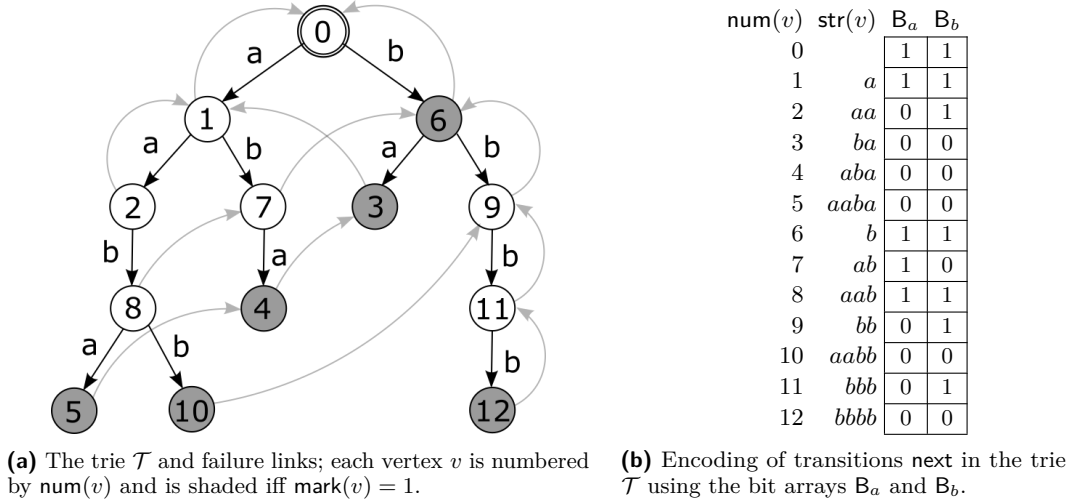
Given a dictionary \mathcal{D} of d patterns, the multiple pattern matching problem is to preprocess \mathcal{D} in order to efficiently find all occurrences of the patterns in an arbitrary given text. In this section we describe for this problem the classical Aho–Corasick solution [1] and its space-efficient version developed by Belazzougui [3].

13:4 Compressed Multiple Pattern Matching

The main component of the Aho–Corasick data structure is the trie \mathcal{T} containing \mathcal{D} . Each vertex v of the trie is augmented with the following structures (see Figure 1a):

- a flag $\text{mark}(v)$ that indicates whether $\text{str}(v) \in \mathcal{D}$: it is so iff $\text{mark}(v) = 1$;
- a hash table $\text{next}(v, \cdot)$ that, for each letter c , either maps c to a vertex $u = \text{next}(v, c)$ such that $\text{str}(u) = \text{str}(v)c$, or returns nil if there is no such u ;
- a link $\text{failure}(v)$ to a vertex such that $\text{str}(\text{failure}(v))$ is the longest proper suffix of $\text{str}(v)$ that can be spelled out on a root-vertex path ($\text{failure}(v)$ is undefined if v is the root);
- a link $\text{report}(v)$ to a vertex such that $\text{str}(\text{report}(v))$ is the longest proper suffix of $\text{str}(v)$ that belongs to \mathcal{D} , or $\text{report}(v) = \text{root}$ if there is no such suffix.

It is well known that the described data structure allows us to find all *occ* occurrences of the patterns from \mathcal{D} in text T in $O(|T| + \text{occ})$ time: we read T from left to right maintaining a “current” vertex v in \mathcal{T} (initially, v is the root) and, when a new letter $T[i]$ arrives, we put $v = \text{next}(v, T[i])$ for the first u in the series $v, \text{failure}(v), \text{failure}(\text{failure}(v)), \dots$ for which $\text{next}(u, T[i])$ is not nil, or put $v = \text{root}$ if there is no such u , then we report all patterns ending at position i using the report links and the flag $\text{mark}(v)$ (see details in, e.g., [8]).



■ **Figure 1** The trie \mathcal{T} with $\mathcal{D} = \{aaba, aabb, aba, b, ba, bbbb\}$.

It is easy to see that the data structure occupies $O(m \log m)$ bits of space, where m is the number of edges in \mathcal{T} . Let us describe now how Belazzougui could reduce the space consumption of this solution with no slowdown (see [3] for a more detailed explanation).

First, he assigned to each vertex v of \mathcal{T} a unique number $\text{num}(v) \in [0..m]$ so that, for any two vertices u and v , we have $\text{num}(u) < \text{num}(v)$ iff $\text{str}(u)^r < \text{str}(v)^r$ lexicographically (see Figure 1a). As it was pointed out by Hon et al. [21], this subtle numbering scheme corresponds to the numbering of vertices in the so-called XBW of \mathcal{T} (see [13]), a generalization of the classical Burrows–Wheeler transform (BWT) [4] for tries. It turns out that the numbering allows us to organize fast navigation in the trie in small space, simulating the tables next in a manner that resembles the so-called FM-indexes [14] based on the BWT.

For each letter $c \in [0..\sigma)$, define a bit array $B_c[0..m]$ such that, for any vertex v , we have $B_c[\text{num}(v)] = 1$ iff $\text{next}(v, c) \neq \text{nil}$ (see Figure 1b). Let $\text{rank}(i, B_c)$ be an operation on B_c , called *partial rank*, that returns nil if $B_c[i] \neq 1$, and returns the number of ones in $B_c[0..i]$ otherwise. By standard arguments, one can show that $\text{num}(\text{next}(v, c)) = \text{rank}(\text{num}(v), B_c) + e_{<c}$, provided $\text{next}(v, c) \neq \text{nil}$, where $e_{<c}$ is the number of edges in the trie with labels smaller than c

(see [3, 13]). Let us concatenate $B_0, B_1, \dots, B_{\sigma-1}$, thus obtaining a new bit array B of length $(m+1)\sigma$. Since $e_{<c}$ is equal to the number of ones in the arrays B_0, B_1, \dots, B_{c-1} , we obtain $\text{num}(\text{next}(v, c)) = \text{rank}(c(m+1) + \text{num}(v), B)$ and $\text{rank}(c(m+1) + \text{num}(v), B) = \text{nil}$ iff $\text{next}(v, c) = \text{nil}$. In order to support rank in $O(1)$ time, we equip B with the following data structure, which, in addition, supports the operation $\text{select}(i, B)$ that returns the position of the i th one in B (select is used below to compute the parent of v by number $\text{num}(v)$; see [3]).

► **Lemma 2** (see [33]). *Every bit array of length n with m ones has an encoding that occupies $\log \binom{n}{m} + o(m)$ bits and supports select and partial rank queries in $O(1)$ time.*

Since B contains exactly m ones, the space occupied by B encoded as in Lemma 2 is $\log \binom{(m+1)\sigma}{m} + o(m)$, which in [3] was estimated by $m \log \sigma + 1.443m + o(m)$ (here $1.443 \approx \log e$). To encode mark , Belazzougui constructs a bit array of length $m+1$ containing exactly d ones at positions $\text{num}(v)$ for all d vertices v with $\text{mark}(v) = 1$, and stores the array as in Lemma 2 (for this array, we need only access queries, which can be simulated by rank), thus occupying $\log \binom{m+1}{d} + o(d) \leq O(d \log \frac{m}{d})$ bits of space. It remains to encode failure and report links.

Belazzougui noticed that the failure links form a tree on the vertices of \mathcal{T} in which, for vertex v , $\text{failure}(v)$ returns the parent of v ; more importantly, the numbering num corresponds to the order of vertices in a depth first traversal of this tree. This allows us to represent the tree of failure links in $2m + o(m)$ bits using the following lemma.

► **Lemma 3** (see [31]). *Every tree with m edges can be encoded in $2m + o(m)$ bits with the support of the operation $\text{parent}(v)$, which returns the parent of vertex v , in constant time, provided all the vertices are represented by their numbers in a depth first traversal.*

An analogous observation is also true for the report links and the tree induced by them has at most d internal vertices, which allows to spend only $O(d \log \frac{m}{d}) + o(m)$ bits for it.

► **Lemma 4** (see [23]). *Every tree with m edges and d internal vertices can be encoded in $d \log \frac{m}{d} + O(d) + o(m)$ bits with the support of the operation $\text{parent}(v)$ in constant time, provided all the vertices are represented by their numbers in a depth first traversal.*

Thus, the total space consumed by the succinct versions of the structures mark , next , failure , and report is $m \log \sigma + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits.

4 Compression Boosting

In [3] it was noticed that if each array $B_0, B_1, \dots, B_{\sigma-1}$ is encoded separately using Lemma 2, then they altogether occupy $\sum_{0 \leq c < \sigma} (\log \binom{m+1}{n_c} + o(n_c))$ bits, where n_c is the number of labels c in \mathcal{T} , which is upper bounded by $\sum_{0 \leq c < \sigma} (n_c \log \frac{m}{n_c} + 1.443n_c) + o(m) = mH_0(\mathcal{T}) + 1.443m + o(m)$ (here we apply the inequalities $\log \binom{m+1}{n} \leq n \log \frac{m+1}{n} + n \log e$ and $\log e < 1.443$, and estimate $n_c \log \frac{m+1}{n_c}$ as $n_c \log \frac{m}{n_c} + o(n_c)$; see [3]). Using such separated array encodings and some auxiliary data structures, one can reduce the space of the whole data structure to $mH_0(\mathcal{T}) + 3.443m + o(m) + O(d \log \frac{m}{d})$ bits, provided $\sigma \leq m^\delta$ for some constant $\delta < 1$. Hon et al. [21] further developed this idea, applying the compression boosting technique [12].

Compression boosting. Choose an integer k such that $0 \leq k \leq \alpha \log_\sigma m - 1$, where $\alpha \in (0, 1)$ is an arbitrary fixed constant. For $i \in [0..m]$, denote by $\text{istr}(i)$ the string $\text{str}(v)$ such that v is the vertex of \mathcal{T} with $\text{num}(v) = i$. For technical reasons, we introduce a special letter $\$ = \sigma$. Hon et al. [21] partition the set $[0..m]$ into segments in which the strings $\text{istr}(i)$, for i from the same segment, have a common suffix of length k ; then the subarrays of B_c corresponding

to these segments are encoded separately using Lemma 2. More precisely, let $[\ell_\rho..r_\rho]$, for $\rho \in [0..\sigma]^k$, be the set of all $i \in [0..m]$ such that ρ is a suffix of $\$^k \text{istr}(i)$ (the definition of `num` implies that this set forms a segment); each array $B_c[0..m]$, for $c \in [0..\sigma)$, is partitioned into the subarrays $B_c[\ell_\rho..r_\rho]$, where $\rho \in [0..\sigma]^k$ (empty segments $[\ell_\rho..r_\rho]$ are excluded), and each such subarray is encoded separately using Lemma 2. Then, all the encoded subarrays in total occupy $\sum_{0 \leq c < \sigma} \sum_{\rho \in [0..\sigma]^k} (\log \binom{m_\rho}{n_{c,\rho}} + o(n_{c,\rho}))$ bits, where $m_\rho = r_\rho - \ell_\rho + 1$ and $n_{c,\rho}$ is the number of ones in the subarray $B_c[\ell_\rho..r_\rho]$. Hon et al. upper bound this sum by $mH_k^*(\mathcal{T}) + 1.443m + o(m)$, where $H_k^*(\mathcal{T}) = \frac{m+\ell}{m+1}H_k(\mathcal{T}^*)$ is their definition of the k th-order empirical entropy (see [21]; for simplicity, we use our notation H_k to define H_k^*) in which ℓ is the number of leaves in \mathcal{T} and \mathcal{T}^* is the trie obtained from \mathcal{T} by attaching to each leaf an outgoing edge labeled with $\$$. Note that $H_k^*(\mathcal{T}) \geq H_k(\mathcal{T}^*)$.

We believe that the definition of H_k^* by Hon et al. is not satisfactory. The problem is that the inequality $H_k^*(\mathcal{T}) \leq \log \sigma$ (and even $H_k(\mathcal{T}^*) \leq \log \sigma$), which seems to be natural for any proper definition of the empirical entropy, does not necessarily hold; as a corollary, according to the analysis of Hon et al., the encoding of next occupying $m \log \sigma + 1.443m$ bits in the Belazzougui's data structure might "grow" after compression up to $mH_k^*(\mathcal{T}) + 1.443m$ bits (however, there is no growth in reality, just the upper bound of Hon et al. is too rough). For instance, one can observe such behavior in the trie \mathcal{T} of all strings of length h over the alphabet $\{0, 1\}$: while it is straightforward that $H_1(\mathcal{T}) = \log \sigma = 1$, it can be shown that $H_1(\mathcal{T}^*) \approx \log 3$ since, for $c \in \{0, 1\}$, the string \mathcal{T}_c^* consisting of all labels in \mathcal{T}^* with "context" c contains roughly $m/4$ of each of the letters 0, 1, and $\$$, where $m = 2^{h+1} - 2$ is the number of edges in \mathcal{T} (we omit further details as they are straightforward).

For brevity, let us denote the summations $\sum_{\rho \in [0..\sigma]^k}$ and $\sum_{0 \leq c < \sigma}$ in this paragraph by \sum_ρ and \sum_c , respectively. For the sake of completeness, we show that the compression boosting technique allows to achieve the k th-order empirical entropy $H_k(\mathcal{T})$, which, unlike H_k^* , satisfies the inequality $\log \sigma \geq H_k(\mathcal{T})$, i.e., we are to prove that $\sum_c \sum_\rho (\log \binom{m_\rho}{n_{c,\rho}} + o(n_{c,\rho})) \leq mH_k(\mathcal{T}) + 1.443m + o(m)$ (but we do not discuss additional structures of Hon et al. required for navigation; see [21]). First, $\sum_c \sum_\rho (\log \binom{m_\rho}{n_{c,\rho}} + o(n_{c,\rho}))$ is upper bounded by $\sum_\rho \sum_c (n_{c,\rho} \log \frac{m_\rho}{n_{c,\rho}} + 1.443n_{c,\rho}) + o(m) = \sum_\rho \sum_c n_{c,\rho} \log \frac{m_\rho}{n_{c,\rho}} + 1.443m + o(m)$. Denote $n_\rho = \sum_c n_{c,\rho}$. Note that, by definition, we have $mH_k(\mathcal{T}) = \sum_\rho \sum_c n_{c,\rho} \log \frac{n_\rho}{n_{c,\rho}}$. Since the function $\log x$ is concave, we have $\log(x+d) \leq \log x + d(\log x)' = \log x + d \log e/x$, for any $x > 0$ and any real d such that $x+d > 0$. Hence, we deduce $\sum_\rho \sum_c n_{c,\rho} \log \frac{m_\rho}{n_{c,\rho}} = \sum_\rho \sum_c n_{c,\rho} \log(\frac{n_\rho}{n_{c,\rho}} + \frac{m_\rho - n_\rho}{n_{c,\rho}}) \leq \sum_\rho \sum_c n_{c,\rho} (\log \frac{n_\rho}{n_{c,\rho}} + \frac{m_\rho - n_\rho}{n_\rho} \log e) = \sum_\rho \sum_c n_{c,\rho} \log \frac{n_\rho}{n_{c,\rho}} + \sum_\rho (m_\rho - n_\rho) \log e = mH_k(\mathcal{T}) + \log e$; the equality $\sum_\rho (m_\rho - n_\rho) \log e = \log e$ holds since $\sum_\rho m_\rho = m+1$ and $\sum_\rho n_\rho = m$. Finally, we hide the constant $\log e$ under $o(m)$ and obtain the k th-order entropy compression: $\sum_c \sum_\rho (\log \binom{m_\rho}{n_{c,\rho}} + o(n_{c,\rho})) \leq mH_k(\mathcal{T}) + 1.443m + o(m)$.

Fixed block compression boosting. The described partition lacks uniformity and requires relatively complex auxiliary data structures in order to support navigation and queries. Hon et al. [21] indeed organize such an infrastructure using $o(m)$ bits, provided $\sigma \leq m^\delta$ for a constant $\delta < 1$ (the condition $0 \leq k \leq \alpha \log_\sigma m - 1$ plays its role in this part). But it turns out that we can considerably simplify their whole construction using the fixed block boosting by Kärkkäinen and Puglisi [24]. The approach relies on the following lemma.

► **Lemma 5** (see [24, Lemma 4]). *Let $s = s_1 s_2 \dots s_\ell$ be an arbitrary partition of a string s into ℓ substrings. Let $s = s'_1 s'_2 \dots s'_t$ be a different partition of s into t substrings each of which has length at most b . Then, we have $\sum_{i=1}^t |s'_i| H_0(s'_i) \leq \sum_{i=1}^\ell |s_i| H_0(s_i) + (\ell - 1)b$.*

We encode the bit array B , which represents the transitions next, as follows.

► **Lemma 6.** *Provided $\sigma \leq m^\delta$ for a constant $\delta < 1$, B has an encoding that supports select and partial rank in $O(1)$ time and occupies $mH_k(\mathcal{T}) + 1.443m + o(m)$ bits simultaneously for all $k \in [0.. \max\{0, \alpha \log_\sigma m - 2\}]$, where $\alpha \in (0, 1)$ is an arbitrary fixed constant.*

Proof. We first discuss a fixed block encoding of the arrays B_c and prove, by means of Lemma 5, that it achieves the k th-order entropy compression. Then, we describe auxiliary structures that occupy only $o(m)$ bits and are used for queries and navigation in the blocks.

We partition each array $B_c[0..m]$ into $t = \lceil \frac{m+1}{b} \rceil$ blocks of length $b = \sigma \lceil \log^2 m \rceil$ (the last block can be shorter), encode each block using Lemma 2, and concatenate the encodings. Thus, we consume $\sum_{i=1}^t \sum_{0 \leq c < \sigma} (\log \binom{b_i}{n_{c,i}} + o(n_{c,i}))$ bits, where $n_{c,i}$ is the number of ones in the block $B_c[(i-1)b.. \min\{ib-1, m\}]$ and b_i is the length of the i th block (so that $b_i = b$, for $i \in [0..t)$, and $b_t \leq b$). Since $\sum_{i=1}^t \log \binom{b_i}{n_{c,i}} \leq \log \binom{m+1}{n_c}$ for $n_c = \sum_{i=1}^t n_{c,i}$, the result trivially holds for $k = 0$. The sum is upper bounded by $\sum_{i=1}^t \sum_{0 \leq c < \sigma} n_{c,i} \log \frac{b_i}{n_{c,i}} + 1.443m + o(m)$. Let us prove that $\sum_{i=1}^t \sum_{0 \leq c < \sigma} n_{c,i} \log \frac{b_i}{n_{c,i}} \leq mH_k(\mathcal{T}) + o(m)$ for all $k \in (0.. \alpha \log_\sigma m - 2]$.

Fix $k \in (0.. \alpha \log_\sigma m - 2]$. Denote $n_i = \sum_{0 \leq c < \sigma} n_{c,i}$. Using the inequality $\log(x+d) \leq \log x + d \log e/x$, we deduce the following upper bound:

$$\sum_{\substack{1 \leq i \leq t \\ 0 \leq c < \sigma}} n_{c,i} \log \frac{b_i}{n_{c,i}} = \sum_{\substack{1 \leq i \leq t \\ 0 \leq c < \sigma}} n_{c,i} \log \left(\frac{n_i}{n_{c,i}} + \frac{b_i - n_i}{n_{c,i}} \right) \leq \sum_{\substack{1 \leq i \leq t \\ 0 \leq c < \sigma}} n_{c,i} \log \frac{n_i}{n_{c,i}} + \sum_{i=1}^t (b_i - n_i) \log e.$$

First, we have $\sum_{i=1}^t (b_i - n_i) \log e = \log e$ since $\sum_{i=1}^t b_i = m + 1$ and $\sum_{i=1}^t n_i = m$. Second, $\sum_{i=1}^t \sum_{0 \leq c < \sigma} n_{c,i} \log \frac{n_i}{n_{c,i}} = \sum_{i=1}^t |s'_i| H_0(s'_i)$, where s'_i is a string of length n_i formed by concatenating the letters on the edges (u, v) such that $\text{num}(u)$ is inside the i th block, i.e., $\text{num}(u) \in [(i-1)b .. (i-1)b + b_i)$. For $\rho \in [0.. \sigma]^k$, let $[\ell_\rho.. r_\rho]$ be the set of all i such that ρ is a suffix of $\$^k \text{istr}(i)$, and let s_ρ be a string formed by concatenating the letters on the edges (u, v) such that $\text{num}(u) \in [\ell_\rho.. r_\rho]$. By definition, $\sum_{\rho \in [0.. \sigma]^k} |s_\rho| H_0(s_\rho) = mH_k(\mathcal{T})$. Note that at most $\ell = \sum_{i=0}^k \sigma^i = \frac{\sigma}{\sigma-1} (\sigma^k - \frac{1}{\sigma})$ strings s_ρ are nonempty and $\ell \leq 2\sigma^k$ since $\sigma \geq 2$. Let $\rho_1, \rho_2, \dots, \rho_{(\sigma+1)^k}$ be an ordering of all strings $\rho \in [0.. \sigma]^k$ such that $\ell_{\rho_1} \leq \dots \leq \ell_{\rho_{(\sigma+1)^k}}$. The definitions of s'_i and s_ρ imply that the letters in s'_i and s_ρ can be arranged so that $s'_1 s'_2 \dots s'_t = s_{\rho_1} s_{\rho_2} \dots s_{\rho_{(\sigma+1)^k}}$. Therefore, by Lemma 5, we obtain the next inequality:

$$\sum_{i=1}^t |s'_i| H_0(s'_i) \leq \sum_{\rho \in [0.. \sigma]^k} |s_\rho| H_0(s_\rho) + (\ell - 1) \max_{i \in [1..t]} n_i \leq mH_k(\mathcal{T}) + 2\sigma^{k+1}b.$$

As $k \leq \alpha \log_\sigma m - 2$, we have $\sigma^{k+1}b = \sigma^{k+2} \lceil \log^2 m \rceil \leq m^\alpha \lceil \log^2 m \rceil = o(m)$.

It remains to describe the auxiliary data structures that help to answer select and partial rank queries on the (now virtual) array B . First, we store σt pointers to the data structures encoding the blocks $B_c[(i-1)b.. \min\{ib-1, m\}]$, for $c \in [0.. \sigma)$ and $i \in [1..t]$. For rank, we create an array of length σt that stores the number of ones in the subarrays $B[0..(m+1)c + ib - 1]$, for $i \in [0..t)$ and $c \in [0.. \sigma)$. All this takes $O(\sigma t \log m) = O(\frac{m}{\log^2 m} \log m) = o(m)$ bits. For select, we create a bit array S formed by concatenating unary encodings for the number of ones in the blocks: e.g., if the first four blocks (of all σt blocks) contain, resp., 3, 2, 0, and 2 ones, then $S = 11101100110 \dots$; S is encoded using Lemma 2 and, thus, occupies $\log \binom{m+\sigma t}{m} + o(m) = \log \binom{m+\sigma t}{\sigma t} + o(m) \leq O(\sigma t \log m) + o(m) = o(m)$ bits. Using these structures, one can straightforwardly perform select and partial rank on B in $O(1)$ time. ◀

5 Main Data Structure

The encoding of failure links imposes a $2m$ -bit overhead, which, for small alphabet or highly compressible data, might be comparable to the space required for other structures. In this section, we describe a different encoding that uses only $\varepsilon m + o(m)$ bits, for any $\varepsilon \in (0, 2)$.

The key idea is to store the failure links only for some trie vertices. We call a subset W of the vertices of a tree a t -dense subset if each vertex $v \notin W$ has an ancestor $p \in W$ located at a distance less than t edges from v . (Note that the definition implies that W must contain the root.) Now we can formulate the main lemma.

► **Lemma 7.** *Suppose that $\text{failure}(v)$ can be calculated in $O(1)$ time only for $v \in W$, where W is a vertex set that is t -dense in the tree \mathcal{T} ; then there is a modification of the Aho–Corasick algorithm that uses the links $\text{failure}(v)$ only for $v \in W$ and processes any text T in $O(t|T| + \text{occ})$ time, where occ is the number of occurrences of the patterns in T .*

Proof. Our algorithm essentially simulates the Aho–Corasick solution but in the case when the usual algorithm calculates $\text{failure}(v)$ for $v \notin W$, the new one instead finds the nearest ancestor $p \in W$ of v , “backtracks” the input string T accordingly, then computes $\text{failure}(p)$, and continues the execution from this point (if p is the root and, thus, $\text{failure}(p)$ is undefined, we simply skip one letter and continue). The pseudocode is as follows (the omitted code reporting patterns in line 2 simply traverses report links and checks whether $\text{mark}(v) = 1$):

```

1: function  $\text{auto}(v, i, i_{\max})$ 
2:   if  $i > i_{\max}$  then  $i_{\max} \leftarrow i$  and report all patterns ending at position  $i - 1$ ;
3:   if  $i = |T|$  then return  $v$ ;
4:   if  $\text{next}(v, T[i]) \neq \text{nil}$  then return  $\text{auto}(\text{next}(v, T[i]), i + 1, i_{\max})$ ;
5:   for ( $p \leftarrow v$ ;  $p \notin W$ ;  $i \leftarrow i - 1$ ) do                                ▷ C-style loop with three parameters
6:     ( $p, c$ )  $\leftarrow$   $\text{parent\_edge}(v)$ ;                                       ▷  $p$  and  $c$  are such that  $v = \text{next}(p, c)$ 
7:   if  $p$  is root then return  $\text{auto}(\text{root}, i + 1, i_{\max})$ ;
8:   return  $\text{auto}(\text{failure}(p), i, i_{\max})$ ;

```

The execution starts with $\text{auto}(\text{root}, 0, 0)$. The function $\text{parent_edge}(v)$ returns the parent p of v and the letter c such that $v = \text{next}(p, c)$ (note that we only use p); in [3] it was shown that $\text{num}(p) = x \bmod (m + 1)$ and $c = \lfloor x / (m + 1) \rfloor$, where $x = \text{select}(\text{num}(v), \mathbf{B})$ and \mathbf{B} is the bit array that encodes the next transitions (see above). To prove the correctness, let us show by induction on the length of T that $\text{auto}(\text{root}, 0, 0)$ returns the vertex v_T such that $\text{str}(v_T)$ is the longest suffix of T that can be spelled out on a root-vertex path in \mathcal{T} .

The base $|T| = 0$ is trivial. Suppose that the claim holds for all lengths smaller than $|T|$. We are to show that $\text{auto}(\text{root}, 0, 0) = v_T$. By the inductive hypothesis, when $\text{auto}(v, i, i_{\max})$ is called with $i = |T| - 1$ for the first time, the string $\text{str}(v)$ is the longest suffix of $T[0..|T| - 2]$ that can be read on a root-vertex path of \mathcal{T} . Therefore, if $\text{next}(v, T[|T| - 1]) \neq \text{nil}$, $\text{next}(v, T[|T| - 1])$ obviously is equal to v_T and we return it in line 4. Now suppose that $\text{next}(v, T[|T| - 1]) = \text{nil}$. In lines 5–6 we find the nearest ancestor p of v belonging to W or put $p = v$ if $v \in W$, and backtrack accordingly to $T[0..i]$, for $i \in [0..|T|]$, such that $\text{str}(p)$ is a suffix of $T[0..i - 1]$. Observe the following claims: (i) for any vertex w such that $\text{str}(w)$ is a suffix of $T[0..i - 1]$, the result of $\text{auto}(w, i, i_{\max})$ is the same as the result of $\text{auto}(\text{root}, 0, 0)$ with $T := T[i - |\text{str}(w)| .. |T| - 1]$; (ii) $\text{str}(v_T)$ is either the empty string or a proper suffix of $\text{str}(v)$ concatenated with $T[|T| - 1]$. When p is not the root, the claim (ii) and the fact that $\text{str}(\text{failure}(p))$ is the longest proper suffix of $\text{str}(p)$ present in the trie imply that $\text{str}(v_T)$ is a suffix of $T[i - |\text{str}(\text{failure}(p))| .. |T| - 1]$. Then, by the claim (i) and the inductive hypothesis, the recursion $\text{auto}(\text{failure}(p), i, i_{\max})$ in line 8 returns v_T . When p is the root, (i) and (ii) analogously imply that the call to $\text{auto}(\text{root}, i + 1, i_{\max})$ in line 7 returns v_T .

Let us analyze the time complexity. The algorithm maintains two indices: i and $k = i - |\text{str}(v)|$. Each call to $\text{auto}(v, i, i_{\max})$ with $i \neq |T|$ either increases i in line 4 or increases k in lines 7 or 8. Since W is a t -dense subset, the loop 5–6 can decrease i by at most t before increasing k . Therefore, i can be decreased by at most $t|T|$ in total and, hence, the running time of the whole algorithm is $O(t|T|)$ plus $O(\text{occ})$ time to report pattern occurrences.


The presented solution explicitly stores T (or at least its last m letters) in order to support “backtracking” during the calculations. We, however, cannot afford to allocate the $m \log \sigma$ bits for T and desire to fit the additional space within an $o(m)$ bound. To this end, we maintain only a substring $T[i..i']$ such that $i' - i \leq 2\sqrt{m}$. While decreasing i in the loop 5–6, we grow this substring to the left using the letters c returned by $\text{parent_edge}(v)$. Once $i' - i$ becomes larger than $2\sqrt{m}$, we simply decrement i' . Once i becomes larger than i' and $i' < i_{\max}$, we must somehow restore the letters $T[i], T[i+1], \dots$. Denote by P the set of all positions $j \in [0..i_{\max}]$ such that j is a multiple of $\lceil \sqrt{m} \rceil$. For each $j \in P \cup \{i_{\max}\}$, we store a vertex v_j that was the vertex v in the function auto when we reached the position j for the first time (so that $\text{str}(v_j)$ is a suffix of $T[0..j-1]$). Since the loop 5–6 cannot make i smaller than $i - |\text{str}(v)|$, one can easily show that we always have $i \geq j - |\text{str}(v_j)|$ for each $j \in P \cup \{i_{\max}\}$. Thus, once $i > i'$, we compute in $O(1)$ time the position $j = \min\{\{i_{\max}\} \cup \{j \in P : j \geq i + \sqrt{m}\}\}$, then put $i' = j - 1$, and restore the string $T[i..i']$ in $O(\sqrt{m})$ time iteratively applying the function parent_edge to the vertex v_j .

Since $i_{\max} - i$ cannot exceed m , it suffices to store v_j only for the $\lceil \sqrt{m} \rceil + 1$ largest positions from P . One can maintain these v_j in a straightforward way using a deque on circular array of length $O(\sqrt{m})$, so that access to arbitrary v_j takes $O(1)$ time. Therefore, the additional space used is $O(\sqrt{m} \log m) = o(m)$ bits. By standard arguments, one can show that the time $O(\sqrt{m})$ required to restore $T[i..j-1]$ is amortized among at least \sqrt{m} increments of i that were performed to make $i > i'$. Thus, the total running time is $O(t|T|)$ as in the version that stores T explicitly. ◀

To perform $\text{failure}(v)$, for $v \in W$, and to check whether $v \in W$, we use the next lemma.

► **Lemma 8.** *Let W be a subset of vertices of a rooted tree \mathcal{F} with m edges. There is an encoding of \mathcal{F} that occupies $2|W| \log \frac{m+1}{|W|} + O(|W|) + o(m)$ bits and, for any vertex v , allows to determine whether $v \in W$ in $O(1)$ time and to compute the parent of v if $v \in W$ in $O(1)$ time, provided all the vertices are represented by their numbers in a depth first traversal of \mathcal{F} .*

Proof. Let $\text{num}(v)$ be a vertex numbering that corresponds to a depth first traversal of \mathcal{F} (the numbers are from the range $[0..m]$). To check whether $v \in W$, we create a bit array A of length $m + 1$ such that, for each vertex v , we have $A[\text{num}(v)] = 1$ iff $v \in W$. The array A is encoded in $\log \binom{m+1}{|W|} + o(|W|) = |W| \log \frac{m+1}{|W|} + O(|W|)$ bits in the data structure of Lemma 2 that supports access to $A[i]$ in $O(1)$ time using partial rank queries.



■ **Figure 2** The transformation of Lemma 8 for the tree of failure links from Figure 1: the left tree is original, the right tree is transformed; vertices from W are gray, important vertices are bold.

Let us now describe a transformation of the tree \mathcal{F} that preserves the function $\text{parent}(v)$ for $v \in W$ (see Figure 2). We call a vertex v *important* if either v is the root or v has a child from W . In the transformed tree, for each vertex u , the parent of u is the nearest

13:10 Compressed Multiple Pattern Matching

important ancestor of u in \mathcal{F} . Obviously, there are at most $|W| + 1$ internal vertices in this defined tree. Further, it is easy to see that the vertex numbering num corresponds to a depth first traversal of the transformed tree. Therefore, we can encode the new tree in $|W| \log \frac{m+1}{|W|} + O(|W|) + o(m)$ bits using the data structure of Lemma 4; for $v \in W$, the query $\text{parent}(v)$ on this structure returns the number num of the parent of v in the tree \mathcal{F} . \blacktriangleleft

Combining Lemmas 7 and 8, we prove the main theorem.

► Theorem 9. *Let $\varepsilon \in (0, 2)$ be an arbitrary constant and let \mathcal{D} be a set of d patterns over the alphabet $[0..\sigma]$ such that $\sigma \leq m^\delta$, for some constant $\delta < 1$, where m is the number of edges in the trie \mathcal{T} containing \mathcal{D} . Then, there is a data structure that allows to find all occ occurrences of the patterns in any text T in $O(|T| + \text{occ})$ time and occupies $mH_k(\mathcal{T}) + 1.443m + \varepsilon m + O(d \log \frac{m}{d})$ bits simultaneously for all $k \in [0..\max\{0, \alpha \log_\sigma m - 2\}]$, where $\alpha \in (0, 1)$ is an arbitrary constant.*

Proof. Let us construct a small t -dense subset for the tree \mathcal{T} . Observe that the set W_i that consists of the root and all vertices of \mathcal{T} with height $h \equiv i \pmod{t}$ is a t -dense subset. Obviously, there exists $j \in [0..t)$ such that $|W_j| \leq \lceil \frac{m+1}{t} \rceil$. We apply Lemma 8 to the tree \mathcal{F} of failure links and the subset W_j , provided all vertices are represented by the numbering num defined in Section 3 (in [3] it was shown that this numbering corresponds to a depth first traversal of \mathcal{F}). By Lemma 7, this allows us to solve the dictionary matching problem in $O(t|T| + \text{occ})$ time with only $2\lceil \frac{m+1}{t} \rceil \log t + O(\frac{m}{t}) + o(m)$ bits used for the failure links, which is upper bounded by $m \frac{c \log t}{t}$, for an appropriate constant $c > 0$. We add to the failure links the data structures for **mark**, **next**, and **report** described in Sections 3 and 4 (Lemma 6), which consume $mH_k(\mathcal{T}) + 1.443m + o(m) + O(d \log \frac{m}{d})$ bits (simultaneously for all $k \in [0..\max\{0, \alpha \log_\sigma m - 2\}]$), and choose t in such a way that $\frac{c \log t}{t} \leq \varepsilon/2$, so that the failure links take only $\varepsilon m/2$ bits. Solving the equation, we obtain $t = \Theta(\varepsilon^{-1} \log \varepsilon^{-1})$. Since ε is constant, the additive terms $o(m)$ in space can be upper bounded by $\varepsilon m/2$ and t in $O(t|T| + \text{occ})$ can be hidden under the big-O, so that the total space is $mH_k(\mathcal{T}) + 1.443m + \varepsilon m + O(d \log \frac{m}{d})$ bits and the processing time is $O(|T| + \text{occ})$. \blacktriangleleft

Note that, as it follows from the proof, the big-O notation in Theorem 9 hides a slowdown of the processing time to $O(|T|\varepsilon^{-1} \log \varepsilon^{-1} + \text{occ})$, which is the price of the space improvement.

Let us now show that our solution is, in a sense, close to space optimal when $d = o(m)$. It is easy to see that any data structure solving the multiple pattern matching problem implicitly encodes the trie \mathcal{T} containing the dictionary \mathcal{D} of patterns: if the data structure is a black box, then we can enumerate all possible strings of length $\leq m$ over the alphabet $[0..\sigma]$ and check which of them are recognized by the black box, thus finding all the patterns from \mathcal{D} . It is known that the number of tries with m edges over an alphabet of size σ is at least $\frac{1}{\sigma(m+1)+1} \binom{\sigma(m+1)}{m+1}$ (see [19, eq. 7.66] and [6, Thm 2.4]). Note that $\binom{\sigma(m+1)}{m+1} \geq \binom{\sigma(m+1)}{m}$, for $\sigma \geq 2$. Therefore, $\log(\frac{1}{\sigma(m+1)+1} \binom{\sigma(m+1)}{m+1}) \geq \log(\binom{\sigma(m+1)}{m}) - O(\log(\sigma m))$ is a lower bound for the worst-case space consumption of any solution for the multiple pattern matching.

► Theorem 10. *For any constant $\varepsilon \in (0, 2)$ and any set \mathcal{D} of d patterns over the alphabet $[0..\sigma]$ such that $d = o(m)$, where m is the number of edges in the trie containing \mathcal{D} , there is a data structure that allows to find all occ occurrences of the patterns in text T in $O(|T| + \text{occ})$ time and occupies $L + \varepsilon m$ bits of space, where $L = \log(\binom{\sigma(m+1)}{m}) - O(\log(\sigma m))$ is a lower bound on the worst-case space consumption for any such data structure.*

Proof. Since $d = o(m)$, we have $d = \frac{m}{f(m)}$, where $f(m) \xrightarrow{m \rightarrow \infty} +\infty$, and therefore $d \log \frac{m}{d} = \frac{m}{f(m)} \log f(m) = o(m)$. Thus, applying Theorem 9 for $\frac{\varepsilon}{2}$, we obtain a data structure occupying $mH_k + 1.443m + \frac{\varepsilon}{2}m + o(m)$ bits. Further, applying the simple encoding from Lemma 2 for the bit array B representing the transitions next, we obtain a solution occupying $\log \binom{\sigma(m+1)}{m} + \frac{\varepsilon}{2}m + o(m)$ bits. Since ε is constant, we have $o(m) \leq \frac{\varepsilon}{2}m$; hence, the result follows. ◀

References

- 1 Alfred V. Aho and Margaret J. Corasick. Efficient string matching: an aid to bibliographic search. *Communications of the ACM*, 18(6):333–340, 1975. doi:10.1145/360825.360855.
- 2 Jarno Alanko and Tuukka Norri. Greedy shortest common superstring approximation in compact space. In *Proc. SPIRE*, volume 10508 of *LNCS*, pages 1–13. Springer, 2017. doi:10.1007/978-3-319-67428-5_1.
- 3 Djamal Belazzougui. Succinct dictionary matching with no slowdown. In *Proc. CPM*, volume 6129 of *LNCS*, pages 88–100. Springer, 2010. doi:10.1007/978-3-642-13509-5_9.
- 4 Michael Burrows and David J. Wheeler. A block-sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation, Palo Alto, California, 1994.
- 5 Ho-Leung Chan, Wing-Kai Hon, Tak-Wah Lam, and Kunihiko Sadakane. Dynamic dictionary matching and compressed suffix trees. In *Proc. SODA*, pages 13–22. SIAM, 2005.
- 6 David Clark. *Compact pat trees*. PhD thesis, University of Waterloo, 1997.
- 7 Raphaël Clifford, Allyx Fontaine, Ely Porat, Benjamin Sach, and Tatiana Starikovskaya. Dictionary matching in a stream. In *Proc. ESA*, volume 9294 of *LNCS*, pages 361–372. Springer, 2015. doi:10.1007/978-3-662-48350-3_31.
- 8 Maxime Crochemore and Wojciech Rytter. *Jewels of stringology*. World Scientific Publishing Co. Pte. Ltd., 2002.
- 9 Vassilis Dimopoulos, Ioannis Papaefstathiou, and Dionisios Pnevmatikatos. A memory-efficient reconfigurable Aho–Corasick FSM implementation for intrusion detection systems. In *Proc. IC-SAMOS*, pages 186–193. IEEE, 2007. doi:10.1109/IC-SAMOS.2007.4285750.
- 10 Guy Feigenblat, Ely Porat, and Ariel Shiftan. Linear time succinct indexable dictionary construction with applications. In *Proc. DCC, 2016*, pages 13–22. IEEE, 2016. doi:10.1109/DCC.2016.70.
- 11 Guy Feigenblat, Ely Porat, and Ariel Shiftan. A grouping approach for succinct dynamic dictionary matching. *Algorithmica*, 77(1):134–150, 2017. doi:10.1007/s00453-015-0056-0.
- 12 Paolo Ferragina, Raffaele Giancarlo, Giovanni Manzini, and Marinella Sciortino. Boosting textual compression in optimal linear time. *Journal of the ACM*, 52(4):688–713, 2005. doi:10.1145/1082036.1082043.
- 13 Paolo Ferragina, Fabrizio Luccio, Giovanni Manzini, and S. Muthukrishnan. Structuring labeled trees for optimal succinctness, and beyond. In *Proc. FOCS*, pages 184–193. IEEE, 2005. doi:10.1109/SFCS.2005.69.
- 14 Paolo Ferragina and Giovanni Manzini. Opportunistic data structures with applications. In *Proc. FOCS*, pages 390–398. IEEE, 2000. doi:10.1109/SFCS.2000.892127.
- 15 Travis Gagie. Large alphabets and incompressibility. *Information Processing Letters*, 99(6):246–251, 2006. doi:10.1016/j.ipl.2006.04.008.
- 16 Simon Gog. SDSL - succinct data structure library, 2015. URL: <https://github.com/simongog/sdsl-lite>.
- 17 Simon Gog, Juha Kärkkäinen, Dominik Kempa, Matthias Petri, and Simon J. Puglisi. Faster, minuter. In *Proc. DCC, 2016*, pages 53–62. IEEE, 2016. doi:10.1109/DCC.2016.94.
- 18 Shay Golan and Ely Porat. Real-time streaming multi-pattern search for constant alphabet. In *Proc. ESA*, volume 87 of *LIPICs*, pages 41:1–41:15. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017. doi:10.4230/LIPICs.ESA.2017.41.
- 19 Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics*. Massachusetts: Addison-Wesley, 1989.

- 20 Dan Gusfield. *Algorithms on strings, trees and sequences: computer science and computational biology*. Cambridge university press, 1997. doi:10.1017/CB09780511574931.
- 21 Wing-Kai Hon, Tsung-Han Ku, Rahul Shah, Sharma V. Thankachan, and Jeffrey S. Vitter. Faster compressed dictionary matching. *Theoretical Computer Science*, 475:113–119, 2013. doi:10.1016/j.tcs.2012.10.050.
- 22 Wing-Kai Hon, Tak-Wah Lam, Rahul Shah, Siu-Lung Tam, and Jeffrey S. Vitter. Compressed index for dictionary matching. In *Proc. DCC*, pages 23–32. IEEE, 2008. doi:10.1109/DCC.2008.62.
- 23 Jesper Jansson, Kunihiko Sadakane, and Wing-Kin Sung. Ultra-succinct representation of ordered trees with applications. *Journal of Computer and System Sciences*, 78(2):619–631, 2012. doi:10.1016/j.jcss.2011.09.002.
- 24 Juha Kärkkäinen and Simon J. Puglisi. Fixed block compression boosting in FM-indexes. In *Proc. SPIRE*, volume 7024 of *LNCS*, pages 174–184. Springer, 2011. doi:10.1007/978-3-642-24583-1_18.
- 25 Tsvi Kopelowitz, Ely Porat, and Yaron Rozen. Succinct Online Dictionary Matching with Improved Worst-Case Guarantees. In *Proc. CPM*, volume 54 of *LIPICs*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2016. doi:10.4230/LIPICs.CPM.2016.6.
- 26 S. Rao Kosaraju and Giovanni Manzini. Compression of low entropy strings with Lempel–Ziv algorithms. *SIAM Journal on Computing*, 29(3):893–911, 1999. doi:10.1137/S0097539797331105.
- 27 Hung-Jen Liao, Chun-Hung R. Lin, Ying-Chih Lin, and Kuang-Yuan Tung. Intrusion detection system: A comprehensive review. *Journal of Network and Computer Applications*, 36(1):16–24, 2013. doi:10.1016/j.jnca.2012.09.004.
- 28 Veli Mäkinen and Gonzalo Navarro. Compressed full-text indexes. *ACM Computing Surveys*, 39(1):2, 2007. doi:10.1145/1216370.1216372.
- 29 Giovanni Manzini. An analysis of the Burrows–Wheeler transform. *Journal of the ACM*, 48(3):407–430, 2001. doi:10.1145/382780.382782.
- 30 Gonzalo Navarro. *Compact data structures: A practical approach*. Cambridge University Press, 2016. doi:10.1017/CB09781316588284.
- 31 Gonzalo Navarro and Kunihiko Sadakane. Fully functional static and dynamic succinct trees. *ACM Transactions on Algorithms*, 10(3):16, 2014. doi:10.1145/2601073.
- 32 Derek Pao, Xing Wang, Xiaoran Wang, Cong Cao, and Yuesheng Zhu. String searching engine for virus scanning. *IEEE Transactions on Computers*, 60(11):1596–1609, 2011. doi:10.1109/TC.2010.250.
- 33 Rajeev Raman, Venkatesh Raman, and S. Srinivasa Rao. Succinct indexable dictionaries with applications to encoding k-ary trees and multisets. *ACM Transactions on Algorithms*, 3(4):43, 2007. doi:10.1145/1290672.1290680.
- 34 Dina Sokol and Marcus Shoshana. Engineering small space dictionary matching. *arXiv preprint*, 2013. arXiv:1301.6428.
- 35 Alan Tam, Edward Wu, Tak-Wah Lam, and Siu-Ming Yiu. Succinct text indexing with wildcards. In *Proc. SPIRE*, pages 39–50. Springer, 2009. doi:10.1007/978-3-642-03784-9_5.
- 36 Sun Wu and Udi Manber. Fast text searching: allowing errors. *Communications of the ACM*, 35(10):83–91, 1992. doi:10.1145/135239.135244.

A Implementation and Experiments

We implemented the data structure described in this paper in C++ and compared its runtime and memory consumption with the Belazzougui’s solution [3] and a naive algorithm. We could not find implementations of the Belazzougui’s data structure and implemented it too. In [34] we found a different data structure based on compressed suffix trees but it showed a very poor performance, so we decided to exclude it from the tests.

The experiments were performed on a machine equipped with six 1.8 GHz Intel Xeon E5-2650L v3 CPUs with 30 MiB L3 cache and 16 GiB of RAM. The OS was Ubuntu 16.04.3 LTS, 64bit running kernel 4.4.0. All programs were compiled using `g++ v5.4.0` with `-O3 -march=x86-64` options. The source codes of all tested algorithms are available at <https://bitbucket.org/umqra/multiple-pattern-matching>. At the same URL one can find the 6 texts and 7 dictionaries on which the experiments were run.

The texts are as follows (see also Table 1):

- `program`: a sample binary file generated by a simple algorithm (see the URL above);
- `chr1.dna`: Human first chromosome genome sequence in FASTA format¹;
- `wiki.en/wiki.ja/wiki.ru/wiki.zh`: the first 200Mb of the dump of all English/Japanese/Russian/Chinese wikipedia articles².

The dictionaries are as follows (see also Table 1):

- `dna.dict`: reads generated for `chr1.dna` by `wgsim` simulator³ with read length ~ 100 ;
- `urls.dict`: URLs from the `.su/.nu` zones⁴ and the Alexa database of popular URLs⁵;
- `virus.dict`: virus signatures from the `main.cvd` file of the ClamAV database⁶;
- `t1.en.dict/t1.ja.dict/t1.ru.dict/t1.zh.dict`: the titles (in lower case) of some English/Japanese/Russian/Chinese wikipedia articles with length at least 3 letters⁷ (the dictionaries were truncated to reduce temporary space used in the index construction).

Our implementations use the SDSL library by Gog [16]. We tested the following algorithms:

- `blz`: the original Belazzougui's compressed data structure [3] in which the bitvectors $B_0, \dots, B_{\sigma-1}$ were implemented using `sd_vector` from the SDSL;
- `cblz`: our algorithm with fixed block compression boosting;
- `cblz8`: the same as `cblz` but with failure links sparsified using an 8-dense vertex subset;
- `smp`: a simple $O(nm)$ -time algorithm that uses only two components of the `blz` data structure: the bitvectors $B_0, \dots, B_{\sigma-1}$ and an array of length $m+1$ encoding the mark flags; they all are implemented using `sd_vector` from the SDSL.

For each algorithm, we ran 7 tests: in each test we search the patterns from a dictionary of Table 1 in the corresponding text from the same table row (note that `wiki.en` is used in two dictionaries). The results are present in Figure 3. The Aho–Corasick data structure required too much memory in our experiments, so we do not include it.

All dictionaries can be split into two unequal groups: dictionaries with long patterns (in our case it is only `dna.dict`) and short patterns (all other dictionaries). On `dna.dict` the fastest algorithms are `blz` and `cblz`; both `cblz8` and `smp` are about two times slower than them. Not surprisingly, `smp` is the fastest algorithm on short patterns and the algorithms `blz`, `cblz`, and `cblz8` are 2–3 times slower than `smp`. Because of some implementation details (we use a specially tailored version of `sd_vector`), `cblz` is faster than the simpler algorithm `blz` in our tests. We were unable to explain why `cblz8` often works faster than the simpler `cblz` algorithm. To sum up, the three `blz` algorithms have acceptable running times, but it really makes sense to use them only on relatively long dictionary patterns.

¹ ftp://ftp.ncbi.nih.gov/genomes/H_sapiens/Assembled_chromosomes/seq/hs_alt_CHM1_1.1_chr1.fa.gz

² <https://dumps.wikimedia.org/>

³ <https://github.com/lh3/wgsim>

⁴ <https://zonedata.iis.se/>

⁵ <http://s3.amazonaws.com/alexa-static/top-1m.csv.zip>

⁶ <https://www.clamav.net/downloads>


⁷ <https://dumps.wikimedia.org/>

13:14 Compressed Multiple Pattern Matching

■ **Table 1** Statistics of the test dictionaries of patterns (number of patterns, average pattern length, m , σ) and texts in which the patterns were searched (length, number of found occurrences).

Dictionary	patterns	avg. len.	m	σ	Text	length	found occ.
dna.dict	1,999,911	100.0	178,323,409	5	chr1.dna	228,503,292	157,393
virus.dict	4,059,198	16.0	56,430,521	256	program	104,857,600	1,347,031
urls.dict	3,825,132	16.7	39,385,319	40	wiki.en	209,040,222	3,461,097
t1l.en.dict	3,875,263	18.3	32,255,913	2881	wiki.en	209,040,222	137,071,073
t1l.ja.dict	1,691,693	8.4	7,182,594	7329	wiki.ja	109,018,155	15,196,791
t1l.ru.dict	4,145,276	22.5	39,908,335	1350	wiki.ru	131,898,987	65,428,137
t1l.zh.dict	1,649,209	7.3	6,164,034	12113	wiki.zh	126,021,996	18,925,337

The central chart shows that in most cases the dictionaries are very well compressed (in comparison with $\frac{\log \sigma}{8}$ bytes per letter; see the shaded columns). But, as one can conclude from the low compression ratio of the corresponding trie on the right chart, the dictionary compression is mostly due to the assembling of the patterns in the trie (for instance, while the total length of the patterns $\{a^i b\}_{i=0}^{k-1}$ is $k(k+1)/2$, their trie contains only $m = 2k$ edges). In the right chart, one can observe that the effect of compression boosting in `cb1z` is hidden under the overhead imposed by additional structures, and, because of this, `cb1z` occupies about the same space as `blz`. The space optimized version `cb1z8` lowers the overhead and the resulting data structure occupies about the same space as `smp`; this, surely, is possible only because `cb1z8` is H_k -compressed while `smp` is only H_0 -compressed.



■ **Figure 3** Performance and space consumption; shaded columns correspond to $\frac{\log \sigma}{8}$.