

# Compressed Full-Text Indexes

Gonzalo Navarro

Dept. of Computer Science, University of Chile

Veli Mäkinen

Technische Fakultät, Universität Bielefeld

---

Full-text indexes provide fast search for patterns over large text collections. A serious problem of these indexes has traditionally been their space consumption. A recent trend is to develop indexes that exploit the compressibility of the text, so that their size is a function of the compressed text length. This concept has evolved into *self-indexes*, which in addition contain enough information to reproduce any text portion, so they *replace* the text. The exciting possibility of an index that takes space close to that of the compressed text, replaces it, and in addition provides fast search over it, has triggered a wealth of activity and produced surprising results in very short time, radically changing the status of this area in less than five years. The most successful indexes nowadays are able to obtain almost optimal space and search time simultaneously.

In this paper we present the main concepts underlying self-indexes. We explain the relationship between text entropy and regularities that show up in index structures, which permit compressing them. Then we cover the most relevant self-indexes up to date, focusing on the essential aspects on how they exploit the text compressibility and how they solve efficiently the different search problems. We aim at giving the theoretical background to understand and follow the developments in this area.

Categories and Subject Descriptors: F.2.2 [**Analysis of algorithms and problem complexity**]: Nonnumerical algorithms and problems—*Pattern matching, Computations on discrete structures, Sorting and searching*; H.2.1 [**Database management**]: Physical design—*Access methods*; H.3.2 [**Information storage and retrieval**]: Information storage—*File organization*; H.3.3 [**Information storage and retrieval**]: Information search and retrieval—*Search process*

General Terms: Algorithms

Additional Key Words and Phrases: Text indexing, text compression, entropy.

---

First author funded by Millennium Nucleus Center for Web Research, Grant P04-067-F, Mideplan, Chile. Second author funded by the Deutsche Forschungsgemeinschaft (BO 1910/1-3) within the Computer Science Action Program, and by the Academy of Finland.

Address: Gonzalo Navarro. Blanco Encalada 2120, Santiago, Chile. E-mail: [gnavarro@dcc.uchile.cl](mailto:gnavarro@dcc.uchile.cl). Web: <http://www.dcc.uchile.cl/~gnavarro>. Veli Mäkinen. Postfach 10031, 33510 Bielefeld, Germany. Email: [veli@cebitec.uni-bielefeld.de](mailto:veli@cebitec.uni-bielefeld.de). Web: <http://www.cebitec.uni-bielefeld.de/u/veli>.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works, requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept, ACM Inc., 1515 Broadway, New York, NY 10036 USA, fax +1 (212) 869-0481, or [permissions@acm.org](mailto:permissions@acm.org).

## 1. INTRODUCTION

The amount of digitally available information is growing at an exponential rate. The volume of new material published in the Web in the period 1999–2003 exceeds what has been produced during the whole previous history of Mankind. A large part of these data is formed by *text*, that is, sequences of symbols representing not only natural language, but also music, program code, signals, multimedia streams, biological sequences, time series, and so on. If we exclude strongly structured data such as relational tables, text is the medium to convey information where retrieval by content is best understood. The recent boom on XML advocates the use of text as the medium to express structured and semistructured data as well, boosting the potential of text as the favorite format for information storage, exchange, and retrieval.

Each scenario where text is used to express information requires a different form of retrieving such information from the text. On natural language text, for example, information retrieval and natural language processing techniques aim at discovering the meaning of text passages and their relevance to the information need of a human. The main tools are statistical analysis of occurrences and variants of words. Another prominent example is computational biology, which works on DNA, protein, or gene sequences to extract information on evolutionary history, biochemical function, chemical structure, and so on. The main tool used is sequence comparison to determine similarities according to diverse biological models.

There is a basic search task, however, that underlies all the applications that aim at extracting information from text. *String matching* is the process of finding the occurrences of a short string (called the *pattern*) inside a (usually much longer) string called the *text*. Virtually every text managing application builds on basic string matching to implement more sophisticated functionality such as finding frequent words or retrieving sequences most similar to a sample. Significant developments in basic string matching have a wide impact on most applications. In the rest of the paper we will refer to *text searching* to denote the more restricted task of string matching, which will be our focus.

Text searching can be carried out in two forms. *Sequential* text searching requires no preprocessing of the text, but rather traverses it sequentially to point out every occurrence of the pattern. *Indexed* text searching builds a data structure (*index*) on the text beforehand, which permits finding all the occurrences of any pattern without traversing the whole text. Indexing is the choice when (i) the text is so large that a sequential scanning is prohibitively costly, (ii) the text does not change so frequently that the cost of building and maintaining the index outweighs the savings on searches, and (iii) there is sufficient *storage space* to maintain the index and provide efficient access to it.

While the first two considerations refer to the convenience of indexing compared to sequentially scanning the text, the last one is a necessary condition to consider indexing at all. In a first sight, the storage issue might not seem significant given the wide availability of massive storage. The real problem, however, is efficient access. In the last two decades, the CPU speeds have been duplicating every 18 months, while the disk access times have stayed basically unchanged. Computer caches are many times faster than their standard main memories. On the other

hand, if we exclude some simple applications on natural language texts that can be addressed with a lightweight index called *inverted index*, the classical indexes for text searching require from 4 to 20 times the text size. This means that, even when we may have enough main memory to hold a text, we may need to use the disk to store the index. Moreover, most existing indexes are not designed to work in secondary memory, so using them from disk is extremely inefficient. As a result, indexes are usually confined to the case where the text is so small that even the index fits in main memory, and those cases are less interesting for indexing given consideration (i). For such a small text, a sequential scanning can be preferable for its simplicity and better cache usage compared to an index.

*Text compression* is a technique to represent a text using less space. Given the relation between main and secondary memory access times, it is advantageous to store a large text that does not fit in main memory in compressed form, so as to reduce disk transfer time, and then decompress it by chunks in main memory prior to sequential searching. Moreover, a text may fit in main memory once compressed, so compression may completely remove the need to access the disk. Some developments in recent years have focused on improving this even more by avoiding the decompression but directly searching the compressed text instead.

Several attempts to reduce the space requirements of text indexes were made in the past with moderate success, and some of them even considered the relation with text compression. Three important concepts have emerged.

**Definition 1** *A succinct index is an index that provides fast search functionality using a space proportional to that of the text itself (say, two times the text size). A stronger concept is that of a compressed index, which takes advantage of the regularities of the text to operate in space proportional to that of the compressed text (say, 3 times the zero-order entropy of the text). An even more powerful concept is that of a self-index, which is a compressed index that, in addition to providing search functionality, contains enough information to efficiently reproduce any text substring. A self-index can therefore replace the text.*

Classical indexes such as suffix trees and arrays are not succinct. On a text of  $n$  characters over an alphabet of size  $\sigma$ , those indexes require  $O(n \log n)$  bits of space, whereas the text requires  $n \log \sigma$  bits. The first succinct index we know of is by Kärkkäinen and Ukkonen [1996a], which uses Ziv-Lempel compression concepts to achieve  $O(n \log \sigma)$  bits of space. Indeed, theirs is a compressed index achieving space proportional to the  $k$ -th order entropy of the text. However, it was not until this decade that the first self-index appeared [Ferragina and Manzini 2000] and the potential of the relationship between text compression and text indexing was fully realized, in particular regarding the correspondence between the entropy of a text and the regularities in some widely used indexing data structures. Several others succinct and self-indexes appeared almost simultaneously [Mäkinen 2000; Grossi and Vitter 2000; Sadakane 2000]. The exciting concept of a self-index that requires space close to that of the compressed text, provides fast searching on it, and moreover replaces the text, has triggered much interest on this issue and produced surprising results in very few years.

At this point, there exist indexes that require space close to that of the best ex-

isting compression techniques, and provide search and text recovering functionality with almost optimal time complexity [Grossi et al. 2003; Ferragina et al. 2004b]. Yet, several of these developments are still in a theoretical stage, and significant algorithm engineering is necessary to implement them. Furthermore, the developments are very immature with respect to several other aspects such as construction time and space, management of secondary storage, searching for more complex patterns, updating upon text changes, etc. Overall, this is an extremely exciting research area, with encouraging results of theoretical and practical interest, and a long way ahead.

The aim of this survey is to give the theoretical background needed to understand and follow the developments in this area. We first explain the relationship between the compressibility of a text and the regularities that show up in its indexing structures. Those regularities can be exploited to compress the indexes. Then, we cover the most relevant self-indexes that have been proposed, focusing on the conceptual foundations underlying them, and explaining how they exploit the regularities of the indexes to compress them and at the same time search them efficiently.

We do not give experimental results in this paper. Doing this carefully and thoroughly would require many more pages. Some comparisons can be found, for example, by Mäkinen and Navarro [2005a]. All the implementations we refer to across the survey can be found from our Web pages. A repository of texts and standardized implementations of succinct full-text indexes will soon be available at `mirrors.pizzachili.dcc.uchile.cl` and `pizzachili.di.unipi.it`.

## 2. NOTATION AND BASIC CONCEPTS

We introduce the notation we use in the rest of the paper. A *string*  $S$  is a sequence of *characters*. Each character is an element of a finite set called *alphabet*. The alphabet is usually called  $\Sigma$  and its size  $|\Sigma| = \sigma$ , and it is assumed to be totally ordered. Sometimes, for technical convenience, we assume  $\Sigma = [1, \sigma] = \{1, 2, \dots, \sigma\}$ . The *length* (number of characters) of  $S$  is denoted  $|S|$ . Let  $n = |S|$ , then the characters of  $S$  are indexed from 1 to  $n$ , so that the  $i$ -th character of sequence  $S$  is  $S_i$ . A *substring* of  $S$  is written  $S_{i,j} = S_i S_{i+1} \dots S_j$ . A *prefix* of  $S$  is a substring of the form  $S_{1,j}$ , and a *suffix* is a substring of the form  $S_{i,n}$ . If  $i > j$  then  $S_{i,j} = \varepsilon$ , the empty string of length  $|\varepsilon| = 0$ . The *concatenation* of strings  $S$  and  $S'$ , written  $SS'$ , is obtained by appending sequence  $S'$  at the end of  $S$ . It is also possible to concatenate string  $S$  and character  $c$ , as in  $cS$  or  $Sc$ . This is the same as if  $c$  were a string of length 1.

The *lexicographical order* “ $<$ ” among strings is defined as follows. Let  $a$  and  $b$  be characters and  $X$  and  $Y$  be strings. Then  $aX < bY$  if  $a < b$ , or if  $a = b$  and  $X < Y$ . On the other hand,  $\varepsilon < X$  for any  $X \neq \varepsilon$ .

The search problem we focus on in this paper is defined as follows.

**Definition 2** *Given a (long) text  $T_{1,n}$  and a (comparatively short) pattern  $P_{1,m}$ , both strings over alphabet  $\Sigma$ , find all the occurrences of  $P$  in  $T$ , that is, compute the set of positions where  $P$  occurs in  $T$ ,  $O = \{1 + |X|, \exists Y, T = XPY\}$ . Two search tasks are of interest: (1) count the number of occurrences, that is, return  $occ = |O|$ ; (2) locate the occurrences, that is, return set  $O$  in any order. When text*

$T$  is not explicitly available, a third task of interest is (3) display text substrings, that is, return  $T_{l,r}$  given  $l$  and  $r$ .

In our search problem, we adopt for technical convenience the assumption that  $T$  is terminated by  $T_n = \$$ , which is a character from  $\Sigma$  that is smaller than all the others and appears nowhere else in  $T$  nor in  $P$ .

Finally, logarithms in this paper are in base 2 unless otherwise stated.

### 3. BASIC TEXT INDEXES

In this section we cover the most popular full-text indexes, which are in particular those that are turned into compressed indexes later in this paper. We are interested in general indexes that work for any kind of text. This excludes, in particular, the popular inverted indexes, which can be used only when the text is composed of “natural language”, that is, syntactically distinguishable words following some statistical rules that govern the size and distribution of the vocabulary (set of different words). Inverted indexes can not be used, for example, for DNA, protein, or gene sequences; MIDI, audio, and other multimedia signals; Chinese, Japanese, and other Eastern languages; source and binary program code; numeric sequences of diverse kinds; and even texts written in languages where long words are formed from particles such as Finnish, German, and many others. There are good books on indexing and compressing “natural language” text [Bell et al. 1990; Witten et al. 1999; Baeza-Yates and Ribeiro 1999].

Given the focus of this paper, we are not covering either the different text indexes that have been designed with space reduction in mind, yet they have no relation to text compression nor they are self-indexes. In general these indexes have had some, but not spectacular, success in lowering the large space requirements of text indexes [Blumer et al. 1987; Andersson and Nilsson 1995; Kärkkäinen 1995; Irving 1995; Colussi and de Col 1996; Kärkkäinen and Ukkonen 1996b; Crochemore and V´erin 1997; Kurtz 1998; Giegerich et al. 1999].

#### 3.1 Tries or Digital Trees

A digital tree or *trie* [Fredkin 1960; Knuth 1973] is a data structure that stores a set of strings. It can search for a string in the set in time proportional to the length of the string sought, and independent on the set size.

**Definition 3** *A trie for a set of distinct strings  $S^1, S^2, \dots, S^N$  is a tree where each node corresponds to a distinct prefix in the set. The root node corresponds to the empty prefix  $\varepsilon$ . Node  $v$  representing prefix  $S^v$  is a child of node  $u$  representing prefix  $S^u$  iff  $S^v = S^u c$  for some character  $c$ , which will label the tree edge from  $u$  to  $v$ .*

We assume that all strings are terminated by “\$”. Under this assumption, no string  $S^i$  is a prefix of another, and thus the trie has exactly  $N$  leaves, each corresponding to a distinct string. Fig. 1 illustrates.

A trie for  $S^1, S^2, \dots, S^N$  is easily built in time  $O(|S^1| + |S^2| + \dots + |S^N|)$  by successive insertions. Any string  $S$  can be searched for in the trie in time  $O(|S|)$  by following from the trie root the path labeled with the characters of  $S$ . Two outcomes are possible: (i) at some point there is no edge labeled  $S_i$  to follow, which means




Fig. 1. A trie for the set {"alabar", "a", "la", "alabarda"}. In general trie nodes may have arity up to  $\sigma$ .

that  $S$  is not in the set, (ii) we reach a leaf corresponding to  $S$  (assume that  $S$  is also terminated with character "\$").

Actually, the above complexities assume that the alphabet size is a constant. For general  $\sigma$ , we must multiply the above complexities by  $O(\log \sigma)$ , which accounts for the overhead of searching the correct character to follow inside each node. This can be made  $O(1)$  by using at each node a table of size  $\sigma$ , but in this case the size and construction cost must be multiplied by  $O(\sigma)$  to allocate and initialize the tables at each node. Alternatively, perfect hashing permits  $O(1)$  search time  $O(1)$  space factor at each node, yet the construction cost is multiplied by  $O(\sigma^2)$  [Raman 1996].

Note that a trie can be used for *prefix searching*, that is, to find every string in the set prefixed by  $S$ . In this case,  $S$  is not terminated with "\$". If we can follow the trie path corresponding to  $S$ , then the internal node reached corresponds to all the strings  $S^i$  in the set prefixed by  $S$ . We can traverse all the leaves of the subtree to find the answers.

### 3.2 Suffix Tries and Suffix Trees

Let us now consider how tries can be used for text indexing. Given text  $T_{1,n}$  (terminated with  $T_n = \$$ ),  $T$  defines  $n$  suffixes  $T_{1,n}, T_{2,n}, \dots, T_{n,n}$ .

**Definition 4** *The suffix trie of a text  $T$  is a trie data structure built over all the suffixes of  $T$ .*

The suffix trie of  $T$  makes up an index for fast text searching. Given a pattern  $P_{1,m}$  (not terminated with "\$"), every occurrence of  $P$  in  $T$  is a substring of  $T$ , that is, the prefix of a suffix of  $T$ . Entering the suffix trie with the characters of  $P$  leads us to a node that corresponds to all the suffixes of  $T$  prefixed by  $P$  (or, if we do not arrive at any trie node, then  $P$  does not occur in  $T$ ). This permits counting the occurrences of  $P$  in  $T$  in  $O(m)$  time, by simply recording the number of leaves that descend from each suffix tree node. It also permits finding all the *occ* occurrences of  $P$  in  $T$  in  $O(m + occ)$  time by traversing the whole subtree (some additional pointers threading the leaves and connecting each internal node to its first leaf are necessary to ensure this complexity).

In practice, the trie is pruned at a node as soon as there is only a unary path from the node to a leaf. Instead, a pointer to the text position where the corresponding suffix starts is stored. In this case, it is possible that we arrive at a leaf during the search for  $P$ , without having read the complete suffix of  $T$ . This means that  $P$  has at most one occurrence in  $T$ . We have to follow the pointer and compare the rest of  $P$  with the rest of the suffix of  $T$  to determine whether there is one occurrence or none.

With or without the trick of pruning nodes, the trie of a text  $T_{1,n}$  requires space and construction time  $O(n^2)$  in the worst case. Albeit this worst case is unlikely, and on average the trie requires  $O(n)$  space [Sedgewick and Flajolet 1996], there exist equally powerful structures that guarantee this in the worst case [Morrison 1968; Apostolico 1985].

**Definition 5** *The suffix tree of a text  $T$  is a suffix trie where each unary path is converted into a single edge. Those edges are labeled by strings obtained by concatenating the characters of the replaced path. The leaves of the suffix tree indicate the text position where the corresponding suffixes start.*

Since there are  $n$  leaves and no unary nodes, it is easy to see that suffix trees require  $O(n)$  space (the strings at the edges are represented with pointers to the text). Moreover, they can be built in  $O(n)$  time [Weiner 1973; McCreight 1976; Ukkonen 1995]. Fig. 2 shows an example.

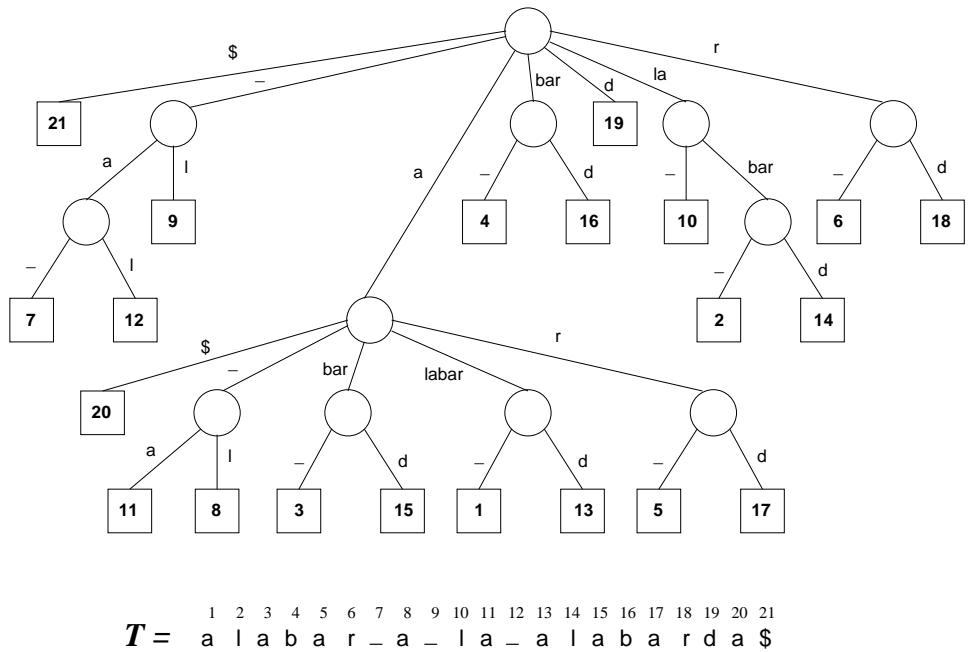


Fig. 2. The suffix tree of the text "alabar a la alabarda\$". The white space is written as an underscore for clarity, and it is lexicographically smaller than the characters "a"-"z".

The search for  $P$  in the suffix tree of  $T$  is similar to a trie search. Now we may use more than one character of  $P$  to traverse an edge, but all edges leaving from a node have different first characters. The search can finish in three possible ways: (i) at some point there is no edge leaving from the current node that matches the characters that follow in  $P$ , which means that  $P$  does not occur in  $T$ ; (ii) we read all the characters of  $P$  and end up at an internal node (or in the middle of an edge), in which case all the answers are in the subtree of the reached node (or edge); or (iii) we reach a leaf of the suffix tree without having read the whole  $P$ , in which case there is at most one occurrence of  $P$  in  $T$ , which must be checked by going to the text pointed to by the leaf and comparing the rest of  $P$  with the rest of the suffix. In any case the process takes  $O(m)$  time and suffices for counting queries.

Suffix trees permit  $O(m + occ)$  locating time without need of further pointers to thread the leaves, since the subtree with  $occ$  leaves has  $O(occ)$  nodes. The real problem of suffix trees is their high space consumption, which is  $O(n \log n)$  bits and at the very least 10 times the text size in practice [Kurtz 1998].

### 3.3 Suffix Arrays

A *suffix array* [Manber and Myers 1993; Gonnet et al. 1992] is simply a permutation of all the suffixes of  $T$  so that the suffixes are lexicographically sorted.

**Definition 6** *The suffix array of a text  $T_{1,n}$  is an array  $A[1, n]$  containing a permutation of the interval  $[1, n]$ , such that  $T_{A[i],n} < T_{A[i+1],n}$  for all  $1 \leq i < n$ , where “ $<$ ” between strings is the lexicographical order.*

The suffix array can be obtained by collecting the leaves of the suffix tree in left-to-right order (assuming that the children of the suffix tree nodes are lexicographically ordered left-to-right by the edge labels). However, it is much more practical to build them directly. In principle, any sorting algorithm can be used, as it is a matter of sorting the  $n$  suffixes of the text. There are several more sophisticated algorithms, from the original  $O(n \log n)$  time [Manber and Myers 1993] to the latest  $O(n)$  time algorithms [Kim et al. 2003; Ko and Aluru 2003; Kärkkäinen and Sanders 2003]. In practice, the best current algorithms are not yet linear time [Larsson and Sadakane 1999; Manzini and Ferragina 2004; Schürmann and Stoye 2005].

Fig. 3 shows our example suffix array. Note that each subtree of the suffix tree corresponds to the subinterval in the suffix array encompassing all its leaves (in the figure we have highlighted the interval corresponding to the suffix tree node that descends from the root by character “a”).

The suffix array plus the text contain enough information to search for patterns. Since the result of a suffix tree search is a subtree, the result of a suffix array search must be an interval. This is also obvious if one considers that all the suffixes prefixed by  $P$  are lexicographically contiguous in the sorted array  $A$ . Thus, it is possible to search for the interval in  $A$  of suffixes prefixed by  $P$  via two binary searches in  $A$ . The first binary search determines the starting position  $sp$  for the suffixes lexicographically larger than or equal to  $P$ . The second binary search determines the ending position  $ep$  for suffixes that start with  $P$ . Then the answer is the interval  $A[sp, ep]$ . A counting query needs to report just  $ep - sp + 1$ . A






Fig. 3. The suffix array of the text "alabar a la alabarda\$". We have shown explicitly where the suffixes starting with "a" point to.

locating query enumerates  $A[sp]$ ,  $A[sp + 1]$ ,  $\dots$ ,  $A[ep]$ .

Note that each step of the binary searches requires a lexicographical comparison between  $P$  and some suffix  $T_{A[i],n}$ , which requires  $O(m)$  time in the worst case. Hence the search is worst case time  $O(m \log n)$  (this can be lowered to  $O(m + \log n)$  by using more space to store longest common prefix information among consecutive suffixes [Manber and Myers 1993; Abouelhoda et al. 2002]). A locating query requires additional  $O(occ)$  time to report the  $occ$  occurrences. Fig. 4 gives the pseudocode.

---

**Algorithm** SASearch( $P_{1,m}, A[1,n], T_{1,n}$ )

- (1)  $sp \leftarrow 1$ ;  $st \leftarrow n$ ;
  - (2) **while**  $sp < st$  **do**
  - (3)      $s \leftarrow \lfloor (sp + st)/2 \rfloor$ ;
  - (4)     **if**  $P > T_{A[s], A[s]+m-1}$  **then**  $sp \leftarrow s + 1$  **else**  $st \leftarrow s$ ;
  - (5)  $ep \leftarrow sp - 1$ ;  $et \leftarrow n$ ;
  - (6) **while**  $ep < et$  **do**
  - (7)      $e \leftarrow \lceil (ep + et)/2 \rceil$ ;
  - (8)     **if**  $P = T_{A[e], A[e]+m-1}$  **then**  $ep \leftarrow e$  **else**  $et \leftarrow e - 1$ ;
  - (9) **return**  $(sp, ep)$ ;
- 

Fig. 4. Algorithm to search for  $P$  in suffix array  $A$  over text  $T$ .  $T$  is assumed to be terminated by "\$", but  $P$  is not. From the returned data one can answer the counting query  $ep - sp + 1$  or the locating query  $A[i]$ , for  $sp \leq i \leq ep$ .

All the space/time tradeoffs offered by the different succinct full-text indexes in this paper will be expressed in the form of tabular theorems, where the meaning of  $n$ ,  $m$ ,  $\sigma$ , and  $H_k$ , is implicit (see Section 4 for the definition of  $H_k$ ). For illustration and comparison, and because suffix arrays are the main focus of compressed indexes, we give now the corresponding theorem for suffix arrays. As the suffix array is not a self-index, the space in bits includes a final term  $n \log \sigma$ . The time to count refers to counting the occurrences of  $P_{1,m}$  in  $T_{1,n}$ . The time to locate refers to giving the text position of a single occurrence after counting has completed. The time to display refers to displaying one contiguous text substring of  $\ell$  characters. In this case (not a self-index), this is trivial as the text is readily available. We show two tradeoffs, the second referring to storing longest common prefix information.

**Theorem 1** *The Suffix Array of Manber and Myers [1993] (SA) offers the following space/time tradeoffs.*

Space in bits	$n \log n + n \log \sigma$
Time to count	$O(m \log n)$
Time to locate	$O(1)$
Time to display $\ell$ chars	$O(\ell)$
Space in bits	$2n \log n + n \log \sigma$
Time to count	$O(m + \log n)$
Time to locate	$O(1)$
Time to display $\ell$ chars	$O(\ell)$

#### 4. SUFFIX ARRAY REGULARITIES AND TEXT COMPRESSIBILITY

Suffix arrays are not random permutations. When the text alphabet size  $\sigma$  is smaller than  $n$ , not every permutation of  $[1, n]$  is the suffix array of some text (as there are more permutations than texts of length  $n$ ). Moreover, the  $k$ -th order empirical entropy of  $T$  is reflected in regularities that appear on its suffix array  $A$ . In this section we show how some measures of suffix array regularities are related to measures of text compressibility. Those relationships are used later to compress suffix arrays.

The analytical results in this section are presented at an intuitive level only. We refer the reader to the original sources to find the detailed proofs. Moreover, we sometimes deviate slightly from the original definitions, changing technical details to allow for a simpler exposition.

##### 4.1 $k$ -th Order Empirical Entropy

Opposed to the classical notion of  $k$ -th order entropy [Bell et al. 1990], which can only be defined for infinite sources, the  *$k$ -th order empirical entropy* defined by Manzini [2001] applies to finite texts. It coincides with the statistical estimation of the entropy of a source taking the text as a finite sample of the infinite source (actually, the same formula of Manzini [2001] is used by Grossi et al. [2003], yet it is interpreted in this latter sense). The definition is especially useful because it can be applied to any text without resorting to assumptions on its distribution.

**Definition 7** Let  $T_{1,n}$  be a text over an alphabet  $\Sigma$ . The zero-order empirical entropy of  $T$  is defined as

$$H_0 = H_0(T) = \sum_{c \in \Sigma} \frac{n_c}{n} \log \frac{n}{n_c},$$

where  $n_c$  is the number of occurrences of character  $c$  in  $T$ . The sum includes only those characters  $c$  that occur in  $T$ , so that  $n_c > 0$ .

**Definition 8** Let  $T_{1,n}$  be a text over an alphabet  $\Sigma$ . The  $k$ -th order empirical entropy of  $T$  is defined as

$$H_k = H_k(T) = \sum_{s \in \Sigma^k} \frac{|T_s|}{n} H_0(T_s), \quad (1)$$

where  $T_s$  is the subsequence of  $T$  formed by all the characters that occur followed by the context  $s$ . More precisely, if the occurrences of  $s$  in  $T_{2,n}$  start at positions

$p_1, p_2, \dots$ , then  $T_s = T_{p_1-1}T_{p_2-1} \dots$ . Again, we consider only contexts  $s$  that do occur in  $T$ .

We note that it is customary to define  $T_s$  regarding the characters *preceded* by  $s$ , rather than *followed* by  $s$ . We use the reverse definition for technical convenience. Although the empirical entropy of  $T$  and its reverse do not necessarily match, one can always work on the reverse text if this is an issue.

The empirical entropy of a text  $T$  provides a lower bound to the number of bits needed to compress  $T$  using any compressor that encodes each character considering only the context of  $k$  characters that follow it in  $T$ . Many self-indexes state their space requirement as a function of the empirical entropy of the indexed text. This is useful because it gives a measure of the index size with respect to the size the best  $k$ -th order compressor would achieve, thus relating the index size with the compressibility of the text.

We note that the classical entropy defined over infinite streams is always constant and can be zero. In contrast, the definition of  $H_k$  we use for finite texts is always positive, yet it can be  $o(1)$  on compressible texts. For an extreme example, consider  $T = (ab)^{n/2}\$,$  where  $H_0(T) = 1$  and  $H_k(T) \leq (1 + \log n)/n$  for  $k \geq 1$ . The technical consequence of the finiteness of the text is the presence of the unique “\$” character.

## 4.2 Self-Repetitions in Suffix Arrays

Consider again the suffix tree for our example text  $T = \text{"alabar a la alabarda\$"}$  depicted in Fig. 2. Observe, for example, that the subtree rooted at "abar" contains leaves  $\{3, 15\}$ , while the subtree rooted at "bar" contains leaves  $\{4, 16\}$ , that is, the same positions shifted by one. The reason is simple: every occurrence of "bar" in  $T$  is also an occurrence of "abar". Actually, the chain is longer: If one looks at subtrees rooted at "alabar", "labar", "abar", "bar", "ar", and "r", the same phenomenon occurs, and positions  $\{1, 13\}$  become  $\{6, 18\}$  after five steps. The same does not occur, for example, with the subtree rooted at " a", whose leaves  $\{7, 12\}$  do not repeat as  $\{8, 13\}$  inside another subtree. The reason is that not all occurrences of "a" belong to occurrences of " a" in  $T$ , and thus there are many more leaves rooted by "a" in the suffix tree, apart from 8 and 13.

Those repetitions show up in the suffix array  $A$  of  $T$ , depicted in Fig. 3. For example,  $A[18, 19] = A[10, 11] + 1$ . There are also longer regularities that do not correspond to a single subtree of the suffix tree, for example  $A[18, 21] = A[10, 13] + 1$ . Still, the text property responsible for the regularity is the same: All the text suffixes in  $A[10, 13]$  start with "a", while those in  $A[18, 21]$  are the same suffixes with the initial "a" excluded. The regularity appears because, for each pair of consecutive suffixes  $aX$  and  $aY$  in  $A[10, 13]$ , the suffixes  $X$  and  $Y$  are contiguous in  $A[18, 21]$ , that is, there is no other suffix  $Z$  such that  $X < Z < Y$  elsewhere in the text. This motivates the definition of *self-repetition* initially devised by Mäkinen [2000, 2003].

**Definition 9** *Given a suffix array  $A$ , a self-repetition is an interval  $[i, i + \ell]$  of  $[1, n]$  such that there exists another interval  $[j, j + \ell]$  such that  $A[j + r] = A[i + r] + 1$  for all  $0 \leq r \leq \ell$ . For technical convenience, cell  $A[1] = n$  is considered as a self-repetition of itself, of length 1.*

Note that  $A[1] = n$  because  $T_{n,n} = \$$ , which is lexicographically the smallest

suffix and hence it is pointed from  $A[1]$ .

A measure of the amount of regularity in a suffix array is how many self-repetitions we need to cover the whole array. This is captured by the following definition [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Definition 10** *Given a suffix array  $A$ , we define  $n_{sr}$  as the minimum number of self-repetitions necessary to cover the whole  $A$ . This is the minimum number of nonoverlapping intervals  $[i_s, i_s + \ell_s]$  that cover the interval  $[1, n]$  such that, for any  $s$ , there exists  $j_s$  such that  $A[j_s + r] = A[i_s + r] + 1$  for all  $0 \leq r \leq \ell_s$  (except for  $i_1 = 1$ , where  $A[i_1] = n$ , and thus  $\ell_1 = 0$  and  $j_1 = 1$ ).*

Self-repetitions are best highlighted through the definition of function  $\Psi$ , which tells where in the suffix array lies the pointer following the current one [Grossi and Vitter 2000; Sadakane 2000].

**Definition 11** *Given suffix array  $A[1, n]$ , function  $\Psi : [1, n] \rightarrow [1, n]$  is defined so that, for all  $1 \leq i \leq n$ ,  $A[\Psi(i)] = A[i] + 1$ . The exception is  $A[1] = n$ , in which case we require  $A[\Psi(1)] = 1$  so that  $\Psi$  is actually a permutation.*

Function  $\Psi$  is heavily used in most compressed suffix arrays, as seen later. There are several properties of  $\Psi$  that makes it appealing to compression. A first one establishes that  $\Psi$  is monotonically increasing in the areas of  $A$  that point to suffixes starting with the same character [Grossi and Vitter 2000; Sadakane 2000].

**Lemma 1** *Given a text  $T_{1,n}$ , its suffix array  $A[1, n]$ , and the corresponding function  $\Psi$ , it holds  $\Psi(i) < \Psi(i + 1)$  whenever  $T_{A[i]} = T_{A[i+1]}$ .*

To see that the lemma holds, assume that  $T_{A[i],n} = cX$  and  $T_{A[i+1],n} = cY$ , so  $cX < cY$  and then  $X < Y$ . Thus  $T_{A[i+1],n} = T_{A[\Psi(i)],n} = X$  and  $T_{A[i+1]+1,n} = T_{A[\Psi(i+1)],n} = Y$ . So  $T_{A[\Psi(i)],n} < T_{A[\Psi(i+1)],n}$ , and thus  $A[\Psi(i)]$  must occur before  $A[\Psi(i + 1)]$  in  $A$ , so  $\Psi(i) < \Psi(i + 1)$ .

Another interesting property is a special case of the above: how does  $\Psi$  behave inside a self-repetition  $A[j+r] = A[i+r] + 1$  for  $0 \leq r \leq \ell$ . Note that  $\Psi(i+r) = j+r$  in all the interval. Fig. 5 illustrates (for now consider only the first two arrays). This motivates the definition of *runs in  $\Psi$*  [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Definition 12** *A run in  $\Psi$  is a maximal interval  $[i, i + \ell]$  in sequence  $\Psi$  such that  $\Psi(i + r + 1) - \Psi(i + r) = 1$  for all  $0 \leq r < \ell$ .*

Note that the number of runs in  $\Psi$  is  $n$  minus the number of positions  $i$  such that  $\Psi(i + 1) - \Psi(i) = 1$ . The following lemma should not be surprising [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Lemma 2** *The number of self-repetitions  $n_{sr}$  to cover a suffix array  $A$  is equal to the number of runs in the corresponding  $\Psi$  function.*



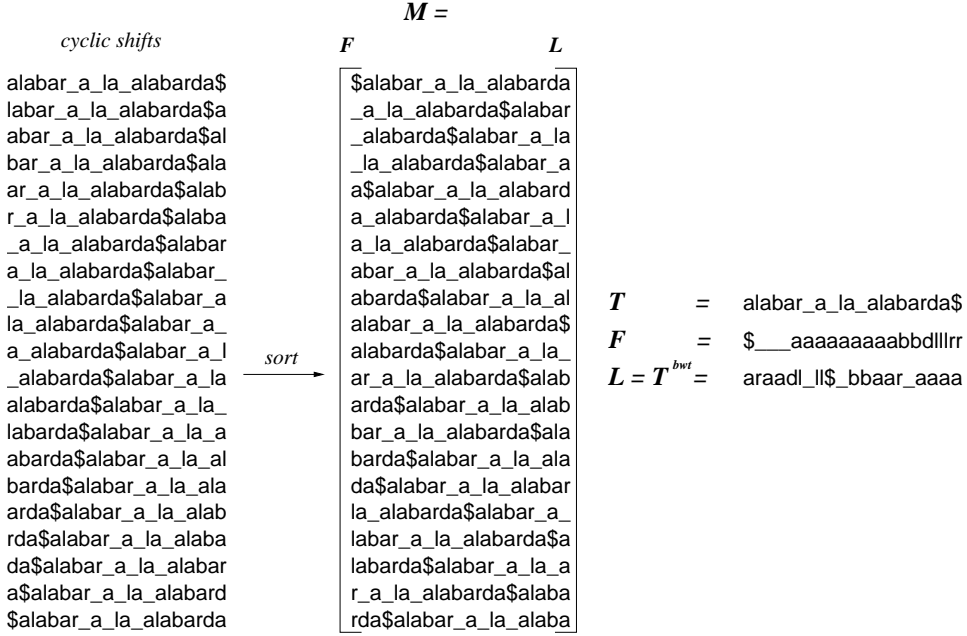


Fig. 6. Obtaining the BWT for the text "alabar a la alabarda\$". Note that  $M[i]$  is essentially  $T_{A[i],n}$  of Fig. 3.

character. Note, however, that all the occurrences of "a" in  $F$  are sorted according to the suffix that follows the "a". Likewise, all the occurrences of "a" in  $L$  are also sorted accordingly to the suffix that follows that "a". Therefore, equal characters preserve the same order in  $F$  and  $L$ . As there are 4 characters smaller than "a" in  $T$  and 6 occurrences of "a" in  $L[1, 18]$ , we have that  $L[18] = \text{"a"}$  occurs at  $F[4 + 6] = F[10]$ , that is,  $LF(18) = 10$ .

The following lemma gives the formula for the LF-mapping [Burrows and Wheeler 1994; Ferragina and Manzini 2000].

**Lemma 3** *Let  $T_{1,n}$  be a text and  $F$  and  $L$  be the result of its BWT. Let  $C : \Sigma \rightarrow [1, n]$  and  $Occ : \Sigma \times [1, n] \rightarrow [1, n]$ , such that  $C(c)$  is the number of occurrences in  $T$  of characters alphabetically smaller than  $c$ , and  $Occ(c, i)$  is the number of occurrences of character  $c$  in  $L[1, i]$ . Then, it holds  $LF(i) = C(L[i]) + Occ(L[i], i)$ .*

With this mapping, reversing the BWT is simple, as  $L[i]$  always precedes  $F[i]$  in  $T$ . Since  $T_n = \$$ , the first cyclic shift is  $T_n T_{1,n-1}$ , and therefore  $T_{n-1} = L[1]$ . We now compute  $i = LF(1)$  to know that  $T_{n-1}$  is at  $F[i]$ . Thus  $T_{n-2} = L[i]$  precedes  $F[i] = T_{n-1}$ . With  $i' = LF(i)$  we know that  $T_{n-2}$  is at  $F[i']$  and thus  $T_{n-3} = L[i']$ , and so on,  $T_{n-k} = L[LF^{k-1}(1)]$ .

We finally notice that  $A[LF(i)] = A[i] - 1$ , and therefore the following observation is immediate, although crucial to understand the relation between different kinds of existing self-indexes [Sadakane 2000; Ferragina and Manzini 2000].

**Lemma 4** *Functions  $LF$  and  $\Psi$  are one the inverse of the other.*

#### 4.4 Relation between Regularities and Compressibility

We start by pointing out a simple but essential relation between  $T^{bwt}$  and  $H_k$ . Note that, for each text context of length  $k$ , all the suffixes starting with that context appear consecutively in  $A$ . Therefore, the characters that precede each context appear consecutively in  $T^{bwt}$ . The following lemma [Ferragina and Manzini 2004; Ferragina et al. 2004a] shows that it suffices to compress the characters of each context with a zero-order compressor to achieve  $k$ -th order compression.

**Theorem 2** *Given  $T_{1,n}$  over an alphabet of size  $\sigma$ , with  $k$ -th order empirical entropy  $H_k$ , if we divide  $T^{bwt}$  into (at most)  $\sigma^k$  pieces according to the text context that follows each character in  $T$ , and then compress each piece  $T_s$  corresponding to context  $s$  using  $c|T_s|H_0(T_s) + f(|T_s|)$  bits, where the second derivative of  $f$  is nonpositive, then the representation for the whole  $T^{bwt}$  requires at most  $cnH_k + \sigma^k f(n/\sigma^k)$  bits.*

To see that the theorem holds, it is enough to recall Eq. (1), as we are representing the characters  $T_s$  followed by each of the contexts  $s \in \Sigma^k$  using space proportional to  $|T_s|H_0(T_s)$ . The extra space,  $\sigma^k f(n/\sigma^k)$ , is just the worst case of the sum of  $\sigma^k$  values  $f(|T_s|)$  where the values  $|T_s|$  add up  $n$ .

Thus, it is enough to encode each portion of  $T^{bwt}$  with a zero-order compressor in order to obtain a  $k$ -th order compressor for  $T$ , for any  $k$ . The price of using a longer context (larger  $k$ ) is paid in the extra  $\sigma^k f(n/\sigma^k)$  term. This can be thought of as the price to manage the model information, and it can easily dominate the overall space if  $k$  is not small enough.

Let us now consider the number of equal-letter runs in  $T^{bwt}$ . This will be related both to self-repetitions in suffix arrays and to the empirical entropy of  $T$  [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Definition 15** *Given a Burrows-Wheeler transformed text  $T^{bwt}$ ,  $n_{bw}$  is the number of equal-letter runs in  $T^{bwt}$ , that is,  $n$  minus the number of positions  $j$  such that  $T_{j+1}^{bwt} = T_j^{bwt}$ .*

There is a close relationship between the number of runs in  $\Psi$  (or the number of self-repetitions in  $A$ ),  $n_{sr}$ , and the number of equal-letter runs in  $T^{bwt}$ ,  $n_{bw}$  [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Lemma 5** *The following relation between  $n_{sr}$  and  $n_{bw}$  holds for any text  $T$  on an alphabet of size  $\sigma$ :  $n_{sr} \leq n_{bw} \leq n_{sr} + \sigma$ .*

To see why the lemma holds, consider Fig. 3 and 5. Let us focus on the longest self-repetition,  $A[10, 13] = \{1, 13, 5, 17\}$ . All those suffixes (among others) start with "a". The self-repetition occurs in  $\Psi([10, 13]) = [18, 21]$ , that is,  $A[18, 21] = \{2, 14, 6, 18\}$ . All those suffixes are preceded by "a", because all  $\{1, 13, 5, 17\}$  start with "a". Hence there is a run of characters "a" in  $T_{18,21}^{bwt}$ .

In all cases where all the suffixes of a self-repetition start with the same character, there will be an equal-letter run in  $T^{bwt}$ . Note that, because of the lexicographical

ordering, consecutive suffixes change their first character at most  $\sigma$  times throughout  $A[1, n]$ . Thus, save at most  $\sigma$  exceptions, every time  $\Psi(i+1) - \Psi(i) = 1$ , there will be a distinct  $j = \Psi(i)$  such that  $T_{j+1}^{bwt} = T_j^{bwt}$ . Thus  $n_{bw} \leq n_{sr} + \sigma$ . (There is one such exception in Fig. 5, where the self-repetition  $A[16, 17] + 1 = A[5, 6]$  does not correspond to an equal-letter run in  $T_{5,6}^{bwt}$ .)

On the other hand, every time  $T_{j+1}^{bwt} = T_j^{bwt}$ , we know that suffix  $T_{A[j], n} = X$  is followed by  $T_{A[j+1], n} = Y$ , and both are preceded by the same character  $c$ . Hence suffixes  $cX$  and  $cY$  must also be contiguous, at positions  $i$  and  $i+1$  so that  $\Psi(i) = j$  and  $\Psi(i+1) = j+1$ , thus it holds  $\Psi(i+1) - \Psi(i) = 1$  for a distinct  $i$  every time  $T_{j+1}^{bwt} = T_j^{bwt}$ . Therefore,  $n_{sr} \leq n_{bw}$ .

We finally relate the  $k$ -th order empirical entropy of  $T$  with  $n_{bw}$  [Mäkinen and Navarro 2004a, 2005a, 2005b].

**Theorem 3** *Given a text  $T_{1,n}$  over an alphabet of size  $\sigma$ , with  $k$ -th order empirical entropy  $H_k$ , and given its Burrows-Wheeler transformed text  $T^{bwt}$  with  $n_{bw}$  equal-letter runs, it holds  $n_{bw} \leq nH_k + \sigma^k$  for any  $k \geq 0$ . In particular, for any  $0 < \alpha < 1$ , it holds  $n_{bw} \leq nH_k + O(n^\alpha) = nH_k + o(n)$  for any  $k \leq \alpha \log_\sigma n$ . The bounds are obviously valid for  $n_{sr} \leq n_{bw}$  as well.*

We only attempt to give a flavor of why the theorem holds. The idea is to partition  $T^{bwt}$  according to the contexts of length  $k$ . Following Eq. (1),  $nH_k$  is the sum of zero-order entropies over all the contexts. It can then be shown that, within a single context  $T_{i,j}^{bwt}$ , the number of equal-letter runs in  $T_{i,j}^{bwt}$  can be upper bounded in terms of the zero-order entropy of the string  $T_{i,j}^{bwt}$ . A constant  $f(|S|) = 1$  in the upper bound is responsible for the  $\sigma^k$  overhead, which is the number of possible contexts of length  $k$ . Thus the rest is a consequence of Theorem 2.

## 5. BASIC COMPACT DATA STRUCTURES

We will learn later that nearly all approaches to represent suffix arrays in compressed form take advantage of *compressed representations of binary sequences*. That is, we are given a bit vector (or bit string)  $B_{1,n}$ ,  $B_i \in \{0, 1\}$  for  $1 \leq i \leq n$ , and we want to compress it while at the same time supporting several operations on it. Typical operations that are required are the following:

- $B_i$ : Accesses the  $i$ -th element of the compressed sequence.
- $rank_b(B, i)$ : Returns the number of times bit  $b$  appears in the prefix  $B_{1,i}$ .
- $select_b(B, j)$ : Returns the position  $i$  of the  $j$ -th appearance of bit  $b$  in  $B_{1,n}$ .

Other useful operations are  $prev_b(B, i)$  and  $next_b(B, i)$ , which give the position of the previous/next bit  $b$  from position  $i$ . However, these operations can be expressed via  $rank$  and  $select$ , and hence are usually not considered separately. Notice also that  $rank_0(B, i) = i - rank_1(B, i)$ , so considering  $rank_1(B, i)$  will be enough. However, the same duality does not hold for  $select$ , and we have to consider both  $select_0(B, j)$  and  $select_1(B, j)$ . We call a representation of  $B$  *complete* if it supports all the listed operations in constant time. Representation is *partial* if it supports the listed operations only for 1-bits, that is, it supports  $rank_b(B, i)$  only if  $B_i = 1$  and it only supports  $select_1(B, j)$ .



The study of succinct representations of various structures, including bit vectors, was initiated by Jacobson [1989]. The main motivation to study these operations came from the possibility to simulate tree traversals in small space: It is possible to represent the shape of a tree as a bit vector, and then the traversal from a node to a child and vice versa can be expressed via constant number of *rank* and *select* operations. Jacobson showed that attaching a dictionary of size  $o(n)$  to the bit vector  $B_{1,n}$  is sufficient to support *rank* operation in constant time on a RAM model. He also studied *select* operation, but for the RAM model the solution was not yet optimal. Later, Munro [1996] and Clark [1996] obtained constant time *select* operation on the RAM model, using the same  $o(n)$  extra space as for *rank*.

Although the  $n + o(n)$  solutions are asymptotically optimal for incompressible binary sequences, one can obtain more space-efficient representations for compressible ones. Consider, for example,  $select_1(B, i)$  operation on a bit vector containing  $\ell = o(n/\log n)$  1-bits. One can store all answers explicitly using  $O(\ell \log n) = o(n)$  bits.

Pagh [1999] was the first to study compressed representations of bit vectors supporting more than just access to  $B_i$ . He gave a representation of bit vector  $B_{1,n}$  that uses  $\lceil \log \binom{n}{\ell} \rceil + o(\ell) + O(\log \log n)$  bits, where  $\ell$  is the number of 1-bits in  $B$ . In principle this representation supported only  $B_i$  queries, yet it also supported *rank* queries for sufficiently dense bit vectors,  $n = O(\ell \text{polylog}(\ell))$ . Notice that  $\log \binom{n}{\ell} = nH_0(B) + O(\log n)$ .

This result was later enhanced by Raman, Raman, and Rao [2002], who developed a representation with similar space complexity,  $nH_0(B) + o(\ell) + O(\log \log n)$  bits, supporting *rank* and *select*. However, this representation is partial: they support  $select_1(B, j)$  but not  $select_0(B, j)$ , and  $rank_1(B, i)$  only when  $B_i = 1$ . They also provide a new complete representation requiring  $nH_0(B) + O(n \log \log n / \log n)$  bits.

The missing piece, therefore, is achieving a complete representation that needs  $o(\ell)$  instead of  $o(n)$  extra space when  $\ell$  is significantly smaller than  $n$ . This, however, seems not possible, as there is a matching lower bound [Miltersen 2005].

In the rest of this section we explain the most intuitive of these results, to give a flavor of how some of the solutions work. We also show how results extend to non-binary sequences.

### 5.1 Basic $n + o(n)$ -bit Solutions for Binary Sequences

We start by explaining the  $n + o(n)$  bits solution [Jacobson 1989; Munro 1996; Clark 1996] supporting  $rank_1(B, i)$  and  $select_1(B, i)$  in constant time. Then  $rank_0(B, i)$  is obtained automatically and  $select_0(B, i)$  is symmetric.

Let us start with *rank*. The structure is composed of a two-level dictionary for partial solutions, plus a global table storing answers for every possible short binary sequence. The answer to a *rank* query is formed by summing values from these dictionaries and tables.

For clarity of presentation we assume  $n$  is a power of four. The general case is handled by considering floors and ceilings when necessary. We assume all divisions  $x/y$  to give the integer  $\lfloor x/y \rfloor$ .

Let us start from the last level. Consider a substring  $I$  of  $B_{1,n}$  of length  $t = \frac{\log n}{2}$ . This case is handled by the so-called four-Russians technique [Arlazarov et al. 1975]: We build a table  $G[0, \sqrt{n} - 1][0, t]$  storing all answers to *rank* queries for

all binary sequences of length  $t$  (note that  $2^t = \sqrt{n}$ ). Then  $rank_1(I, i) = G[I][i]$  is obtained in constant time. To index  $G$ ,  $I$  is regarded as an integer in the usual way. Note that this can be extended to substrings of length  $c \log n$ , which would be solved in at most  $2c$  accesses to table  $G$ . For example, if  $I$  is of length  $\log n$ , then  $rank_1(I, t + 3) = G[I^1, t] + G[I^2, 3]$ , where  $I^1$  and  $I^2$  are the two halves of  $I$ .

We could complete the solution by dividing  $B$  into *blocks* of length  $\log n$  and explicitly storing *rank* answers for block boundaries, in a table  $D[0, \frac{n}{\log n} - 1]$ , such that  $D[q] = rank_1(B, q \log n)$  for  $0 \leq q < \frac{n}{\log n}$ . Then, given  $i = q \log n + r$ ,  $0 \leq r < \log n$ , we have  $rank_1(B, i) = D[q] + rank_1(B_{q \log n + 1, q \log n + \log n}, r)$ . As the latter  $rank_1$  query is answered in constant time using table  $G$ , we have constant time *rank* queries on  $B$ .

The problem with  $D$  is that there are  $\frac{n}{\log n}$  blocks in  $B$ , each of which requires  $\log n$  bits in  $D$ , for  $n$  total bits. To obtain  $o(n)$  extra space, we build a *superblocks* dictionary  $D^s[0, \frac{n}{\log^2 n} - 1]$  such that  $D^s[q'] = rank_1(B, q' \log^2 n)$  for  $0 \leq q' < \frac{n}{\log^2 n}$ . We replace the blocks dictionary  $D$  with  $D^b$ , such that  $D^b$  stores *relative* answers, inside the corresponding superblock of  $D^s$ . That is,  $D^b[q] = D[q] - D^s[\lfloor \frac{q}{\log n} \rfloor]$  for  $0 \leq q < \frac{n}{\log n}$ . It follows that, for  $i = q' \log^2 n + r' = q \log n + r$ ,  $0 \leq r' < \log^2 n$ ,  $0 \leq r < \log n$ ,

$$rank_1(B, i) = D^s[q'] + D^b[q] + rank_1(B_{q \log n + 1, q \log n + \log n}, r),$$

where the last  $rank_1$  query is answered in constant time using table  $G$ .

The values stored in  $D^b$  are in the range  $[0, \log^2 n - 1]$ , hence table  $D^b$  takes  $O(n \log \log n / \log n)$  bits. Table  $D^s$  takes  $O(n / \log n)$  bits, and finally table  $G$  takes  $O(\sqrt{n} \log^2 n)$  bits. We have obtained the claimed  $n + o(n)$  bits representation of  $B_{1,n}$  supporting constant time *rank*. Fig. 7 illustrates the structure.

<b>B</b>	0 1 0 1 1 0 1 0 0 <b>1 1 1</b> 0 1 0 0	<b>G</b>	1	2
<b>D<sup>s</sup></b>	0	<b>00</b>	0	0
<b>D<sup>b</sup></b>	0	<b>01</b>	0	1
	2	<b>10</b>	1	1
	4	<b>11</b>	1	2
	7			
	0			

$$\begin{aligned}
 rank_1(B, 11) &= D^s[0] + D^b[2] + G[01, 2] + G[11, 1] \\
 &= 0 + 4 + 1 + 1 = 6
 \end{aligned}$$

Fig. 7. An example of constant time *rank* computation using  $n + o(n)$  bits of space.

Extending the structure to provide constant time *select* queries is more complicated. We explain a simple version here which shares the essential idea of the real solutions [Munro 1996; Clark 1996].

We partition the space  $[1, n]$  of possible *arguments* of *select*, that is, the values  $j$  of queries  $select_1(B, j)$ . We cut  $[1, n]$  into blocks of  $\log^2 n$  bits. A dictionary  $D^1[j]$  answers  $select_1(B, j \log^2 n)$ , requiring  $O(n / \log n)$  bits.

Some of those blocks may span a large extent in  $B$  (with many 0-bits and just  $\log^2 n$  1-bits). A fundamental problem of using blocks and superblocks for *select* is that there is no guarantee that relative answers inside blocks do not require  $\log n$  bits anyway. A block is called *long* if it spans more than  $\log^4 n$  positions in  $B$ , and *short* otherwise. Note that there cannot be more than  $n/\log^4 n$  long blocks. As long blocks are problematic, we simply store all their  $\log^2 n$  answers explicitly. As each answer requires  $\log n$  bits, this accounts for other  $n/\log n$  bits overall.

The short blocks contain  $n' = \log^2 n$  1-bits (arguments for  $select_1$ ) and span at most  $\log^4 n$  bits in  $B$ . We divide them again into *miniblocks* of  $\log^2 n' = O((\log \log n)^2)$  arguments. A miniblock directory  $D^2[j]$  will store the relative answer to  $select_1$  inside the short block,  $D^2[j] = select_1(B, j \log^2 n') - D^1[\frac{j \log^2 n'}{\log^2 n}]$ . Values in  $D^2$  are in the range  $[1, \log^4 n]$  and thus require  $O(\log \log n)$  bits. Thus  $D^2$  requires  $O(n/\log \log n)$  bits. A miniblock will be called *long* if it spans more than  $\log n$  bits in  $B$ . For long miniblocks, we will again store all their answers explicitly. There are at most  $n/\log n$  long miniblocks overall, so storing all the  $\log^2 n'$  answers of all long miniblocks requires  $O(n(\log \log n)^3/\log n)$  bits. Finally, short miniblocks span at most  $\log n$  bits in  $B$ , so a precomputed table analogous to  $G$  gives their answers using  $O(\sqrt{n} \log^2 n)$  bits. This completes the solution. These structures have to be duplicated for  $select_0(B, j)$ .

**Theorem 4** *Bit vector  $B_{1,n}$  can be represented using  $n + o(n)$  bits of space so that  $B_i$ ,  $rank_b(B, i)$ , and  $select_b(B, j)$ , can be answered in constant time.*

There exist implementations of these solutions which work well in practice [González et al. 2005].

## 5.2 More Sophisticated $nH_0$ -bits Solutions

We explain now how to improve the representation of previous section for compressible sequences, so as to obtain complete representations requiring  $nH_0(B) + o(n)$  bits of space [Raman et al. 2002]. Our presentation is simplified and will only cover  $rank_1(B, i)$  queries.

We cut  $B$  into blocks of fixed length  $t = \frac{\log n}{2}$ . Each such block  $I = B_{ti+1, ti+t}$  with  $\ell$  bits set will belong to class  $\ell$  of  $t$ -bitmaps. For example, if  $t = 4$ , then class 0 is  $\{0000\}$ , class 1 is  $\{0001, 0010, 0100, 1000\}$ , class 2 is  $\{0011, 0101, 0110, 1001, 1010, 1100\}$ , and so on until class  $t = 4$ ,  $\{1111\}$ . As class  $\ell$  contains  $\binom{t}{\ell}$  elements, an index to identify a  $t$ -bitmap within its class requires only  $\log \binom{t}{\ell}$  bits. Instead of using  $t$  bits to represent a block, we use two components  $(\ell, r)$ : a class identifier  $0 \leq \ell \leq t$ , using  $\lceil \log(t+1) \rceil$  bits, and an index  $r$  within that class, using  $\lceil \log \binom{t}{\ell} \rceil$  bits. Then  $B$  is represented as a sequence of  $\lceil n/t \rceil$  such descriptions.

The class identifiers amount to  $O(n \log t/t) = O(n \log \log n / \log n)$  bits overall. The interesting part is the sequence of indexes. Let  $I^i$  be the  $i$ -th block, with  $\ell_i$  bits set. Let us also call  $\ell_B$  the total number of bits set in  $B$ . The number of bits required by all the blocks  $I^1 \dots I^{n/t}$  is [Pagh 1999]

$$\left\lceil \log \binom{t}{\ell_1} \right\rceil + \dots + \left\lceil \log \binom{t}{\ell_{\lceil n/t \rceil}} \right\rceil \leq \log \left( \binom{t}{\ell_1} \times \dots \times \binom{t}{\ell_{\lceil n/t \rceil}} \right) + n/t$$

$$\leq \log \binom{n}{\ell_1 + \dots + \ell_{\lceil n/t \rceil}} + n/t = \log \binom{n}{\ell_B} + n/t \leq nH_0(B) + O(n/\log n),$$

where the second inequality holds because the ways to choose  $\ell_i$  bits from each block of  $t$  bits are included in the ways to choose  $\ell_B$  bits out of  $n$ . Thus, we have represented  $B$  with  $nH_0(B) + o(n)$  bits.

We need more structures to answer queries on  $B$ . The same directories  $D^s$  and  $D^b$  used in Section 5.1, with block size  $t$ , are used. As the descriptions  $(\ell, r)$  have varying length, we need position directories  $P^s$  and  $P^b$ , which work like  $D^s$  and  $D^b$  to give the position in the compressed  $B$  where the description of each block starts.

In order to complete the solution with table  $G$ , we must obtain the bitmap  $I$  from its description. For each class  $\ell$  we store an array  $E_\ell$ , so that  $E_\ell[r]$  is the bitmap with index  $r$  of class  $\ell$  (in our example,  $E_2[4] = 1001$ ). Therefore, instead of  $G[I, i]$  as in Section 5.1, we use  $G[E_\ell[r], i]$ , as  $I$  is not directly available but just its description  $(\ell, r)$ . Table  $E$  needs overall  $t^2 2^t = O(\sqrt{n} \log^2 n)$  bits. (The real solution does not use  $G$  but stores the answers to queries directly in  $E$ , instead of the bitmaps.) Thus all the extra structures still require  $o(n)$  bits. Fig. 8 illustrates.

<b>B</b>	0 1 0 1 1 0 1 0 0 1 1 1 0 1 0 0	<b>E<sub>0</sub></b>	1 2
		<b>00</b>	0 0
<b>D<sup>s</sup></b>	0	8	<b>E<sub>1</sub></b>
			1 2
<b>D<sup>b</sup></b>	0	1	2
			3
			4
			5
			7
			8
			0
<b>B'</b>	(1,0)(1,0)(1,1)(1,1)(1,0)(2,0)(1,0)(0,0)	<b>E<sub>2</sub></b>	1 2
		<b>11</b>	1 2
rank <sub>1</sub> (B,11) = D <sup>s</sup> [0] + D <sup>b</sup> [5] + E <sub>2</sub> [11,1]			
= 0 + 5 + 1 = 6			

Fig. 8. An example of constant time *rank* computation using  $nH_0(B) + o(n)$  bits of space. Here  $B'$  is the sequence of pairs  $(\ell, r)$ . In a real implementation, this sequence is represented as a binary string, each pair occupying a variable number of bits, and the values in  $P^s$  and  $P^b$  point to the corresponding positions in the binary representation of  $B'$ . Note also that we have integrated the answers of table  $G$  into table  $E$ , as in the real implementation.

For *select*, the solution is again more complicated, but it also uses the strategy of dividing blocks into short and long. We omit the details here.

**Theorem 5** *Bit vector  $B_{1,n}$  can be represented using  $nH_0(B) + O(n \log \log n / \log n)$  bits of space so that  $B_i$ ,  $rank_b(B, i)$ , and  $select_b(B, j)$ , can be answered in constant time.*

### 5.3 Handling General Sequences and Wavelet Trees

Consider now a sequence (or string)  $S_{1,n}$  from an alphabet of size  $\sigma$ . We wish to support  $rank_c(S, i)$  and  $select_c(S, j)$  for all alphabet symbols  $c$ :  $rank_c(S, i)$  gives the number of times character  $c$  appears in  $S_{1,i}$  and  $select_c(S, j)$  gives the position of the  $j$ -th  $c$  in  $S$ . Analogously to the binary case, we call a representation of  $S$  *complete* if it supports access to  $S$ ,  $rank_c$ , and  $select_c$ , in constant time for all symbols  $c$ .

We can easily obtain a complete representation for  $S$  using the results from previous section: Consider indicator bit vectors  $B_{1,n}^c$  such that  $B_i^c = 1$  iff  $S_i = c$ . Then  $rank_c(S, i) = rank_1(B^c, i)$  and  $select_c(S, j) = select_1(B^c, j)$ . Using Theorem 5, the representations of vectors  $B^c$  take overall  $\sum_{c \in \Sigma} (nH_0(B^c) + o(n)) = H_0(S) + O(n) + o(\sigma n)$  bits.

The  $O(n)$  extra term can be removed with a more careful design [Ferragina et al. 2004c]. Essentially, one can follow the development leading to Theorem 5 on a general sequence. Now binomials become multinomials and the scheme is a bit more complicated, but the main idea does not change. This leads to the next observation.

**Lemma 6** *Sequence  $S_{1,n}$  over an alphabet of size  $\sigma$  can be represented using  $nH_0(S) + O(\sigma n \log \log n / \log_\sigma n)$  bits of space so that  $B_i$ ,  $rank_c(S, i)$ , and  $select_c(S, j)$ , can be answered in constant time.*

This complete representation of  $S$  takes sublinear space,  $o(n \log \sigma)$ , only if  $\sigma = o(\log n / \log \log n)$ .

A completely different technique is the *wavelet tree* [Grossi et al. 2003]. Fig. 21, on page 44, illustrates this structure. The wavelet tree is a perfect binary tree of height  $\lceil \log \sigma \rceil$ , built on the alphabet symbols, such that the root represents the whole alphabet and each leaf represents a distinct alphabet symbol. If a node  $v$  represents alphabet symbols in the range  $\Sigma^v = [i, j]$ , then its left child  $v_l$  represents  $\Sigma^{v_l} = [i, \frac{i+j}{2}]$  and its right child  $v_r$  represents  $\Sigma^{v_r} = [\frac{i+j}{2} + 1, j]$ .

We associate to each node  $v$  the subsequence  $S^v$  of  $S$  formed by the characters in  $\Sigma^v$ . However, sequence  $S^v$  is not really stored at the node. Instead, we store a bit sequence  $B^v$  telling whether characters in  $S^v$  go left or right, that is,  $B_i^v = 1$  iff  $S_i^v \in \Sigma^{v_l}$ .

All queries on  $S$  are easily answered in  $O(\log \sigma)$  time with the wavelet tree, provided we have complete representations of the bit vectors  $B^v$ . To determine  $S_i$ , we check  $B_i^{root}$  to decide whether to go left or right. If we go left, we now seek character  $rank_0(B^{root}, i)$  in the left child of the root, otherwise we seek character  $rank_1(B^{root}, i)$  in its right child. We continue recursively until reaching a leaf corresponding to a single character, which is the original  $S_i$ . Similarly, to answer  $rank_c(S, i)$ , we go left or right, adapting  $i$  accordingly. Once we arrive at a leaf, the current  $i$  value is the answer. Fig. 21 gives an example for  $rank_{\text{"a"}}(S, 15) = 5$ .

Finally, to answer  $select_c(S, j)$ , we start at the leaf corresponding to  $c$  and move upwards in the tree. If the leaf is a left child, then the position corresponding to  $j$  in its parent  $v$  is  $select_0(B^v, j)$ , otherwise it is  $select_1(B^v, j)$ . When we reach the root, the current  $j$  value is the answer. For example, in Fig. 21,  $select_{\text{"a"}}(S, 5)$  starts with the leaf for "a". It is a right child, so in its parent the position is  $select_1(1110001101111, 5) = 8$ . This in turn is a left child, so in its parent (the root), the final position is  $select_0(010011011001100100000, 8) = 15$ .

If we use the complete representation of Theorem 4 for the bit vectors  $B^v$ , the overall space is  $n \log \sigma (1 + o(1))$ , that is, essentially the same space to store  $S$  (and we do not need to also store  $S$ ). Yet, by using the representation of Theorem 5, the sum of entropies of all bit vectors simplifies to  $nH_0(S)$  and the extra terms add up  $O(n \log \log n / \log_\sigma n) = o(n \log \sigma)$  [Grossi et al. 2003].

**Theorem 6** *Sequence  $S_{1,n}$  over an alphabet of size  $\sigma$  can be represented using the wavelet tree in  $nH_0(S) + O(n \log \log n / \log_\sigma n)$  bits of space, so that  $S_i$ ,  $\text{rank}_c(S, i)$ , and  $\text{select}_c(S, j)$ , can be answered in  $O(\log \sigma)$  time.*

It is possible to combine the representations of Lemma 6 and Theorem 6. The former gives a complete representation (constant query time) with sublinear extra space, for  $\sigma = o(\log n / \log \log n)$ . The latter works for any alphabet  $\sigma$  but it pays  $O(\log \sigma)$  query time. By using  $r$ -ary wavelet trees, with suitable  $r$ , one can obtain a complete representation that works for  $\sigma = \text{polylog}(n)$  [Ferragina et al. 2004c, 2004b]. The idea is that, instead of bitmaps at each node, we represent sequences over an alphabet of size  $r$  (to represent the tree branch chosen by each character). Those sequences are handled with Lemma 6.

**Lemma 7** *Sequence  $S_{1,n}$  over an alphabet of size  $\sigma$  can be represented using an  $r$ -ary wavelet tree, for  $2 \leq r \leq \min(\sigma, \sqrt{n})$ , in  $nH_0(S) + O(\sigma \log n) + O(rn \log \log n / \log_\sigma n)$  bits of space, so that  $S_i$ ,  $\text{rank}_c(S, i)$  and  $\text{select}_c(S, j)$ , can be answered in time  $O(\log_r \sigma)$ .*

The number of levels of the wavelet tree can be made constant when  $\sigma = O(\log^d n)$ , by choosing  $r$  of the form  $\sigma^{1/\kappa}$ , for constant  $\kappa > d/(1 - \epsilon)$ , with any  $0 < \epsilon < 1$  of our choice. The result is a structure with constant access time.

**Theorem 7** *Sequence  $S_{1,n}$  over an alphabet of size  $\sigma = O(\text{polylog}(n))$  can be represented using a multi-ary wavelet tree in  $nH_0(S) + O(n / \log^\epsilon n)$  bits of space, for any constant  $0 < \epsilon < 1$ , so that  $S_i$ ,  $\text{rank}_c(S, i)$ , and  $\text{select}_c(S, j)$ , can be answered in constant time  $O(\frac{1}{1-\epsilon})$ .*

For larger  $\sigma = O(n^\beta)$ ,  $0 < \beta < 1$ , the scheme still improves over the binary wavelet tree, by choosing  $r = O(\log n / (\log \log n)^2)$  [Ferragina et al. 2004b]. Actually the result holds for any  $\sigma = o(n / \log \log n)$ .

**Theorem 8** *Sequence  $S_{1,n}$  over an alphabet of size  $\sigma = o(n / \log \log n)$ , can be represented using a multi-ary wavelet tree in  $nH_0(S) + O(n \lceil \log \sigma / \log \log n \rceil)$  bits of space, so that  $S_i$ ,  $\text{rank}_c(S, i)$ , and  $\text{select}_c(S, j)$ , are answered in  $O(\lceil \log \sigma / \log \log n \rceil)$  time.*

## 6. COMPRESSED SUFFIX ARRAYS

The first type of compressed indexes we are going to review can be considered as the result of the *abstract optimization* of the suffix array data structure. That is, the search algorithm remains essentially as in Fig. 4, but suffix array  $A$  is taken as an abstract data type that gives access to the array in some way. This abstract data type is implemented using as little space as possible. This is the case of the *Compact Suffix Array* of Mäkinen [2000, 2003] (MAK-CSA) and the *Compressed Suffix Array* of Grossi and Vitter [2000, 2005] (GV-CSA). Both ideas appeared simultaneously and independently during 2000, and they are based on different ways of exploiting the regularities that appear on the suffix arrays of compressible texts. Those structures are still not self-indexes, as they need the text  $T$  to operate.

### 6.1 MAK-CSA: Mäkinen's Compact Suffix Array

The *Compact Suffix Array* of Mäkinen [2000, 2003] (MAK-CSA) was aimed at explicitly exploiting the self-repetitions of the suffix array  $A$  in order to store it in a more compact form. It was conceived as a succinct index. Only later [Mäkinen and Navarro 2004a] it was shown that its size was related to the text entropy, thus becoming a compressed index. We describe here a version similar in spirit to the original proposal and equally efficient, yet somewhat cleaner.

The idea of the MAK-CSA is, essentially, to represent  $A$  as the minimal sequence of the  $n_{sr}$  self-repetitions covering  $A$  (recall Definition 9). Assume that the partition is  $A[i_1, i_1 + \ell_1]$ ,  $A[i_2, i_2 + \ell_2]$ , and so on, so that  $i_{s+1} = i_s + \ell_s + 1$ ,  $i_1 = 1$ ,  $i_{n_{sr}} + \ell_{n_{sr}} = n$ . Let  $j_s$  be so that  $A[j_s, j_s + \ell_s] = A[i_s, i_s + \ell_s] + 1$ . Let  $p_s$  be the partition number where  $j_s$  lies, that is,  $i_{p_s} \leq j_s < i_{p_s+1}$ , and let  $o_s = j_s - i_{p_s}$  be the corresponding offset. Then, the MAK-CSA for  $A$  is an array of  $n_{sr}$  blocks of the form  $B_s = (p_s, o_s, \ell_s, A[i_s])$ . That is, for each interval  $[i_s, i_s + \ell_s]$ ,  $B_s$  tells the partition where the interval repeats, the offset from the origin of that partition, the length of the repetition, and the explicit  $A$  value of the first cell of the block. Sequence  $B_s$  contains enough information to reconstruct  $A$ .

Fig. 9 illustrates the MAK-CSA for our example text. For example,  $B_9 = (8, 2, 1, 4)$  means that block 9 (that is,  $A[14, 15] = \{4, 16\}$ ) is obtained by copying 2 (1 + 1) cells starting at offset 2 of block 8 (and subtracting 1 from them), and it also records  $A[i_9] = A[14] = 4$ . Note that a pointer from a block may span several other blocks (such as  $B_8$ ), and that it may point at the middle of a block (as the same  $B_9$ ).




Fig. 9. The MAK-CSA of the text "alabar a la alabarda\$". We show the suffix array  $A$  and the MAK-CSA  $B$ . The minimal covering of self-repetitions is shown on the bottom of  $A$ , each region being an entry in  $B$ . We also show graphically some pointers related to self-repetitions.

The figure also illustrates partially how can we *decompress* the second cell of block  $B_{11}$ . The entry leads us to decompress the second cell of  $B_7$ , then the second cell of  $B_9$ , then the fourth cell of  $B_8$ , and finally the first cell of  $B_{13}$ . First cells are explicitly known, in this case its value is 18. Since we made 4 steps to reach  $B_{13}$ , the value we wanted is  $18 - 4 = 14$ . This path must finish at some moment because  $B_1$  is at the end of any path and it is of length 1.

Fig. 10 gives the pseudocode to decompress a cell value. The value  $i = i_s + \delta$  we wish to obtain is represented by *reference pair*  $(s, \delta)$ . Note that in lines 1–3 we

consider the possibility that  $\delta > \ell_s$ , in which case the pair  $(s, \delta)$  is said to be not *canonical*, so we advance in  $B$  until finding the correct canonical reference pair. This *canonization* process is necessary because, even if we ensure  $\delta \leq \ell_s$  in the first call, recursive calls with offset  $o_s + \delta$  may need canonization. Note that the search finishes when the offset  $\delta = 0$ , as we explicitly store all first cells of blocks.

---

**Algorithm** Mak-CSA-decompress( $s, \delta, B$ )

- (1) **while**  $\delta > \ell_s$  **do**
- (2)  $\delta \leftarrow \delta - \ell_s - 1$
- (3)  $s \leftarrow s + 1$
- (4) **if**  $\delta = 0$
- (5) **then return**  $A[i_s]$ ;
- (5) **else return** Mak-CSA-decompress( $p_s, o_s + \delta, B$ );

---

Fig. 10. Algorithm to obtain  $A[i_s + \delta]$  from the MAK-CSA structure.

Although one could implement a basic binary search (Fig. 4) using  $B$ , it is more clever to first binary search on the first cells of the blocks, which are explicitly stored. After  $O(\log n_{sr})$  steps, one knows that the answer lies within  $B_s$  for some  $s$ . At this point, one continues the binary search inside  $B_s$ , this time using algorithm Mak-CSA-decompress() to obtain each needed value in  $A[i_s, i_s + \ell_s]$ . Thus the total cost of the binary search is  $O(m \log n)$  as in the basic suffix array, plus that of decompressing some cells of a  $B$  block.

The time to decompress a cell is limited with two parameters  $C$  and  $D$ . The former is the maximum length of a block,  $\ell_s \leq C$ , and it is enforced by cutting longer self-repetitions into several ones. Parameter  $D$  limits the length of a chain of successive self-repetitions, so as to ensure that we will reach an explicit cell after  $D$  recursive invocations of Mak-CSA-decompress(). This is ensured by forcing explicit cells at the  $D$ -th position in any self-repetition path. With these parameters, the extra cost of the binary search inside a block is  $O(CD \log C)$  (the  $C$  factor comes from canonizing the references).

For locating queries, we must decompress several cells of  $B$  to obtain the whole interval  $A[sp, ep]$ . In the worst case this may require a full decompression for each of the  $occ$  occurrences, at  $O(CD \text{ occ})$  cost.

Given Theorem 3, it is clear that the size of the MAK-CSA is  $O(n_{sr} \log n) = O(nH_k \log n)$  bits, thus it is a compressed index. Yet, it is not a self-index, as it still needs  $T$  to operate. Forcing a maximum block length  $C$  introduces at most  $n/C$  extra blocks. The same occurs by cutting referencing paths at depth  $D$ : every new cut introduced correctly bounds  $D$  paths, so  $n/D$  cuts suffice. The following theorem assumes  $C = D = \log n / \log \log n$  so that the extra space associated to  $C$  and  $D$  is  $O(n \log \log n)$ , but other tradeoffs are possible. In addition, this makes each entry of  $B$  require only  $\log n + 2 \log \log n$  bits (as  $o_s$  and  $\ell_s$  are  $O(\log n)$ ).

**Theorem 9** *The Compact Suffix Array of Mäkinen [2000, 2003] (MAK-CSA) offers the following space/time tradeoff.*



Space in bits	$2nH_k \log n + O(n \log \log n) + n \log \sigma$ (for the text)
Time to count	$O(m \log n + \log^2 n / \log \log n)$
Time to locate	$O(\log^2 n / (\log \log n)^2)$
Time to display $\ell$ chars	$O(\ell)$ (text is available)

*In practice.* The implementation of MAK-CSA follows closely the description above, except that the blocks are decompressed in pieces taking advantage of the the common part of the search path of consecutive elements. This does not improve the worst case complexity, but is significant speedup in practice.

## 6.2 GV-CSA: Grossi and Vitter's Compressed Suffix Array

The *Compressed Suffix Array* of Grossi and Vitter [2000, 2005] (GV-CSA) is a succinct index based on the idea of providing efficient access to  $A[i]$  without representing  $A$ , so that the search algorithm is exactly as in Fig. 4. The text  $T$  is maintained explicitly.

The GV-CSA uses a hierarchical decomposition of  $A$  based on the  $\Psi$  function (Definition 11). Let us focus on the first level of that hierarchical decomposition. Let  $A_0 = A$  be the original suffix array. A bit vector  $B_0[1, n]$  is defined so that  $B_0[i] = 1$  iff  $A[i]$  is even. Let also  $\Psi_0[1, \lceil n/2 \rceil]$  contain the sequence of values  $\Psi(i)$  for values  $i$  where  $B_0[i] = 0$ . Finally, let  $A_1[1, \lfloor n/2 \rfloor]$  be the subsequence of  $A_0[1, n]$  formed by the even  $A_0[i]$  values, divided by 2.

Then,  $A = A_0$  can be represented using only  $\Psi_0$ ,  $B_0$ , and  $A_1$ . To retrieve  $A[i]$ , we first see if  $B_0[i] = 1$ . If it is, then  $A[i]$  is (divided by 2) somewhere in  $A_1$ . The exact position depends on how many 1's are there in  $B_0$  up to position  $i$ , that is,  $A[i] = 2 \cdot A_1[\text{rank}_1(B_0, i)]$ . If  $B_0[i] = 0$ , then  $A[i]$  is odd and not represented in  $A_1$ . However,  $A[i] + 1 = A[\Psi(i)]$  has to be even and thus represented in  $A_1$ . Since  $\Psi_0$  collects only the  $\Psi$  values where  $B_0[i] = 0$ , we have  $A[\Psi(i)] = A[\Psi_0[\text{rank}_0(B_0, i)]]$ . Once we compute  $A[\Psi(i)]$  (for even  $\Psi(i)$ ), we simply obtain  $A[i] = A[\Psi(i)] - 1$ .

Fig. 11 illustrates. For example, to obtain  $A[11]$ , we verify that  $B_0[11] = 0$ , thus it is not represented in  $A_1$ . So we need to obtain  $\Psi(11)$ , which is  $\Psi_0[\text{rank}_0(B_0, 11)] = \Psi_0[8] = 19$ . We must then obtain  $A[19]$ . Not surprisingly,  $B_0[19] = 1$ , so  $A[19]$  is at  $A_1[\text{rank}_1(B_0, 19)] = A_1[8] = 7$ . Thus  $A[19] = 2 \cdot 7 = 14$  and finally  $A[11] = A[19] - 1 = 13$ . (Note the exception that  $A[1]$  is odd and  $A[\Psi(1)] = A[10]$  is odd too, but this is not a problem because we know that  $A[1] = n$  always.)

The idea can be used recursively: Instead of representing  $A_1$ , we replace it with  $B_2$ ,  $\Psi_2$ , and  $A_2$ . This is continued until  $A_h$  is small enough to be represented explicitly. Fig. 12 gives the pseudocode to extract an entry  $A[i]$  from the recursive structure. The complexity is clearly  $O(h)$  assuming a constant-time implementation for  $\text{rank}$  (Section 5.1).

It is convenient to use  $h = \lceil \log \log n \rceil$ , so that the  $n/2^h$  entries of  $A_h$ , each of which requires  $O(\log n)$  bits, take overall  $O(n)$  bits. All the  $B_\ell$  arrays add up at most  $2n$  bits (as their length is halved from each level to the next), and their additional  $\text{rank}$  structures add  $o(n)$  extra bits (Section 5.1). The only remaining problem is how to represent the  $\Psi_\ell$  arrays.

For a compact representation of  $\Psi_0$ , we remind that  $\Psi$  is increasing within the area of  $A$  that points to suffixes starting with the same character (Lemma 1). Although Grossi and Vitter [2000] do not detail how to use this property to represent

$A_0 =$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
	21	7	12	9	20	11	8	3	15	1	13	5	17	4	16	19	10	2	14	6	18

$B_0 =$	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
	0	0	1	0	1	0	1	0	0	0	0	0	0	1	1	0	1	1	1	1	1

$\Psi_0 =$	1	2	3	4	5	6	7	8	9	10	11
	10	7	17	3	14	15	18	19	20	21	5

$A_1 =$	1	2	3	4	5	6	7	8	9	10
	6	10	4	2	8	5	1	7	3	9

Fig. 11. The first level of the GV-CSA recursive structure, for the text "alabar a la alabarda\$". We show the original suffix array  $A = A_0$  and structures  $B_0$ ,  $\Psi_0$ , and  $A_1$ . We show  $\Psi_0$  and  $A_1$  in scattered form to ease understanding but they are obviously stored contiguously. Recall that  $A_0$  is not really stored but replaced by the other three structures.

---

**Algorithm** GV-CSA-lookup( $i, \ell$ )

- (1) if  $\ell = h$  then return  $A_h[i]$ ;
  - (2) if  $B_\ell[i] = 1$
  - (3) then return  $2 \cdot \text{GV-CSA-lookup}(\text{rank}_1(B_\ell, i), \ell + 1)$ ;
  - (4) else return  $\text{GV-CSA-lookup}(\Psi_\ell[\text{rank}_0(B_\ell, i)], \ell) - 1$ ;
- 

Fig. 12. Algorithm to obtain  $A[i]$  from GV-CSA recursive structure with  $h$  levels. It is invoked as  $\text{GV-CSA-lookup}(i, 0)$ .

$\Psi$  in little space, an elegant solution is found in later work by Sadakane [2000, 2003]. Essentially, Sadakane shows that  $\Psi$  can be encoded differentially ( $\Psi(i+1) - \Psi(i)$ ) within the areas where it is increasing, using Elias- $\delta$  coding [Elias 1975; Witten et al. 1999]. The number of bits this representation requires is  $nH_0 + O(n \log \log \sigma)$ . For  $\Psi_0$ , since only odd text positions are considered, the result is the same as if we had a text  $T'_{1, n/2}$  formed by bigrams of  $T$ ,  $T'_i = T_{2i-1, 2i}$ . Since the zero-order entropy of  $T$  taken as a sequence of  $2^\ell$ -grams is  $H_0^{(2^\ell)} \leq 2^\ell H_0$  [Sadakane 2003],  $\Psi_0$  requires  $|T'|H_0^{(2)} + O(|T'| \log \log(\sigma^2)) \leq (n/2)(2H_0) + O((n/2)(1 + \log \log \sigma))$ . In general,  $\Psi_\ell$  requires  $(n/2^\ell)(2^\ell H_0) + O((n/2^\ell)(\ell + \log \log \sigma)) = nH_0 + O(n\ell/2^\ell) + O((n \log \log \sigma)/2^\ell)$  bits. Overall, the  $h$  levels require  $hnH_0 + O(n \log \log \sigma)$  bits.

In order to access the entries of these compressed  $\Psi_\ell$  arrays in constant time [Sadakane 2003], absolute  $\Psi$  values are inserted every  $\Theta(\log n)$  bits, so this adds  $O(n)$  bits. To extract an arbitrary position of  $\Psi$ , we go to the nearest absolute sample before that position and then sequentially advance summing up differences until reaching the desired position. By maintaining a precomputed table with the total number of differences encoded inside every possible chunk of  $\frac{\log n}{2}$  bits, we can process each such chunk in constant time, so the  $\Theta(\log n)$  bits of differences can be processed in constant time. The size of that table is only  $O(\sqrt{n} \log^2 n) = o(n)$  bits. Note the similarity with the other four-Russians technique for constant time *rank* (Section 5.1).

What we have, overall, is a structure using  $nH_0 \log \log n + O(n \log \log \sigma)$  bits of space, which encodes  $A$  and permits retrieving  $A[i]$  in  $O(\log \log n)$  time.

A tradeoff with  $\frac{1}{\epsilon} nH_0 + O(n \log \log \sigma)$  bits of space and  $O(\log^\epsilon n)$  retrieval time,

for any constant  $0 < \epsilon < 1$ , can be obtained as follows. Given the  $h = \lceil \log \log n \rceil$  levels, we only keep three of them: level 0, level  $\lfloor h/2 \rfloor$ , and level  $h$ . Bit vectors  $B_0$  and  $B_{\lfloor h/2 \rfloor}$  indicate which entries are represented in levels  $\lfloor h/2 \rfloor$  and  $h$ , respectively. The space for  $\Psi_0$ ,  $\Psi_{\lfloor h/2 \rfloor}$ , and  $A_h$ , is only  $2nH_0 + O(n \log \log \sigma)$  bits. However, we cannot move from one level to the next in constant time. We must use  $\Psi_\ell$  several times until reaching an entry that is sampled at the next level. The number of times we must use  $\Psi_\ell$  is at most  $2^{h/2} = O(\sqrt{\log n})$ . If, instead of 3 levels, we use a constant number  $1 + 1/\epsilon$  of levels 0,  $h\epsilon$ ,  $2h\epsilon$ ,  $\dots$ ,  $h$ , the time is  $O(\log^\epsilon n)$ . By using the search algorithm of Fig. 4 and the usual way to locate occurrences with  $A$ , we get the following results.

**Theorem 10** *The Compressed Suffix Array of Grossi and Vitter [2000] (GV-CSA), in the variant proposed by Sadakane [2000, 2003], offers the following space/time tradeoffs.*

Space in bits	$nH_0 \log \log n + O(n \log \log \sigma) + n \log \sigma$ (for the text)
Time to count	$O(m \log n \log \log n)$
Time to locate	$O(\log \log n)$
Time to display $\ell$ chars	$O(\ell)$ (text is available)
Space in bits	$\frac{1}{\epsilon} nH_0 + O(n \log \log \sigma) + n \log \sigma$ (for the text)
Time to count	$O(m \log^{1+\epsilon} n)$
Time to locate	$O(\log^\epsilon n)$
Time to display $\ell$ chars	$O(\ell)$ (text is available)
Conditions for all	$0 < \epsilon \leq 1$ is an arbitrary constant

We have described the solution of Sadakane [2000, 2003] to represent  $\Psi$  in little space and constant access time. The solution of the original authors is about to appear [Grossi and Vitter 2005] and it is slightly different. They also use the fact that  $\Psi$  is piecewise increasing in a different way, achieving  $\frac{1}{2}n \log \sigma$  bits at each level instead of  $nH_0$ . Furthermore, they take  $T$  as a binary string of  $n \log \sigma$  bits, which yields essentially  $n \log \sigma \log \log_\sigma n$  bits for the first version and  $(1 + 1/\epsilon)n \log \sigma$  bits for the second version of the theorem. They actually use  $h = \lceil \log \log_\sigma n \rceil$ , which adds up  $n \log \sigma$  extra space for  $A_h$  and slightly reduces the time to access  $A[i]$  in the first variant of the above theorem (to  $O(h)$ ).

Grossi and Vitter [2000, 2005] show how the *occ* occurrences can be located more efficiently in batch when  $m$  is large enough. They also show how to modify a compressed suffix tree [Munro et al. 2001] so as to obtain  $O(m/\log_\sigma n + \log^\epsilon n)$  search time, for any constant  $0 < \epsilon < 1$ , using  $O(n \log \sigma)$  bits of space. This is obtained by modifying the compressed suffix tree [Munro et al. 2001] in two ways: First, using perfect hashing to allow traversing  $O(\log_\sigma n)$  tree nodes downwards in one step, and second, replacing the suffix array required by the compressed suffix tree with the GV-CSA. We do not abound in details because in this paper we are more interested in indexes smaller than  $O(n \log \sigma)$  bits. In this sense, we are not interested in the GV-CSA by itself, but as a predecessor of other self-indexes that appeared later.

A generalization of this structure (but still not a self-index) is presented by Rao [2002], where they index a binary text using  $O(nt \log^{1/t} n)$  bits and retrieve  $A[i]$  in

$O(t)$  time, for any  $1 \leq t \leq \log \log n$ .

## 7. TURNING COMPRESSED SUFFIX ARRAYS INTO SELF-INDEXES

Further development over the compressed suffix arrays of Section 6 lead to self-indexes, which can operate without the text. The idea is to replace  $T$  by an abstract data type that gives access to any substring of it. The first index in this line was the *Compressed Suffix Array* of Sadakane [2000, 2002, 2003] (SAD-CSA). This was followed by the *Compressed Suffix Array* of Grossi, Gupta, and Vitter [2003, 2004] (GGV-CSA), and by the *Compressed Compact Suffix Array* of Mäkinen and Navarro [2004a] (MN-CCSA).

### 7.1 Sad-CSA: Sadakane's Compressed Suffix Array

Sadakane [2000, 2003] showed how the GV-CSA can be converted into a self-index, and at the same time optimized it in several ways. The resulting index was also called *Compressed Suffix Array* and will be referred to as SAD-CSA in this paper.

The SAD-CSA does not give, as the GV-CSA, direct access to  $A[i]$ , but rather to any prefix of  $T_{A[i],n}$ . With this value one can still use the basic search algorithm of Fig. 4. The SAD-CSA represents  $A$  and  $T$  using the full function  $\Psi$  (Definition 11), and a special representation of function  $C$  (Lemma 3).

Imagine we wish to compare  $P$  against  $T_{A[i],n}$ . For the binary search, we need to extract enough characters from  $T_{A[i],n}$  so that its lexicographical relation to  $P$  is clear. Since  $T_{A[i]}$  is the first character of the suffix pointed to by  $A[i]$  (or, alternatively,  $T_{A[i]} = F[i]$  in Fig. 6), it is the character  $c$  such that  $C(c) < i \leq C(c+1)$  (assuming  $C(\sigma+1) = n$ ). Once we determine  $T_{A[i]} = c$  in this way, we need to obtain the next character,  $T_{A[i+1]}$ . But  $T_{A[i+1]} = T_{A[\Psi(i)]}$ , so we simply move to  $i' = \Psi(i)$  and keep extracting characters with the same method, as long as necessary. Note that at most  $m$  characters suffice to decide any comparison to  $P$ .

In order to find quickly the  $c$  such that  $C(c) < i \leq C(c+1)$ , we represent  $C$  with a bit vector  $D[1, n]$ , so that  $D[i+1] = 1$  iff  $C(c) = i < n$  for some  $c \in [1, \sigma]$ , and a string  $S$  where the (at most  $\sigma$ ) distinct characters of  $T$  are concatenated in alphabetical order. Therefore, the  $c$  such that  $C(c) < i \leq C(c+1)$  is precisely  $c = S[\text{rank}_1(D, i)]$ . Using the succinct data structures of Section 5.1 we can compute  $c$  in constant space using only  $n + o(n)$  bits for  $D$  and  $\sigma \log \sigma$  bits for  $S$ .

Fig. 13 illustrates. To obtain, say,  $T_{A[11],n}$  we see that  $S[\text{rank}_1(D, 11)] = S[3] = \text{"a"}$ . Then we move to  $19 = \Psi(11)$ . The second character is thus  $S[\text{rank}_1(D, 19)] = S[6] = \text{"1"}$ . Then we move to  $9 = \Psi(19)$  and get third character  $S[\text{rank}_1(D, 9)] = S[3] = \text{"a"}$ , and so on. Note that we are, implicitly, walking the text in forward direction. Note also that we do not know where we are in the text, that is, we never know  $A[i]$ , just  $T_{A[i],n}$ .

Thus the SAD-CSA implements the binary search in  $O(m \log n)$  worst-case time, which is better than in the GV-CSA structure. Fig. 14 gives the pseudocode to compare  $P$  against a suffix of  $T$ .

Right now we have used  $n + o(n) + \sigma \log \sigma$  bits of space for  $D$  and  $S$ , plus the representation for  $\Psi$  described in Section 6.2,  $nH_0 + O(n \log \log \sigma)$  bits. Note that, since the SAD-CSA does not give direct access to  $A[i]$ , it needs more structures to solve a locating query. That is, although the index knows that the answers are in  $A[sp, ep]$  and thus that there are  $occ = ep - sp + 1$  answers, it does not have

$A =$	21	7	12	9	20	11	8	3	15	1	13	5	17	4	16	19	10	2	14	6	18
-------	----	---	----	---	----	----	---	---	----	---	----	---	----	---	----	----	----	---	----	---	----

$\Psi =$	10	7	11	17	1	3	4	14	15	18	19	20	21	12	13	5	6	8	9	2	16
----------	----	---	----	----	---	---	---	----	----	----	----	----	----	----	----	---	---	---	---	---	----

$D =$	1	1	0	0	1	0	0	0	0	0	0	0	0	1	0	1	1	0	0	1	0
-------	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---

$S =$  **\$\_abdlr**

$T =$  a l a b a r \_ a \_ l a \_ a l a b a r d a \$

$C(\$) = 0$     $C(d) = 15$

$C(_) = 1$     $C(l) = 16$

$C(a) = 4$     $C(r) = 19$

$C(b) = 13$

Fig. 13. The main components of the SAD-CSA structure, for the text "alabar a la alabarda\$". The text  $T$ , the original suffix array  $A$ , and function  $C$ , are shown for illustrative purposes but are not represented in the structure.

---

**Algorithm** Sad-CSA-compare( $P, m, i, \Psi, D, S$ )

- (1)  $j \leftarrow 1$ ;
  - (2) **do**  $c \leftarrow S[\text{rank}_1(D, i)]$ ;
  - (3)     **if**  $P_j < c$  **then return** "<";
  - (4)     **if**  $P_j > c$  **then return** ">";
  - (5)      $i \leftarrow \Psi(i)$ ;
  - (6)      $j \leftarrow j + 1$ ;
  - (7) **while**  $j \leq m$ ;
  - (8) **return** "=";
- 

Fig. 14. Algorithm to compare  $P$  against  $T_{A[i], A[i]+m-1}$ , where  $T$  and  $A$  are represented by  $\Psi$ ,  $D$ , and  $S$ .

enough information to know the text positions pointed to by  $A[i]$ ,  $sp \leq i \leq ep$ . For this sake, the SAD-CSA includes the hierarchical GV-CSA structure (without the text and with  $\Psi$  instead of  $\Psi_0$ , as we already have the more complete  $\Psi$ ). We use the version requiring  $\frac{1}{\epsilon} nH_0 + O(n \log \log \sigma)$  bits of space and computing  $A[i]$  in  $O(\log^\epsilon n)$  time, for any constant  $0 < \epsilon < 1$  (Theorem 10).

Sadakane, however, notices an important fact. The same hierarchical GV-CSA structure can serve to compute the inverse of  $A$ ,  $A^{-1}[j]$ , within the same  $O(\log^\epsilon n)$  time, provided we store explicitly  $A_h^{-1}$  in the last level (in addition to the explicit  $A_h$  stored in the GV-CSA structure). The reason is that, if one knows that  $A^{-1}[j-1] = i$ , then  $A[i] = j-1$ , and therefore  $A[\Psi(i)] = j$ , thus  $A^{-1}[j] = \Psi(i) = \Psi(A^{-1}[j-1])$ . Hence, to obtain  $A_\ell^{-1}[j]$ , we take the largest  $j' < j$  that is represented in the next level,  $j' = \text{select}_1(B_\ell, \text{rank}_1(B_\ell, j-1))$ , and obtain  $i' = A_\ell^{-1}[j']$ . Now we apply  $\Psi$  successively over  $i'$  to obtain  $i = \Psi^{j-j'}(i')$ .

The inverse of  $A$  is useful to implement the additional function that a self-index must provide: given  $1 \leq l \leq r \leq n$ , retrieve  $T_{l,r}$  without having access to  $T$ . We already know how to retrieve  $T_{A[i],n}$  (or any prefix of it) given  $i$ . With the inverse suffix array we first find  $i = A^{-1}[l]$  in  $O(\log^\epsilon n)$  time, and then retrieve  $T_{l,r} = T_{A[i], A[i]+(r-l)}$  in  $O(r-l)$  time.

**Theorem 11** *The Compressed Suffix Array of Sadakane [2000, 2003] (SAD-CSA) offers the following space/time tradeoff.*

Space in bits	$\frac{1}{\epsilon} nH_0 + O(n \log \log \sigma) + \sigma \log \sigma$
Time to count	$O(m \log n)$
Time to locate	$O(\log^\epsilon n)$
Time to display $\ell$ chars	$O(\ell + \log^\epsilon n)$
Conditions	$0 < \epsilon \leq 1$ is an arbitrary constant

It is interesting that not only a local portion of  $T$  can be decompressed, but also the suffix array of such a local segment of  $T$  can be retrieved using  $A^{-1}$  and  $\Psi$ .

We note that the method of inserting absolute samples every  $\Theta(\log n)$  positions in  $\Psi$  and using four-Russian techniques to access between samples in constant time, works with many other compression methods. In particular, if we compress runs in  $\Psi$  (Definition 12) with run-length compression, we can achieve  $nH_k(\log \sigma + \log \log n) + O(n)$  bits of space, for  $k \leq \alpha \log_\sigma n$  and any constant  $0 < \alpha < 1$ , while retaining the same search times [Mäkinen and Navarro 2004b] (recall Theorem 3).

*In practice.* The real implementation of SAD-CSA differs in several aspects from its theoretical description. First, it does not implement the inverse suffix array to locate occurrences. Rather, it samples  $A$  at regular intervals of length  $d$ , explicitly storing  $A[id]$  for all  $i$ . In order to obtain  $A[j]$ , we compute  $\Psi$  repeatedly over  $j$  until obtaining a value  $j' = \Psi^r(j)$  multiple of  $d$ . Then  $A[j] = A[j'] - r$ . Similarly, constant access to  $\Psi$  is not provided. Rather, absolute  $\Psi$  values are sampled every  $L$  positions. To obtain  $\Psi(i)$ , we start from its closest previous absolute sample and decompress the differential encoding until position  $i$ . Finally, instead of the classical suffix array searching, backward searching is used [Sadakane 2002]. This avoids any need to obtain text substrings, and it is described in Section 8.2. Thus,  $d$  and  $L$  give practical tradeoffs between index space and search time.

## 7.2 MN-CCSA: Mäkinen and Navarro's Compressed Compact Suffix Array

The *Compressed Compact Suffix Array* of Mäkinen and Navarro [2004a] (MN-CCSA) is a simple self-index using essentially the structures of the MAK-CSA of Section 6.1 to implement the search method of the SAD-CSA of Section 7.1. Its main idea is twofold: (1) make the MAK-CSA a self-index, and (2) reduce its space by replacing some of its many pointers by smarter data structures based on *rank* and *select* on bit arrays (Section 5.1).

The MN-CCSA uses the basic structure of the MAK-CSA, with the slight difference that all the suffixes belonging to a single block in the MN-CCSA must start with the same character, otherwise the block must be split. This produces at most  $\sigma$  extra splits from the optimal covering.

Just as done in Section 7.1, the text is deleted and replaced by bit array  $D$  and string  $S$ . Yet, since all suffixes in a block start with the same character, the length of  $D$  is just  $n_{sr}$ , with one bit per block. Then, with an equivalent of function  $\Psi$  we can retrieve any text string of the form  $T_{A[i],n}$  given  $i$ . In this case we must keep track all the time of the block number we are at, not only the position in  $A$ .

The equivalent of function  $\Psi$  is given by array  $B_s = (p_s, o_s, \ell_s, A[i_s])$  of Section 6.1. Assume cell  $i$  of  $A$  corresponds to offset  $\delta$  in block  $s$ , that is,  $i = i_s + \delta$

and  $0 \leq \delta \leq \ell_s$ . Then  $i$  is represented by the reference pair  $(s, \delta)$ , and  $\Psi(i)$  is represented by reference pair  $(p_s, o_s + \delta)$ , yet this pair is not necessarily canonical. The MAK-CSA lost  $O(C)$  time in canonizing each reference pair before continuing.

The MN-CCSA uses a more compact representation that in addition performs canonization in constant time. Assume we have a bit vector  $L[1, n]$  aligned to  $A$ , where the block starts are signaled with a bit set. In addition to  $L$ , instead of storing sequence  $B_s = (p_s, o_s, \ell_s, A[i_s])$  we store the sequence  $j_s$  of absolute positions in  $A$  where the copy of each block  $s$  starts. That is,  $j_s$  is the number represented by the reference pair  $(p_s, o_s)$ . From  $j_s$  and  $L$  we can obtain other components in constant time:  $p_s = \text{rank}_1(L, j_s)$ ,  $o_s = j_s - \text{select}_1(L, p_s) - 1$ , and  $\ell_s = \text{select}_1(L, p_s + 1) - \text{select}_1(L, p_s)$ . Furthermore, if current position  $i$  is represented by reference pair  $(s, \delta)$ , then  $\Psi(i)$  corresponds to canonical reference pair  $(p_s, \delta')$ , where  $p_s = \text{rank}_1(L, j_s + \delta)$  and  $\delta' = j_s + \delta - \text{select}_1(L, p_s)$ .

This permits using the same binary search of the SAD-CSA, using the same technique to extract text substrings to compare against  $P$ . (So we do not give explicit access to  $A[i]$  as in the MAK-CSA.) Fig. 15 gives the pseudocode.

---

**Algorithm** MN-CCSA-compare( $P, m, i, L, J, D, S$ )

```

(1)  $j \leftarrow 1$ ;
(2) do  $s \leftarrow \text{rank}_1(L, i)$ ;
(3)    $\delta \leftarrow i - \text{select}_1(L, s)$ ;
(4)    $c \leftarrow S[\text{rank}_1(D, s)]$ ;
(5)   if  $P_j < c$  then return “<”;
(6)   if  $P_j > c$  then return “>”;
(7)    $i \leftarrow J[s] + \delta$ ;
(6)    $j \leftarrow j + 1$ ;
(7) while  $j \leq m$ ;
(8) return “=”;
```

---

Fig. 15. Algorithm to compare  $P$  against  $T_{A[i], A[i]+m-1}$ , where  $T$  and  $A$  are represented by  $L$ ,  $J$ ,  $D$ , and  $S$ . Array  $J[s]$  stores sequence  $j_s$ .

Array  $j_s$  takes  $n_{sr} \log n$  bits, while  $L$  takes  $n + o(n)$  bits using the techniques of Section 5.1. Likewise, array  $D$  takes  $n_{sr}(1 + o(1))$  bits. Adding up all the space complexities and recalling Theorem 3, we have overall  $n(H_k(1 + \log n) + 1) + o(n)$  bits.

In addition, the MN-CCSA uses  $O(n)$  extra bits to locate occurrences and display text contexts. The technique resembles that used for the SAD-CSA in practice (and even more those in the FM-Index [Ferragina and Manzini 2000], Section 8.3). Text positions are sampled at regular intervals  $h = 2\epsilon \log n$ , and the positions in  $A$  that point to  $T_{1,n}$ ,  $T_{h+1,n}$ ,  $T_{2h+1,n}$ , and so on, are sorted and stored in array  $S_p$  (that is, obtain  $A^{-1}[1]$ ,  $A^{-1}[h+1]$ , etc., sort them, and store the result in  $S_p$ ). A bit array  $\text{in}S_p[1, n]$ , aligned to  $A$ , tells which positions in  $A$  turned out to be sampled in  $S_p$ . Then, to find  $A[i]$ , we take  $\Psi(i)$  zero or more times until finding the first  $r$  such that  $\text{in}S_p[\Psi^r(i)] = 1$ , at which point we have  $A[i] = S_p[\text{rank}_1(\text{in}S_p, \Psi^r(i))] - r$ . This costs  $O(\epsilon \log n)$  at worst. To show text contexts, the same  $A$  positions of  $S_p$  are stored in another array  $S_t$ , now sorted by their order in  $T$  (that is,  $S_t[s] = A^{-1}[hs + 1]$ ). To display  $T_{i,r}$  we compute its

immediately preceding sample  $s = \lfloor \frac{l-1}{h} \rfloor$ , and rather obtain  $T_{hs+1,r} = T_{A[S_t[s]],r}$ , so we can use the already known method to extract substrings of the form  $T_{A[i],n}$ . Those structures take  $n(1 + 1/\epsilon) + o(n)$  extra bits.

**Theorem 12** *The Compressed Compact Suffix Array of Mäkinen and Navarro [2004a] (MN-CCSA) offers the following space/time tradeoff.*

Space in bits	$n(H_k(1 + \log n) + 2 + \frac{1}{\epsilon}) + O(n \log \log n / \log_\sigma n)$
Time to count	$O(m \log n)$
Time to locate	$O(\epsilon \log n)$
Time to display $\ell$ chars	$O(\ell + \epsilon \log n)$
Conditions	$\epsilon > 0$ is an arbitrary constant; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

*In practice.* The MN-CCSA is implemented as described, so it has a major difference with respect to the implementation of the SAD-CSA. The MN-CCSA uses the usual suffix array search method, while the SAD-CSA uses backward searching as explained at the end of Section 7.1. This favors MN-CCSA for large  $m$  because the  $O(m \log n)$  worst case complexity is no more than  $O(\log^2 n)$  on average, while the worst and average case complexities are both  $O(m \log n)$  with backward searching.

### 7.3 GGV-CSA: Grossi, Gupta, and Vitter's Compressed Suffix Array

The *Compressed Suffix Array* of Grossi, Gupta, and Vitter [2003, 2004] (GGV-CSA) is an evolution over the GV-CSA and the SAD-CSA. It is a self-index whose space usage depends on the  $k$ -th order entropy of  $T$ . The result is based on a new representation of  $\Psi$  that requires essentially  $nH_k$  bits rather than  $nH_0$ , basically using Theorem 2.

Consider Fig. 3 and 13, and the text context  $s = \text{"1a"}$ . Its occurrences are pointed from  $A[17, 19] = \{10, 2, 14\}$ . The  $\Psi$  values that point to that interval are  $\Psi(4) = 17$ ,  $\Psi(10) = 18$ , and  $\Psi(11) = 19$ . The first argument, 4, corresponds to character  $T_{A[4]} = \text{" "}$  preceding "1a", while the other two correspond to "a" preceding "1a".

Fig. 16 illustrates how  $\Psi$  is partitioned into *lists* ( $\Psi(i)$  belongs to list  $T_{A[i]}$ , that is, the character preceding position  $\Psi(i)$  in  $T$ ) and *contexts* ( $\Psi(i)$  belongs to context  $T_{A[\Psi(i)], A[\Psi(i)]+k-1} = T_{A[i+1], A[i+k]}$ , that is, the text starting at  $\Psi(i)$  in  $T$ ). If we sort the sequence of  $\Psi$  values in the table by list, and by context inside each list, then we have the entries sorted by  $T_{A[i], A[i+k]}$ , that is, an order compatible to that of the suffix array. Thus we recover the original sequence  $\Psi$  if we read the table column by column (left to right), and each column top to bottom (as we ordered lists and contexts lexicographically). We respect the original  $\Psi$  order within cells.

The  $\Psi$  values along each list (column) are increasing because  $\Psi$  is increasing in the area of  $A$  that starts with the same character (Lemma 1). Similarly,  $\Psi$  is increasing inside each particular cell.

Consider now a specific row (corresponding to context  $s$ ). All the  $\Psi$  values found there correspond to the initial positions of occurrences of context  $s$  in  $T$ . Thus, all the  $\Psi$  values found in such row form a contiguous subinterval if  $[1, n]$ . Each cell in the row corresponds to a different character preceding the context, and values



Context	List "\$"	List " "	List "a"	List "b"	List "d"	List "l"	List "r"
"\$"			1,				
" a"			3,				2,
" l"			4,				
"a\$"					5,		
"a "		7,				6,	
"ab"	10,					8, 9,	
"a l"		11,					
"ar"				12, 13,			
"ba"			14, 15,				
"da"							16,
"la"		17,	18, 19,				
"r "			20,				
"rd"			21,				

Fig. 16. Partition of  $\Psi$  into lists (columns) and contexts (rows) in the GGV-CSA, for the text "alabar a la alabarda\$". Array  $\Psi$  is read from leftmost to rightmost column, each column top to bottom.

are increasing inside each cell. If we identify each  $\Psi$  value in the row with the character of the column (list) it belongs to, then the set of all characters form precisely  $T_s$ . Thus, if we manage to encode each row in  $|T_s|H_0(T_s)$  bits, we will have  $nH_k$  bits overall (recall Eq. (1) and Theorem 2). In our previous example, considering context  $s = "la"$  ( $k = 2$ ), we have to represent all the  $\Psi$  values that lie inside that context ([17, 19]) in space proportional to the zero-order entropy of the first characters of the positions in  $\Psi$  that point there (in this case,  $T_{A[4]}$ ,  $T_{A[10]}$ , and  $T_{A[11]}$ ), overall  $T_s = T_{17,19}^{bwt} = "aa"$ .

To obtain  $\Psi(i)$  from this table we first need to determine the row and column  $i$  belongs to, and then the position inside that cell. To know the column, bitmap  $D[1, n]$  of Fig. 13 (corresponding to  $G_k$  of Grossi et al. [2003]) suffices, as  $c = \text{rank}_1(D, i)$ . Using the techniques of Section 5.2,  $D$  can be represented in  $nH_0(D) + o(n) \leq \sigma \log n + o(n)$  bits, answering  $\text{rank}_1$  in constant time. The position of  $i$  inside column  $c$  is  $i' = i - \text{select}_1(D, c) + 1$ . In the example, to retrieve  $\Psi(10)$  (thus  $i = 10$ ), we find  $c = \text{rank}_1(D, 10) = 3$  (third column in the table). Inside the table, we want the 6th value, because  $i' = 10 - \text{select}_1(D, 3) + 1 = 10 - 5 + 1 = 6$ .

A similar technique gives the right cell within the column. A bit array  $L_c$  is maintained for each column  $c$  (this corresponds to  $L_k^y$  of Grossi et al. [2003]).  $L_c$  is aligned to the area of  $A$  where suffixes start with the  $c$ -th character in the alphabet of  $T$ .  $L_c$  contains a 1 every time a new context of length  $k$  appears in column  $c$ . In our example,  $L_3[1, n] = 111101011$ . It is easy to see that  $i'$  corresponds to the  $j'$ -th nonempty cell in column  $c$ , where  $j' = \text{rank}_1(L_c, i')$ . In our example, the cell is at the 5th nonempty row, as  $j' = \text{rank}_1(L_3, 6) = 5$ . Finally, the position we want inside the  $j'$ -th column is  $p = i' - \text{select}_1(L_c, j') + 1$ . In our example this is the first position of the cell, as  $p = 6 - \text{select}_1(L_3, 5) + 1 = 1$ . Using again the techniques from Section 5.2, each  $L_c$  can be stored in  $n_c H_0(L_c) + o(n_c)$  bits (where  $n_c$  is the number of elements in column  $c$ ), and answer those queries in constant time. As there are at most  $\sigma^k$  bits set in  $L_c$ , the space is at most  $\sigma^k \log(n_c) + o(n_c)$  bits. Added over all columns, this is at most  $\sigma^{k+1} \log n + o(n)$ . In the real structure [Grossi et al. 2003], all vectors  $L_c$  are concatenated for technical reasons we omit.

Finally, we are inside a cell and know which position we seek inside it. We need to know which is the first value of the range handled by the row. For example, for  $s = \text{"1a"}$ , row 11, we must know that the interval  $[17, 19]$  starts with 17. We store a bitmap  $F$  (corresponding to  $F_k$  of Grossi et al. [2003]) whose positions are aligned to  $A$ , storing a 1 each time a context change occurs while traversing  $A$ . This is the global version of  $L_c$  bit vectors. In our example,  $F = 110111010011010110011$ . If we know we are in global row  $j$ , then  $\text{select}_1(F, j)$  tells the first element in the interval handled by row  $j$ . In our example,  $j = 11$  and  $\text{select}_1(F, 11) = 17$ . We will add this value to the result we obtain inside our cell. Using again Section 5.2,  $F$  requires  $\sigma^k \log n + o(n)$  bits.

Yet, we know only  $j' = 5$ , not  $j = 11$ . To map each local row number ( $j'$ -th nonempty row at column  $c$ ) to global row number  $j$ , we store a bitmap  $E_c$  for each column  $c$  with one bit per context indicating which are the existing contexts in column  $c$  (this corresponds to  $b_k^y$  of Grossi et al. [2003]). The global row number is  $j = \text{select}_1(E_c, j')$ . In our example,  $E_3 = 1110000010111$  and  $j = \text{select}_1(E_3, 5) = 11$ . All the  $E_c$  vectors, using Section 5.1, require  $\sigma^{k+1}(1 + o(1))$  bits of space.

The final piece is to obtain the  $p$ -th element of cell  $(j, c)$  (or locally, the  $j'$ -th nonempty cell at column  $c$ ). At this point there are different choices. One, leading to Theorem 13, is to store a bitmap  $Z$  for each cell, indicating which elements of the row interval belong to the cell. In our example, for row 11, the interval is  $[17, 19]$ , and we have  $Z = 100$  at column 2 and  $Z = 011$  at column 3 (zeros and ones can be interleaved in general). With  $\text{select}_1(Z, p)$  we obtain the offset of our element in the row interval. Therefore, we finally have  $\Psi(i) = \text{select}_1(F, j) + \text{select}_1(Z, p) - 1$ . In our example,  $\text{select}_1(011, 1) = 2$ , and therefore  $\Psi(10) = 17 + 2 - 1 = 18$ . Fig. 17 gives the pseudocode.

---

**Algorithm** GGV-CSA- $\Psi(i, D, L, E, F, Z)$

- (1)  $c \leftarrow \text{rank}_1(D, i)$ ;
  - (2)  $i' \leftarrow i - \text{select}_1(D, c) + 1$ ;
  - (3)  $j' \leftarrow \text{rank}_1(L_c, i')$ ;
  - (4)  $p \leftarrow i' - \text{select}_1(L_c, j') + 1$ ;
  - (5)  $j \leftarrow \text{select}_1(E_c, j')$ ;
  - (6) **return**  $\text{select}_1(F, j) + \text{select}_1(Z_{j,c}, p) - 1$ ;
- 

Fig. 17. Algorithm to compute  $\Psi(i)$  using the GGV-CSA.

We consider now the space required by the bit vectors  $Z$ . Using again the techniques of Section 5.2, those can be stored in  $|Z|H_0(Z) + o(|Z|)$  bits (note that  $|Z| = |T_s|$  for each cell in the row corresponding to context  $s$ ). Summed over all the lists of the same row, this turns out to be  $|T_s|H_0(T_s) + |T_s| \log e + o(|T_s|)$  [Grossi et al. 2003]. Added over all the contexts, we get  $nH_k + n \log e + o(n)$  (recall Eq. (1) and Theorem 2).

If we add up all the space requirements, we have  $nH_k + n \log e + O(\sigma^{k+1} \log n) + o(n)$ . For  $k + 1 \leq \alpha \log_\sigma n$ , for any constant  $0 < \alpha < 1$  of our choice, the space is  $nH_k + n \log e + o(n)$ .

Remember that we are not representing  $\Psi$ , but  $\Psi_\ell$  for  $0 \leq \ell \leq h$  (Section 6.2). The structure above works verbatim for  $\Psi_0$ , but it also can be used for any level

$\ell$ . The difference is that at level  $\ell$  there are not  $\sigma$  lists, but rather  $\sigma^{2^\ell}$ . The space requirement at level  $\ell$  turns out to be  $nH_k + (n/2^\ell) \log e + O(\sigma^{k+2^\ell} \log n) + o(n)$  bits (remember the space analysis in Section 6.2). To maintain the third term within  $o(n)$ , we must have  $k + 2^\ell \leq \alpha \log_\sigma n$ , for some constant  $0 < \alpha < 1$  of our choice. If we have  $k \leq \alpha' \log_\sigma n$ , then any  $\alpha' < \alpha < 1$  is good. Moreover, let  $h' = \log((\alpha - \alpha') \log_\sigma n) = \Theta(\log \log_\sigma n)$ , so that  $2^\ell \leq (\alpha - \alpha') \log_\sigma n$  when  $\ell \leq h'$ . This level  $h'$  is where the recursive structure must finish to maintain the extra space within  $o(n)$ . The levels between  $h'$  and  $h = \lceil \log \log n \rceil$  (where we store  $A_h$  explicitly) must be omitted, and this means that we must jump directly from level  $h'$  to level  $h$ , just as when we used a constant number of levels in Theorem 10. The number of steps for that jump is not constant but  $2^{h-h'} = O(\log \sigma)$ .

As we can access each  $\Psi_\ell$  in constant time, we first pay  $O(h')$  time to reach level  $h'$  and then pay  $O(\log \sigma)$  to reach level  $h$ . This is  $O(\log \log_\sigma n + \log \sigma) = O(\log \log n + \log \sigma)$  time to access  $A[i]$ . For the space, we have  $h'$  levels taking  $nH_k + (n/2^\ell) \log e + o(n)$  each, plus the final level containing  $A$  and its inverse. This yields the first result of Grossi et al. [2003].

**Theorem 13** *The Compressed Suffix Array of Grossi, Gupta, and Vitter [2003] (GGV-CSA) offers the following space/time tradeoffs.*

Space in bits	$nH_k \log \log_\sigma n + 2(\log e + 1)n + O(n \log \log n / \log n)$
Time to count	$O(\log n(m / \log_\sigma n + \log \log n + \log \sigma))$
Time to locate	$O(\log \log n + \log \sigma)$
Time to display $\ell$ chars	$O(\ell / \log_\sigma n + \log \log n + \log \sigma)$
Space in bits	$\frac{1}{\epsilon} nH_k + 2(\log e + 1)n + O(n \log \log n / \log n)$
Time to count	$O(\log n(m / \log_\sigma n + \log^\epsilon n + \log \sigma))$
Time to locate	$O(\log^\epsilon n + \log \sigma)$
Time to display $\ell$ chars	$O(\ell / \log_\sigma n + \log^\epsilon n + \log \sigma)$
Conditions for all	$0 < \epsilon \leq 1$ is an arbitrary constant; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

Note that, compared to Theorem 11,  $m$  has been replaced by  $m / \log_\sigma n$ , and likewise  $\ell$  by  $\ell / \log_\sigma n$ . This is because they note that the process carried out by the SAD-CSA to extract the text using  $\Psi$  character by character can actually be carried out at  $\Psi_{h'}$ , where  $2^{h'} = O(\log_\sigma n)$  characters are obtained in each step. Thus the counting time is  $O(m \log \sigma + \text{polylog}(n))$ . The  $O(\log n)$  factor multiplying  $m$  in previous approaches becomes now  $O(\log \sigma)$ , although new polylogarithmic terms in  $n$  appear. On the other hand, the version using  $\frac{1}{\epsilon} nH_k$  bits is obtained just as with the GV-CSA, using a constant number  $1 + 1/\epsilon$  of levels in  $[0, h']$ .

We observe that there is still an  $O(n)$  term in the space complexity, which is possible to remove. The extra space appears because we are coding each cell individually, indicating in its  $Z$  vector which elements of the row interval belong to the cell and which do not. Summed over all the row, the positive information amounts to  $|T_s|H_0(T_s) + o(|T_s|)$  bits, but the negative information amounts to the extra  $O(|T_s|)$  bits.

An alternative representation for the row is a wavelet tree (Section 5.3, invented by Grossi et al. [2003] for this purpose). The idea is that, for each row of context  $s$ ,

we encode sequence  $T_s$  with the binary wavelet tree of Theorem 6. (Actually, the characters used are not those in  $\Sigma$  but the column numbers.) In our example, for row  $s = "1a"$ , we encode  $R_j = T_s = "aa"$ . In order to retrieve element  $p$  from column  $c$  in global row  $j$ , we just compute  $select_c(R_j, p)$ , adding it to  $select_1(F, j) - 1$  to return the value in the final line of Fig. 17. In our example,  $select_{"a"}("aa", 1) = 2$ , which added to  $17 - 1$  gives the final answer 18.

The binary wavelet tree requires  $|T_s|H_0(T_s) + o(|T_s|)$  bits of space and answers the queries in  $O(\log \sigma)$  time. Adding over all the contexts  $s$  we get  $nH_k + o(n)$  bits. In exchange, the time increases by an  $O(\log \sigma)$  factor. In level  $\ell$ , the space remains the same but the query time of the wavelet tree is  $O(\log(\sigma^{2^\ell})) = O(2^\ell \log \sigma)$ . The search time added over  $1 + 1/\epsilon$  levels in  $0 \leq \ell \leq h'$  is  $O(\log^{1+\epsilon} n)$ , while the time to move from level  $h'$  to  $h$  is  $O(\log n \log \sigma)$ .

Yet, we still have a problem, namely the extra  $O(n)$  bits due to storing  $A_h$  and  $A_h^{-1}$ . These are converted into  $O(n \log \log n / \log_\sigma n)$  by setting  $h = \lceil \log \log_\sigma n - \log \log n \rceil$ . Now the jump from level  $h'$  to  $h$  takes  $O(\log^2 n / \log \log n)$  time. Thus we have an alternative structure.

**Theorem 14** *The Compressed Suffix Array of Grossi, Gupta, and Vitter [2003] (GGV-CSA) offers the following space/time tradeoffs.*

Space in bits	$\frac{1}{\epsilon} nH_k + O(n \log \log n / \log_\sigma n)$
Time to count	$O(m \log \sigma + \log^3 n / \log \log n)$
Time to locate	$O(\log^2 n / \log \log n)$
Time to display $\ell$ chars	$O(\ell / \log_\sigma n + \log^2 n / \log \log n)$
Conditions	$0 < \epsilon \leq 1$ is an arbitrary constant; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

Note that it is not good idea to use multiary wavelet trees (e.g. Theorem 7) to improve time, as these limit the alphabet size, while in the last levels of the structure the alphabet reaches size  $\sigma^{2^{h'}} = \Theta(n)$ .

More complicated tradeoffs are given by Grossi et al. [2003]. The most relevant ones obtain, roughly,  $O(m / \log_\sigma n + \log^{\frac{2\epsilon}{1-\epsilon}} n)$  counting time with  $\frac{1}{\epsilon} nH_k + o(n \log \sigma)$  space,  $0 < \epsilon < 1/3$ ; or  $O(m \log \sigma + \log^4 n)$  counting time with almost optimal  $nH_k + o(n \log \sigma)$  space.

*In practice.* Some experiments and practical considerations are given by Grossi et al. [2004]. They show that bit vectors  $Z$  can be represented using run-length encoding and then Elias- $\gamma$  [Elias 1975; Witten et al. 1999], so that they take in the worst case  $2|Z|H_0(Z)$  bits (and they may take less if the bits are not uniformly distributed). Note that this result was partially foreseen by Sadakane [2000, 2003] to achieve zero-order encoding of  $\Psi$  in the SAD-CSA (Section 6.2). They do not explain how to do *rank* and *select* in constant time on this representation, but sampling and local decompression (as in the end of Section 7.1) should be practical.

An interesting result of Grossi et al. [2004] is that, since the sum of all the  $\gamma$ -encodings across all the cells adds up  $2nH_k$  bits, we could use the same encoding to code each column in Fig. 16 as a whole. The values within a column are increasing. The total space for this representation is that of the  $\gamma$ -encoding inside the cells

(which overall amounts to  $2nH_k$  bits) plus that of the  $\gamma$ -encoding of the jumps between cells. The latter is shown to be  $o(n)$  as long as  $k \leq \alpha \log_\sigma n$  for some constant  $0 < \alpha < 1$ . Thus, we obtain  $2nH_k + o(n)$  space. Note, and this is the key part, that the sequence of differences we have to represent is the same no matter how the values are split along the rows. That is, the sequence (and its space) is the same  $\Psi$  independently of how long the contexts are. Therefore, this encoding achieves  $2nH_k + o(n)$  *implicitly and simultaneously* for any  $k \leq \alpha \log_\sigma n$ . This is in contrast with their original work [Grossi et al. 2003], where  $k$  had to be chosen at indexing time. Interestingly, this also shows that the Elias- $\delta$  representation of the SAD-CSA (where in fact a column-wise differential representation is used for  $\Psi$ ) actually requires  $nH_k + O(n \log \log \sigma)$  bits of space, improving the analysis by Sadakane [2000, 2003] (contrast with the other  $nH_k$ -like solution at the end of Section 7.1).

## 8. BACKWARD SEARCHING AND THE FM-INDEX FAMILY

Backward searching is a completely different approach to searching using suffix arrays. It matches particularly well with the BWT (Section 4.3), but it can also be applied with compressed suffix arrays based on the  $\Psi$  function, using the fact that  $\Psi$  and  $LF$  are the inverse of each other. The first exponent of this idea was the FM-Index of Ferragina and Manzini [2000], and many others followed. We start by explaining the idea and then go on to describe its different realizations.

### 8.1 The Backward Search Concept

Consider searching for  $P$  in  $A$  as follows. We start determining the range  $[sp_m, ep_m]$  in  $A$  of suffixes starting with  $P_m$ . Since  $P_m$  is a single character, function  $C$  of Lemma 3 can be used to determine  $[sp_m, ep_m] = [C(P_m)+1, C(P_m+1)]$ . Now, given  $[sp_m, ep_m]$ , we want to compute  $[sp_{m-1}, ep_{m-1}]$ , the interval of  $A$  corresponding to suffixes starting with  $P_{m-1,m}$ . This is of course a subinterval of  $[C(P_{m-1})+1, C(P_{m-1}+1)]$ . In the general case, we know the interval  $[sp_{i+1}, ep_{i+1}]$  of  $A$  corresponding to suffixes that start with  $P_{i+1,m}$  and want  $[sp_i, ep_i]$ , which is a subinterval of  $[C(P_i)+1, C(P_i+1)]$ . At the end,  $[sp_1, ep_1]$  is the answer for  $P$ .

The LF-mapping (Definition 14) is the key to obtain  $[sp_i, ep_i]$  from  $[sp_{i+1}, ep_{i+1}]$ . Consider again Fig. 6, and the search for  $P = \text{"ala"}$ . Initially,  $[sp_3, ep_3] = [C(\text{"a"})+1, C(\text{"b"})] = [5, 13]$ . Now, within this subinterval, we wish to know which of those "a"s are preceded by "l", as this would give us the occurrences of "la". That is, we know that all the "l"s in  $L[5, 13]$  appear contiguously in  $F$ . Let  $s$  and  $e$  be the first and last position in  $[5, 13]$  where  $L[s] = L[e] = \text{"l"}$ . In our example,  $s = 6$  and  $e = 9$ . Then,  $LF(s)$  and  $LF(e)$  are the first and last rows of  $M$  that start with "la". In our case,  $[sp_2, ep_2] = [LF(6), LF(9)] = [17, 19]$ . Remind from Lemma 3 that  $LF(6) = C(\text{"l"}) + Occ(\text{"l"}, 6) = 16 + 1$  and that  $LF(9) = C(\text{"l"}) + Occ(\text{"l"}, 9) = 16 + 3$ . With the same mechanism we finally get  $[sp_1, ep_1] = [10, 11]$ .

In general, given  $[sp_{i+1}, ep_{i+1}]$ , we determine  $s$  and  $e$  and then have  $[sp_i, ep_i] = [LF(s), LF(e)]$ . The problem is that we do not know  $s$  and  $e$ . Yet, this is not necessary. Since  $s$  is the position of the first occurrence of  $P_i$  in  $L[sp_{i+1}, ep_{i+1}]$ , it follows that  $Occ(P_i, s) = Occ(P_i, sp_{i+1} - 1) + 1$ . Likewise,  $Occ(P_i, e) = Occ(P_i, ep_{i+1})$  because  $e$  is the last occurrence of  $P_i$  in  $L[sp_{i+1}, ep_{i+1}]$ . The resulting algorithm is

rather simple and is shown in Fig. 18.

---

**Algorithm** FM-search( $P, m, n, C, Occ$ )

- (1)  $sp \leftarrow 1; ep \leftarrow n;$
- (2) **for**  $i \leftarrow m$  **to** 1
- (3)      $sp \leftarrow C(P_i) + Occ(P_i, sp - 1) + 1;$
- (4)      $ep \leftarrow C(P_i) + Occ(P_i, ep);$
- (5)     **if**  $sp > ep$  **then return**  $\emptyset;$
- (6)      $i \leftarrow i - 1;$
- (7) **return**  $[sp, ep];$

---

Fig. 18. Backward search algorithm to find the interval in  $A[1, n]$  of the suffixes that start with  $P_{1,m}$ .

Note that function  $C$  can be implemented in constant time as an array, using just  $\sigma \log n$  bits. A similar approach with  $Occ$  would require  $\sigma n \log n$  bits, which is unacceptable. All the different variants of the backward search concept aim basically at implementing  $Occ$  in little time and space. If we achieve constant time for  $Occ$ , then the backward search needs just  $O(m)$  time, which is better than any compressed suffix array from the previous section.

## 8.2 Backward Searching using $\Psi$

Before entering into the more typical backward search implementations, which are based on the BWT, we show that backward searching can be implemented using function  $\Psi$  [Sadakane 2002]. As mentioned at the end of Section 7.1, this is how the SAD-CSA is actually implemented.

Since  $\Psi$  and  $LF$  are inverse functions, we might binary search values  $LF(s)$  and  $LF(e)$  using  $\Psi$ . Imagine we already know  $[sp_{i+1}, ep_{i+1}]$  and  $[C(P_i) + 1, C(P_i + 1)]$ . Function  $\Psi$  is increasing in the latter interval. Moreover,  $[sp_i, ep_i]$  is the subinterval of  $[C(P_i) + 1, C(P_i + 1)]$  such that  $\Psi(j) \in [sp_{i+1}, ep_{i+1}]$  if  $j \in [sp_i, ep_i]$ . Hence, two binary searches permit obtaining  $[sp_i, ep_i]$  in  $O(\log n)$  time.

Backward search then completes in  $O(m \log n)$  time using the SAD-CSA, just as classical searching. An advantage of backward searching is that it is not necessary at all to obtain text substrings at search time. A disadvantage is that the average and worst case are similar, because we perform  $m$  steps of cost  $O(\log n)$  each. In classical searching, we perform  $O(\log n)$  steps of cost between 1 and  $m$  each, the average case being  $O(\log^2 n)$  rather than  $O(m \log n)$  if  $m > \log n$  (see the end of Section refsec:mn-ccsa). On the other hand, Sadakane [2002] shows how the backward search can be implemented in  $O(m)$  time if  $\sigma = O(\text{polylog}(n))$ , essentially using the same idea we present in Section 8.5.

We also remind (end of Section 7.1) that function  $\Psi$  is implemented via sampling, so the binary searching is performed first over the samples and then completed with a sequential decompression plus search between two samples. If the sampling step is  $L = \Theta(\log n)$  (to maintain the space cost of the samples within  $O(n)$ ), then the binary search complexity remains  $O(\log n)$  time and the overall search time  $O(m \log n)$ , even if we do not use four-Russian techniques for fast decompression. If we use the normal suffix array searching, there would have been  $O(m \log n)$  accesses

to arbitrary positions of  $\Psi$ , so the use of four-Russian techniques is mandatory to avoid a total search cost of  $O(m \log^2 n)$ .

### 8.3 FMI: Ferragina and Manzini's Implementation

The first implementation of backward searching was proposed, together with the concept itself, by Ferragina and Manzini [2000]. This will be called FMI in this paper.

Ferragina and Manzini [2000] show that  $Occ$  can be implemented so that it takes constant time, using space upper bounded by  $5nH_k + o(n)$ . This was the first structure achieving  $O(nH_k)$  bits of space. Essentially,  $Occ$  is implemented as the compressed BWT transformed text  $T^{bwt}$  and some directory information.

They compress  $T^{bwt}$  by applying move-to-front transform, then run-length compression, and finally a variable-length prefix code. Move-to-front [Bentley et al. 1986] consists of keeping a list of characters ranked by recency, that is, the last character seen is first in the list, then the next-to-last, and so on. Every time we see a new character  $c$ , which is at position  $p$  in the list, we output  $p$  and move  $c$  to the beginning of the list. This transform produces small numbers over text zones with few different characters. This is precisely what happens in  $T^{bwt}$ . In particular, there tend to appear runs of equal characters in  $T^{bwt}$  (precisely,  $n_{bw}$  runs, recalling Definition 15), which become runs of 1's after move-to-front. These runs are then captured by the run-length compression. Finally, the prefix code applied is a version of Elias- $\gamma$  with some provisions for the run lengths. Overall, they show using results from Manzini [2001] that this representation compresses  $T^{bwt}$  to at most  $5nH_k + O(\sigma^k \log n)$  bits. The compressed  $T^{bwt}$  will be called  $Z$ .

The directories to answer  $Occ(c, i)$  resemble the solution for *rank* in Section 5.1. We choose a *block length*  $\ell$  and cut the range  $[1, n]$  into *blocks* of length  $\ell$ . Every  $\ell$  consecutive blocks will be grouped into a *superblock* of length  $\ell^2$ . For each superblock  $1 \leq i \leq \lceil n/\ell^2 \rceil$  and each character  $c \in \Sigma$ , we store  $NO[c, i] = Occ(c, (i-1)\ell^2)$ , as well as  $W[i]$ , the bit position of  $T_{(i-1)\ell^2+1}^{bwt}$  in  $Z$ . This requires  $O(n\sigma \log n/\ell^2)$  bits. Similarly, for each block  $1 \leq j \leq \lceil n/\ell \rceil$  we store the same data relative to its superblock  $i = 1 + \lfloor j/\ell \rfloor$ , that is,  $NO'[c, j] = Occ(c, (j-1)\ell) - Occ(c, (i-1)\ell^2)$ , as well as  $W'[j]$ , the bit position of  $T_{(j-1)\ell+1}^{bwt}$  in  $Z$  minus  $W[i]$ . It is easy to see that  $NO'[c, j] \leq \ell^2$ , and similarly it holds  $W'[j] = O(\ell^2 \log \sigma)$ , so these two tables require  $O(n\sigma \log \ell/\ell)$  bits. In addition, for each block  $j$ , the state of the move-to-front transformation (that is, the recency rank of characters) at the beginning of the block is maintained in  $MTF[j]$ . This requires  $O(n\sigma \log \sigma/\ell)$  additional bits.

The final component is a table  $S$  that does not depend on  $Z$ .  $S[c, o, B, M]$  is indexed by a character  $c \in \Sigma$ , an offset inside a block  $o \in [1, \ell]$ , the content of a block (a bit stream) whose length is in the worst case  $\ell' = (1 + 2 \log \sigma)\ell$ , and the state of a move-to-front transformation (a permutation of  $[1, \sigma]$ ). The content of  $S[c, o, B, M]$  is  $Occ(c, o)$  for the text obtained by decompressing  $B$  starting with a move-to-front transform initialized as  $M$ .

It is not hard to see that, given position  $u$ , we compute  $i = 1 + \lfloor u/\ell^2 \rfloor$ ,  $j = 1 + \lfloor u/\ell \rfloor$  and  $o = u - (j-1)\ell$ , and then it holds  $Occ(c, u) = NO[c, i] + NO'[c, j] + S[c, o, Z_{W[i]+W'[j]}, W[i]+W'[j+1]-1, MTF[j]]$  (if  $j = \ell - 1$  then  $W[i] + W'[j+1]$  should actually be  $W[i+1]$ ). This can be computed in constant time on a RAM machine

provided  $|B| = \ell' = O(\log n)$  and  $|M| = O(\log n)$ . The first restriction yields  $\ell = O(\log_\sigma n)$ , while the second becomes  $\sigma = O(\log n / \log \log n)$ .

Fig. 19 illustrates, with  $\ell = 3$ . For example, to compute  $Occ(\text{"a"}, 14)$  we consider superblock  $i = 1 + \lceil 14/\ell^2 \rceil = 2$ , block  $j = 1 + \lceil 14/\ell \rceil = 5$ , and offset  $o = 14 - (j - 1)\ell = 2$ . Then  $Occ(\text{"a"}, 14) = NO[\text{"a"}, 2] + NO'[\text{"a"}, 5] + S[\text{"a"}, 2, \{1, 6, 1\}, \text{"b\$ldar"}] = 3 + 0 + 1 = 4$ .

<i>Superblocks</i>	<i>Blocks</i>	<i>MTF (bwt(T))</i>
$S[\text{"a"}, 2, \{1, 6, 1\}, \text{"b\_ldar"}] = 1$	NO [a,1]=0	NO' [a,1]=0
		NO' [a,2]=2
		NO' [a,3]=3
	NO [a,2]=3	NO' [a,4]=0
		NO' [a,5]=0
		NO' [a,6]=2
	NO [a,3]=6	NO' [a,7]=0

a	3	\$_abdlr
r	7	
a	2	
a	1	ar\$_bdl
d	6	
l	7	
_	6	ldar\$_b
l	2	
l	1	
\$	6	l_dar\$b
_	3	
b	7	
b	1	b_\$ldar
a	6	
a	1	
r	7	ab_\$ldr
_	4	
a	3	
a	1	a_rb\$d
a	1	
a	1	

Fig. 19. The main components of the FMI structure, for the text "alabar a la alabarda\$" using  $\ell = 3$ . The compression of  $T^{bwt} = bwt(T)$  is illustrated only up to the move-to-front (MTF) transform for readability, and we also show the MTF states stored per block. We show the  $NO$  and  $NO'$  values only for character "a", and omit  $W$  and  $W'$ . We show just a single entry of  $S$ .

Adding up all the space complexities we obtain, apart from  $|Z|$ , the upper bound  $O(n\sigma \log n/\ell^2 + n\sigma \log \ell/\ell + n\sigma \log \sigma/\ell + \sigma \ell 2^{\ell'} \sigma! \log \ell)$ . If we choose  $\ell = x \log_\sigma n$  for constant  $0 < x < 1/3$ , then  $\ell' < 3x \log n < \log n$ . Under this setting the extra space is  $O(n\sigma \log \sigma \log \log n / \log n + (\sigma/e)^{\sigma+3/2} n^{3x} \log_\sigma n \log \log n)$ .

In order to locate occurrences, they sample text positions at regular intervals, just as described in Section 7.2 (the concept was invented by Ferragina and Manzini [2000] and reused for the MN-CCSA). Instead of marking one text position every  $\epsilon \log n$ , they mark one text position every  $\log^{1+\epsilon} n$ , for some  $\epsilon > 0$ . To know  $A[i]$ , they find the smallest  $r \geq 0$  such that  $LF^r(i)$  is a marked position, and then  $A[i] = A[LF^r(i)] + r$ . This way, they pay  $O(n/\log^\epsilon n)$  extra space and can locate the occurrences in  $O(occ \log^{1+\epsilon} n)$  steps. A problem is that there is no easy way to know  $T_i^{bwt}$  in order to compute  $LF(i) = C(T_i^{bwt}) + Occ(T_i^{bwt}, i)$  (Lemma 3). They discover  $T_i^{bwt}$  by sequentially searching  $\Sigma$  for the  $c$  such that  $Occ(c, i) \neq Occ(c, i-1)$ . This takes  $O(\sigma)$  time per step. A similar approach permits displaying  $T_{l,r}$  in  $O(\sigma(r-l + \log^{1+\epsilon} n))$  time.



**Theorem 15** *The FM-Index of Ferragina and Manzini [2000] (FMI) offers the following space/time tradeoff.*

Space in bits	$5nH_k + O(n\sigma \log \log n / \log_\sigma n$ $+ (\sigma/e)^{\sigma+3/2} n^\gamma \log_\sigma n \log \log n + n / \log^\epsilon n)$ (e.g. it is $5nH_k + o(n)$ if $\sigma = O(\log^\lambda n)$ for $0 < \lambda < 1$ )
Time to count	$O(m)$
Time to locate	$O(\sigma \log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O(\sigma(\ell + \log^{1+\epsilon} n))$
Conditions	$\sigma = O(\log n / \log \log n)$ ; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$ ; $0 < \gamma < 1$ and $\epsilon > 0$ are arbitrary constants

We note that  $5nH_k$  is actually a rather pessimistic upper bound, and that the technique works with essentially any compressor for  $T^{bwt}$ . Thus the FMI obtains unbeaten counting complexity and attractive space complexity. Its real problem is the alphabet dependence. Apart from  $|Z|$ , the other terms of the space complexity depend heavily on  $\sigma$ , up to the point that they are not sublinear for  $\sigma = \Omega(\log n / (\log \log n)^d)$  for any  $d \geq 1$ . Further work on the FMI have focused on alleviating its dependence on  $\sigma$ .

Some more complicated techniques [Ferragina and Manzini 2000], based on using alphabet  $\Sigma^q$  instead of  $\Sigma$ , permit reducing the  $O(\log^{1+\epsilon} n)$  to  $O(\log^\epsilon n)$ , yet this makes the alphabet dependence of the index even sharper.

*In practice.* A practical implementation of the FMI could not follow the idea of the  $S$  table. Ferragina and Manzini [2001] propose replacing  $S$  with a plain decompression and scanning of block  $B$ , which (according to the theoretical value of  $\ell$ ) takes  $O(\log n)$  time and raises the counting complexity to  $O(m \log n)$ . Some heuristics have also been used to reduce the size of the directories in practice. Also, instead of sampling the text at regular intervals, all the occurrences of some given character are sampled. This removes the need to store a table telling which positions are marked, as this can be deduced from the current character in  $T^{bwt}$ .

Finally, they consider alternative ways of compressing the text. The most successful one is to compress each block with Huffman, with a distinct Huffman tree per block. If we recall Theorem 2, this does not guarantee  $O(nH_k)$  bits of space, but it should be pretty close. The reason is that  $T^{bwt}$  is partitioned into equal-size blocks, not according to contexts of length  $k$ . Such a partitioning will be considered in Section 8.7.

#### 8.4 Huff-FMI: An $O(nH_0)$ Size Implementation

We present now an alternative implementation of the backward search idea that is unable to reach the  $O(nH_k)$  size bound, yet it is an interesting way to remove the alphabet dependence. The *Huffman FM-Index* by Grabowski, Mäkinen, and Navarro [2004] (HUFF-FMI) uses Huffman as a tool to reduce the alphabet of the text to bits.

The idea is to first compress  $T$  using Huffman [Huffman 1952; Bell et al. 1990]. In the resulting bit stream  $T'_{1,n'}$ , of  $n' < n(H_0 + 1)$  bits, mark those bits that start a

codeword. Apply the BWT over  $T'$  to obtain bit stream  $B[1, n'] = (T')^{bwt}$ . Create another bit stream  $Bh[1, n']$  indicating the pointed bits that are marked in  $T'$ . That is, if  $A'[1, n']$  is the suffix array of  $T'$ , then  $Bh[i] = 1$  iff  $T'[A'[i]]$  is marked. No attempt is made to compress  $B$  nor  $Bh$ .

Implementing *Occ* over  $B$  is very easy, because it is a binary stream. Indeed,  $Occ(0, i) = rank_0(B, i)$  and  $Occ(1, i) = rank_1(B, i)$ , which can be answered in constant time using  $n' + o(n')$  bits with the techniques of Section 5.1. Any alphabet dependence has vanished. (We omit some technicalities due to boundary effects because of not having a terminator character for  $T'$ .)

To count the occurrences of  $P_{1,m}$  in  $T$ , we first Huffman-transform  $P$  into a binary pattern  $P'_{1,m'}$ , using the codebook we have created for  $T$ . Then, we use the backward search of Fig. 18 for  $P'$  over  $B$ , finishing after  $O(m')$  constant-time steps with  $[sp', ep']$  as the bounds in  $A'$  for the suffixes that start with  $P'$ . Note that  $m' < m(H_0 + 1)$  if  $P$  has the same zero-order entropy of  $T$ , and in any case  $m' = O(m \log n)$  because the longest Huffman code has  $O(\log n)$  bits [Witten et al. 1999, pp. 397].

Yet, not every occurrence of  $P'$  in  $T'$  corresponds to an occurrence of  $P$  in  $T$ , as we want only the codeword-aligned occurrences. This is where  $Bh$  comes into play. Only the bits set in  $Bh[sp', ep']$  correspond to real occurrences, so we complete the counting query by returning  $rank_1(Bh, ep') - rank_1(Bh, sp' - 1)$ . In order to find the exact entries in  $A'$  corresponding to occurrences, we start with  $b = rank_1(Bh, sp' - 1)$  and use  $select_1(Bh, b + 1)$ ,  $select_1(Bh, b + 2)$ , and so on, taking constant time per occurrence.

Fig. 20 illustrates the search for  $P = \text{"ba"}$ . This is translated into  $P' = 11010$ . The backward search yields  $A'[43, 46]$ . From those, only  $rank_1(Bh, 46) - rank_1(Bh, 42) = 2$  are valid.

To locate the occurrences, the same mechanism of the MN-CCSA (Section 7.2) of marking  $T'$  at regular intervals of the form  $\epsilon \log n$  is employed, yet the marks are slightly moved to align to codeword beginnings. The absolute values stored for the samples correspond to positions in  $T$ , not in  $T'$ . Using the LF-mapping over  $B$  and using  $Bh$  to keep count of how many positions in  $T$  we have moved, we can deduce the text position corresponding to  $A'[i]$  in  $O(\epsilon(H_0 + 1) \log n)$  time. Finally, to display  $T_{l,r}$  we store the same samples, this time in text position order and pointing to  $A'$  entries. After extracting the bits from  $T'$ , we decode them to get  $T_{l,r}$ .

**Theorem 16** *The Huffman FM-Index of Grabowski, Mäkinen, and Navarro [2004] (HUFF-FMI) offers the following space/time tradeoff.*

Space in bits	$n(2H_0 + 3 + \frac{1}{\epsilon}) + O(n \log \log n / \log n + \sigma \log \sigma)$
Time to count	$O(m(H_0 + 1))$ (average case) $O(m \log n)$ (worst case)
Time to locate	$O(\epsilon(H_0 + 1) \log n)$
Time to display $\ell$ chars	$O((H_0 + 1)(\ell + \epsilon \log n))$ (average case) $O((\ell + \epsilon(H_0 + 1)) \log n)$ (worst case)
Conditions	$\epsilon > 0$ is an arbitrary constant

We note that the average-case complexity for counting assumes that  $P$  has the



The idea of the WT-FMI is extremely simple, once in context. Just use the wavelet tree of Section 5.3 over the sequence  $T^{bwt}$ . Hence,  $Occ(c, i) = rank_c(T^{bwt}, i)$  can be answered in  $O(\log \sigma)$  time using the basic wavelet tree (Theorem 6), and in  $O(1)$  time for  $\sigma = O(\text{polylog}(n))$  using the multi-ary one (Theorem 7). The method for locating the occurrences and displaying the text is the same as for the FMI, yet this time we also find  $T_i^{bwt}$  in  $O(\log \sigma)$  or  $O(1)$  time using the same wavelet tree.

Fig. 21 illustrates, using a binary wavelet tree, how to obtain  $Occ("a", 15) = rank_{"a"}(T^{bwt}, 15) = 5$ . By algebraic manipulation we know that "a" is in the first half of the alphabet, so it is a zero at the root. Since  $rank_0(15) = 8$  at the root, we go to the left child and consider now position 8. Now we know that "a" is in the second half of this node, thus it is a one in this bitmap. Since  $rank_1(8) = 5$  at this node, we go to its right child and consider position 5. But the new node is a leaf of all "a"s and then the answer is 5. Fig. 22 gives the pseudocode. Depending on which wavelet tree we use, different tradeoffs are obtained.

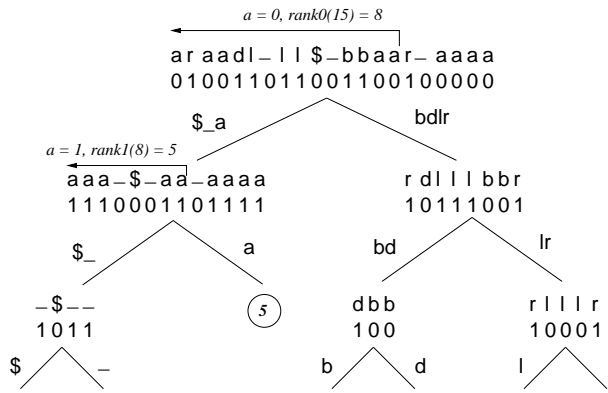


Fig. 21. A binary wavelet tree for the text "alabar a la alabarda\$". This is the main component of the WT-FMI. Only the bitmaps are stored, the texts are shown for clarity.

---

**Algorithm** WT-Occ( $c, i, \sigma_1, \sigma_2, v$ )

- (1) if  $\sigma_1 = \sigma_2$  then return  $i$ ;
  - (2)  $\sigma_m = \lfloor \frac{\sigma_1 + \sigma_2}{2} \rfloor$ ;
  - (3) if  $c \leq \sigma_m$ 
    - (4) then return WT-Occ( $c, rank_0(B_v, i), \sigma_1, \sigma_m, v_l$ );
    - (5) else return WT-Occ( $c, rank_1(B_v, i), \sigma_m + 1, \sigma_2, v_r$ );
- 

Fig. 22. Computing  $Occ(c, i)$  on a binary wavelet tree. It is invoked as WT-Occ( $c, i, 1, \sigma, root$ ). We call  $B_v$  the bit vector at tree node  $v$ ,  $v_l$  its left child, and  $v_r$  its right child.

**Theorem 17** *The Wavelet Tree FM-Index (WT-FMI) offers the following space/time tradeoffs.*

Space in bits	$nH_0 + O(n \log \log n / \log_\sigma n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma)$
Time to locate	$O(\log \sigma \log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma)$
Space in bits	$nH_0 + O(n / \log^\epsilon n)$
Time to count	$O(m)$
Time to locate	$O(\log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O(\ell + \log^{1+\epsilon} n)$
Conditions	$\sigma = O(\text{polylog}(n))$
Space in bits	$nH_0 + O(n / \log_\sigma \log n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma / \log \log n)$
Time to locate	$O(\log^{1+\epsilon} n \log \sigma / \log \log n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma / \log \log n)$
Conditions	$\sigma = o(n / \log \log n)$ , $\sigma = \omega(\text{polylog}(n))$
Conditions for all	$0 < \epsilon < 1$ is an arbitrary constant

Despite of its simplicity, the WT-FMI is the precursor of further research that lead to the best implementations of the backward search concept (Sections 8.6 and 8.7).

*In practice:* The implementation of the WT-FMI uses the binary wavelet tree, preprocessed for *rank* using the simple techniques of Section 5.1, and gives the wavelet tree the shape of the Huffman tree of the text. This way, instead of the theoretical  $nH_0 + o(n)$  bits, we obtain  $n(H_0 + 1)(1 + o(1))$  bits with much simpler means. In addition, the Huffman shape gives the index  $O(mH_0)$  average counting time. The worst case time is  $O(m \log n)$ , but it can be limited to  $O(m \log \sigma)$  without losing the  $O(mH_0)$  average, by pruning the Huffman tree and make it balanced after depth  $(1 + x) \log \sigma$  [Mäkinen and Navarro 2004b].

### 8.6 The Run-Length FM-Index of Mäkinen and Navarro (RL-FMI)

The *Run-Length FM-Index* of Mäkinen and Navarro [2004b, 2004c, 2005a] (RL-FMI) is an improvement over the WT-FMI of Section 8.5, which exploits the equal-letter runs of the BWT (Theorem 3) to achieve  $O(nH_k \log \sigma)$  bits of space. It retains the good search complexities of the FMI, but it is much more resistant to the alphabet size. Actually this was the first index achieving  $O(m)$  search time for  $\sigma = O(\text{polylog}(n))$  and taking simultaneously space proportional to the  $k$ -th order entropy of the text. The idea is to compute  $\text{Occ}(c, i) = \text{rank}_c(T^{bwt}, i)$  using a wavelet tree built over the run-length compressed version of  $T^{bwt}$ .

In Fig. 21 we built the wavelet tree of  $T^{bwt} = \text{"araadl 1l\$ bbaar aaaa"}$ . Assume that we run-length compress  $T^{bwt}$  to obtain  $S = \text{"aradl 1\$ bar a"}$ . By Theorem 3, we have the limit  $|S| \leq nH_k + \sigma^k$  for any  $k$ . Therefore, a wavelet tree (Section 5.3) built over  $S$  would require  $(nH_k + \sigma^k)H_0(S) + o(n)$  bits. The only useful bound we have for the zero-order entropy of  $S$  is  $H_0(S) \leq \log \sigma$ , thus the space bound is  $nH_k \log \sigma + o(n)$  for any  $k \leq \alpha \log_\sigma n$  for any constant  $0 < \alpha < 1$ .

The problem is that *rank* over  $S$  does not give the answers we need over  $T^{bwt}$ . For example, assume we want to compute  $\text{rank}_{\text{"a"}}(T^{bwt}, 19) = 7$ . We need to know that  $T_{19}^{bwt}$  lies at  $S_{14}$ . This is easily solved by defining a bitmap  $B[1, n]$  indicating

the beginnings of the runs in  $T^{bwt}$ . In our case  $B = 111011110111010111000$ . We know that the position of  $T_{19}^{bwt}$  in  $S$  is  $rank_1(B, 19) = 14$ . Yet, this is not sufficient, as  $rank_{"a"}(S, 14) = 4$  just tells us that there are 4 runs of "a"s before and including that of  $T_{19}^{bwt}$ . What we need is to know the total length of those runs, and in which position of its run is  $T_{19}^{bwt}$  (in our case, 2nd).

For this sake, we reorder the runs in  $B$  alphabetically, accordingly to the characters that form the run. Runs of the same character stay in the same relative order. We form bit array  $B'[1, n]$  with the reordered  $B$ . In our case  $B' = 111111010100010111011$ . We also compute array  $C_S$  indexed by  $\Sigma$ , so that  $C_S[c]$  tells the number of occurrences in  $S$  (runs in  $T^{bwt}$ ) of characters smaller than  $c$  (thus  $C_S$  plays for  $S$  the same role  $C$  plays for  $T$  in the FMI). In our example  $C_S["a"] = 4$ . This means that, in  $B'$ , the first  $C_S["a"] = 4$  runs correspond to characters smaller than "a", and then come those of "a", of which  $T_{19}^{bwt}$  is in the 4th because  $rank_{"a"}(S, 14) = 4$ . Fig. 23 illustrates.

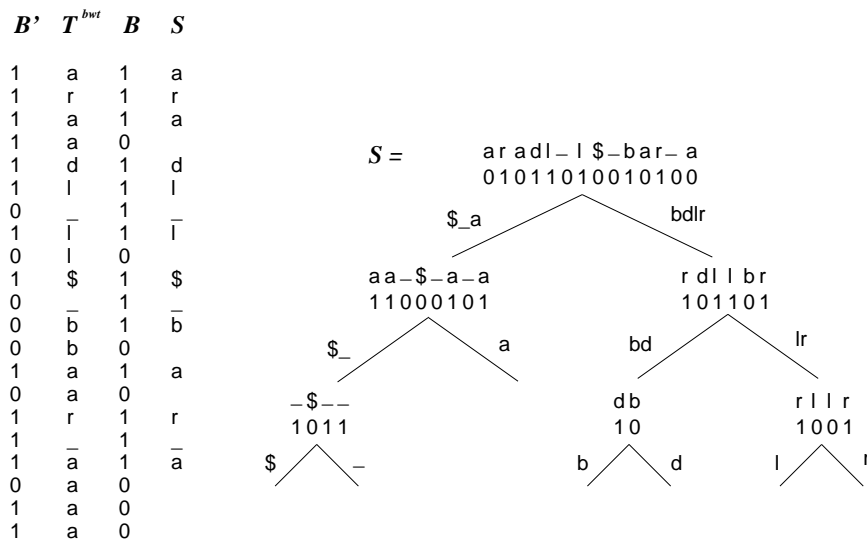


Fig. 23. The main RL-FMI structures for the text "alabar a la alabarda\$". The transformed text  $T^{bwt}$  is shown only for clarity. The wavelet tree on the right, built for  $S$ , stores only the bitmaps, not the texts at each node.

To compute  $rank_c(T^{bwt}, i)$ , we first find  $i' = rank_1(B, i)$ , the position of the run  $T_i^{bwt}$  belongs to in  $S$ . Thus there are  $j' = rank_c(S, i')$  runs of  $c$ 's in  $T_{1,i'}^{bwt}$ . In  $B'$ , the runs corresponding to  $c$  start at  $j = select_1(B', C_S[c] + 1)$ . Now there are two cases. If  $S_{i'} \neq c$ , then the run of  $T_i^{bwt}$  does not belong to  $c$ , and thus we must accumulate the full length of the first  $j'$  runs of  $c$ ,  $select_1(B', C_S[c] + 1 + j') - j$ . If, on the other hand,  $S_{i'} = c$ , then the run of  $T_i^{bwt}$  does belong to  $c$ , and we must be careful on how much of the last run we count. We are sure that the first  $j' - 1$  runs must be fully counted, so we have  $select_1(B', C_S[c] + j') - j$ , but we must add the corresponding part of the last run. This part is clearly  $i - select_1(B, i) + 1$ . Fig. 24 gives the pseudocode.

---

**Algorithm** RLFM-Occ( $c, i, S, B, B', C_S$ )

- (1)  $i' \leftarrow \text{rank}_1(B, i)$ ;
- (2)  $j' \leftarrow \text{rank}_c(S, i')$ ;
- (3)  $j \leftarrow \text{select}_1(B', C_S[c] + 1)$ ;
- (4) **if**  $S_{i'} = c$  **then**
- (5)  $j' \leftarrow j' - 1$ ;
- (6)  $ofs \leftarrow i - \text{select}_1(B, i') + 1$ ;
- (7) **else**  $ofs \leftarrow 0$ ;
- (8) **return**  $\text{select}_1(B', C_S[c] + 1 + j') - j + ofs$ ;

---

Fig. 24. Algorithm to compute  $\text{Occ}(c, i)$  with the RL-FMI.

Thus the RL-FMI takes  $nH_k \log \sigma + 2n + o(n)$  bits of space, and it solves  $\text{Occ}(c, i) = \text{rank}_c(T^{bwt}, i)$  in the time necessary to perform  $\text{rank}_c$  over  $S$ . Depending on which wavelet tree is used to implement this operation is implemented, we have different results. The rest is handled just like the WT-FMI.

**Theorem 18** *The Run-Length FM-Index of Mäkinen and Navarro [2004b, 2004c, 2005a] (RL-FMI) offers the following space/time tradeoffs.*

Space in bits	$nH_k \log \sigma + 2n$ $+O(n \log \log n / \log_\sigma n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma)$
Time to locate	$O(\log \sigma \log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma)$
Space in bits	$nH_k \log \sigma + 2n + O(n / \log^\epsilon n)$
Time to count	$O(m)$
Time to locate	$O(\log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O(\ell + \log^{1+\epsilon} n)$
Conditions	$\sigma = O(\text{polylog}(n))$
Space in bits	$nH_k \log \sigma + 2n + O(n / \log_\sigma \log n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma / \log \log n)$
Time to locate	$O(\log^{1+\epsilon} n \log \sigma / \log \log n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma / \log \log n)$
Conditions	$\sigma = o(n / \log \log n)$ , $\sigma = \omega(\text{polylog}(n))$
Conditions for all	$0 < \epsilon < 1$ is an arbitrary constant; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

*In practice.* The implementation of the RL-FMI, just as that of the WT-FMI (end of Section 8.5), uses binary wavelet trees with Huffman shape, with the bitmaps using the techniques of Section 5.1. This gives at most  $nH_k(H_0(S) + 1)(1 + o(1))$  space, which in the worst case is  $nH_k(\log \sigma + 1)(1 + o(1))$  bits, close to the theoretical version but much simpler to implement. Even when it has not been proved that  $H_0(S)$  is smaller than anything other than  $\log \sigma$ , in practice space and access time are reduced.

### 8.7 The Alphabet-Friendly FM-Index of Ferragina et al. (AF-FMI)

The *Alphabet-Friendly FM-Index* of Ferragina, Manzini, Mäkinen, and Navarro [2004a, 2004c, 2004b] (AF-FMI) is another improvement over the WT-FMI of Section 8.5. It combines the idea with Theorem 2 to achieve  $nH_k + o(n)$  bits of space and the same search time of the WT-FMI.

Theorem 2 tells that, if we split  $T^{bwt}$  into substrings  $T_s$  according to its contexts  $s$  of length  $k$ , and manage to represent each resulting block  $T_s$ , of length  $n_s = |T_s|$ , in  $n_s H_0(T_s) + f(n_s)$  bits, for any convex function  $f$ , then the sum of all bits used is  $nH_k + \sigma^k f(n/\sigma^k)$ . In particular, we can use the binary wavelet tree of Section 8.5 for each block (Theorem 5.3). It requires  $n_s H_0(T_s) + O(n_s \log \log n_s / \log_\sigma n_s)$  bits, so we need overall  $nH_k + O(n \log \log(n/\sigma^k) / \log_\sigma(n/\sigma^k))$  bits, for any  $k$ . If  $k \leq \alpha \log_\sigma n$ , the space is  $nH_k + O(n \log \log n / \log_\sigma n)$ .

Assume  $s$  is the  $j$ -th nonempty context in  $T^{bwt}$ . The wavelet tree of  $T_s$  allows us to solve  $Occ_j(c, i)$ , which is the number of occurrences of  $c$  in  $(T_s)_{1,i}$ . To answer a global  $Occ(c, i)$  query, we must be able to (1) determine which block  $j$  does  $i$  belong to, so as to know which wavelet tree to query, and (2) know how many occurrences of  $c$  are there before  $T_s$  in  $T^{bwt}$ .

We store a bit vector  $B[1, n]$  indicating the beginning of blocks in  $T^{bwt}$ , so that  $j = rank_1(B, i)$  (this is equivalent to vector  $F$  in Section 7.3). Moreover, the  $j$ -th block starts at  $i' = select_1(B, j)$ .

We also store, for each block  $j$ , vector  $C_j[c]$ , which tells the number of occurrences of  $c$  in  $T_{1,i'-1}^{bwt}$ , that is, before block  $j$  (or before substring  $T_s$ ). Since  $C_j[c] = Occ(c, i' - 1)$ , it is clear that  $Occ(c, i) = C_j[c] + Occ_j(c, i - i' + 1)$ . Similarly, to determine  $T_i^{bwt}$ , we obtain  $j$  and  $i'$ , and query the wavelet tree of the  $j$ -th block to find its  $(i - i' + 1)$ -th character.

Using the technique of Section 5.2, bit vector  $B$  requires at most  $\sigma^k \log n$  bits because it has at most  $\sigma^k$  bits set. On the other hand, arrays  $C_j$  require overall  $\sigma^{k+1} \log n$  bits of space. If  $k \leq \alpha \log_\sigma n$  for some constant  $0 < \alpha < 1$ , the overall space stays  $nH_k + o(n)$ .

Fig. 25 illustrates. To determine  $Occ("1", 8)$ , we first find that  $j = rank_1(B, 8) = 3$  is the block number where  $i = 8$  belongs. The first position of block 3 is  $i' = select_1(B, 3) = 5$ . The number of occurrences of "1" in  $T_{1,i'-1}^{bwt}$ , that is, before block  $j = 3$ , is  $C_3("1") = 0$ . Inside block 3, corresponding to substring  $T_s = "d1\ 11\$ \ bb"$  of  $T^{bwt}$ , we need the number of "1"s in  $(T_s)_{1,i-i'+1} = (T_s)_{1,4}$ . This is given by the wavelet tree, which answers  $Occ_3("1", 4) = 2$ . Thus the answer is  $0 + 2 = 2$ .

Yet, Ferragina et al. [2004a] go further. Instead of choosing a fixed  $k$  value in advance, they use a method by Ferragina and Manzini [2004] that, given a space overhead function  $f(n_s)$  on top of  $n_s H_0(T_s)$ , finds the partition of  $T^{bwt}$  that optimizes the final space complexity. Using blocks of fixed context length  $k$  is just one of the possibilities considered in the optimization, so the resulting partition is below  $nH_k + \sigma^k f(n/\sigma^k)$  *simultaneously for all  $k$*  (and it is possibly better than with any fixed  $k$ ). That is, although we have made the analysis assuming a given  $k$ , the construction does not have to choose any  $k$  but it reaches the space bound for any  $k$  of our choice. Thus this index also achieves the independence of  $k$  mentioned at the end of Section 7.3, achieving at the same time the minimum space  $nH_k$ .

The locating of occurrences and displaying of text is handled just as in Section 8.5.



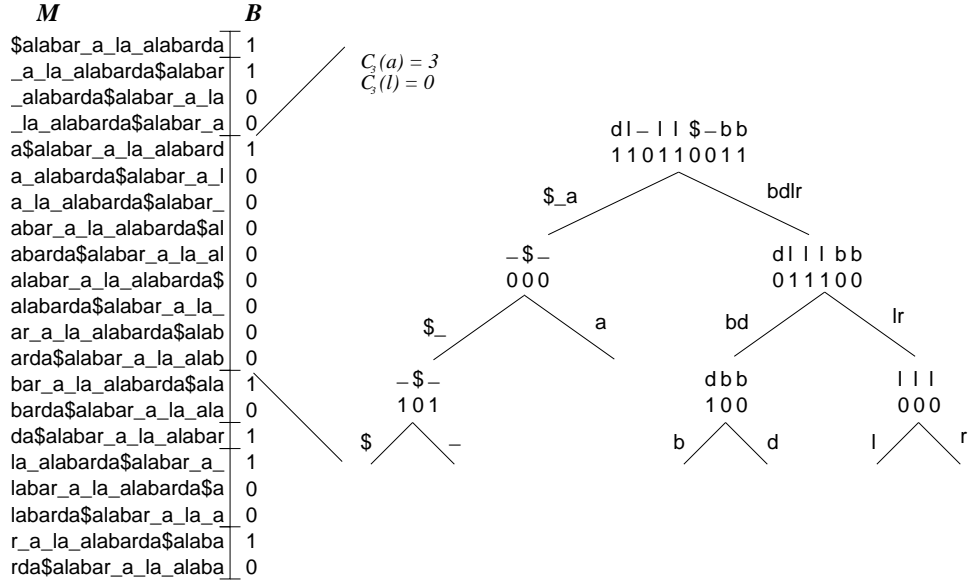


Fig. 25. The main AF-FMI structures for the text "alabar a la alabarda\$" considering contexts of length  $k = 1$ . Matrix  $M$  of the BWT is shown only for clarity. We show only one of the wavelet trees, corresponding to context "a", and only a couple of its  $Occ_j$  values. Note that this wavelet tree corresponds to a substring of that in Fig. 21.

Depending on which wavelet tree we use, we obtain different times for different alphabet sizes.

**Theorem 19** *The Alphabet-Friendly FM-Index of Ferragina, Manzini, Mäkinen, and Navarro [2004a, 2004c, 2004b] (AF-FMI) offers the following space/time tradeoffs.*

Space in bits	$nH_k + O(n \log \log n / \log_\sigma n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma)$
Time to locate	$O(\log \sigma \log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma)$
Space in bits	$nH_k + O(n / \log^\epsilon n)$
Time to count	$O(m)$
Time to locate	$O(\log^{1+\epsilon} n)$
Time to display $\ell$ chars	$O(\ell + \log^{1+\epsilon} n)$
Condition	$\sigma = O(\text{polylog}(n))$
Space in bits	$nH_k + O(n / \log_\sigma \log n + n / \log^\epsilon n + \sigma \log n)$
Time to count	$O(m \log \sigma / \log \log n)$
Time to locate	$O(\log^{1+\epsilon} n \log \sigma / \log \log n)$
Time to display $\ell$ chars	$O((\ell + \log^{1+\epsilon} n) \log \sigma / \log \log n)$
Conditions	$\sigma = o(n / \log \log n), \sigma = \omega(\text{polylog}(n))$
Conditions for all	$0 < \epsilon < 1$ is an arbitrary constant; $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

## 9. ZIV-LEMPER BASED INDEXES

Up to now we have considered different ways of compressing suffix arrays. While this is clearly the most popular trend on compressed indexing, it is worthwhile to know that there exist alternative approaches to self-indexing, based on Ziv-Lempel compression. In particular, one of those achieves  $O(m + occ)$  time to locate the  $occ$  occurrences of  $P$  in  $T$ , with no restriction on  $m$  or  $occ$ . This has not been achieved with other indexes.

### 9.1 Ziv-Lempel Compression

In the seventies, Ziv and Lempel [1977, 1978] presented a new approach to data compression. It was not based on text statistics, but rather on identifying repeated text substrings and replacing repetitions by pointers to their former occurrences in  $T$ . Ziv-Lempel methods produce a *parsing* (or partitioning) of the text into *phrases*.

**Definition 16** *The Ziv-Lempel parsing of text  $T_{1,n}$  is a sequence  $Z[1, n']$  of phrases such that  $T = Z[1] Z[2] \dots Z[n']$ . In particular, the original Ziv-Lempel parsing [Lempel and Ziv 1976] is built as follows. Assume we have already processed  $T_{1,i-1}$  producing sequence  $Z[1, p-1]$ . Then, we find the longest prefix  $T_{i,i'-1}$  of  $T_{i,n}$  which occurs in  $T_{1,i-1}$ . If  $i' > i$  then  $Z[p] = T_{i,i'-1}$  and we continue from text position  $i'$ . Otherwise  $T_i$  has not appeared before and  $Z[p] = T_i$ , continuing from text position  $i + 1$ . The process finishes once we obtain  $Z[n'] = "\$"$ .*

The output of a compressor using this parsing is essentially the position where each new phrase starts (called its *source*) and its length (new characters in  $\Sigma$  that appear are exceptions in this encoding). An important property of this parsing is that every phrase has appeared before, unless it is a new character of  $\Sigma$ . (The original definition [Lempel and Ziv 1976] actually permits the former occurrence of  $T_{i,i'-1}$  to extend beyond position  $i - 1$ , but we ignore this here.)

We will also be interested in a parsing called LZ78 [Ziv and Lempel 1978], where each phrase is formed by an already known phrase concatenated with a new character at the end.

**Definition 17** *The LZ78 parsing is built as follows. The first phrase is  $Z[1] = \varepsilon$ . Assume we have already processed  $T_{1,i-1}$  producing sequence  $Z[1, p-1]$ . Then, we find the longest prefix of  $T_{i,n}$  which is equal to some  $Z[p']$ ,  $1 \leq p' < p$ . Thus  $T_{i,n} = Z[p'] c T_{i',n}$ , where  $Z[p'] c$  does not appear in  $Z$  and  $i' = i + |Z[p']| + 1$ . We define  $Z[p] = Z[p'] c$ , increment  $p$ , and continue processing  $T$  from position  $i'$ . The process finishes once we get  $c = "\$"$ .*

The output of the LZ78 compressor is essentially the sequence of pairs  $(p', c)$  found at each step  $p$  of the algorithm. This parsing has a couple of important properties. First, all the phrases in an LZ78 parsing are different from each other. Second, the prefixes of a phrase are phrases.

Fig. 26 shows the LZ78 parsing of our example text. The figure also illustrates the *Ziv-Lempel trie*  $LZTrie$ , which is the trie storing the set of strings  $Z$ . The trie has  $n'$  nodes (as there is one per string in  $Z$ ). If  $Z[p] = Z[p'] c$ , then node  $p$  is a child of  $p'$  by edge labeled  $c$ .  $LZTrie$  is used in LZ78 compression to produce the

Ziv-Lempel parse in  $O(n)$  time. The other trie, *RevTrie*, is used by some indexes we review soon. Depending on the case, we can associate some extra information to these tries, such as the text position of the phrases, the lexicographical rank of the phrases (shown in the figure), and so on.

**Definition 18** *The Ziv-Lempel trie of  $T$  is a trie storing all the phrases of the Ziv-Lempel parsing of  $T$ . Each node corresponds to a distinct phrase.*

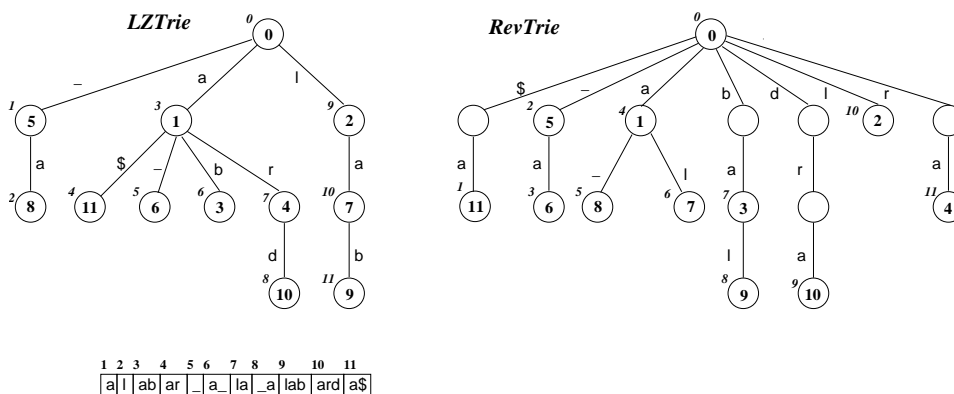


Fig. 26. On the bottom, the LZ78 parsing for the text "alabar a la alabarda\$". On the left we show *LZTrie*, the trie of the phrases  $Z[p]$ , and on the right *RevTrie*, the trie indexing the reverse phrases. We use the phrase identifiers as node labels. The smaller numbers in italics outside the nodes are the lexicographic rank of the corresponding strings.

An important property of the Ziv-Lempel parsings is that the number of phrases they produce is  $n' \leq n / \log_\sigma n$  in the worst case, and in general the size of the Ziv-Lempel compressed text (slowly) converges to the entropy of the source [Cover and Thomas 1991]. Which has more direct relevance to us, it has been shown that  $n'$  is related to the definition of  $H_k(T)$  we use in this paper [Kosaraju and Manzini 1999; Ferragina and Manzini 2002].

**Lemma 8** *Let  $n'$  the number of phrases produced by Ziv-Lempel parsing of text  $T_{1,n}$ . Then  $n' \log n = nH_k(T) + O(n \log \log n / \log n)$  for any  $k \leq \alpha \log_\sigma n$ , for any constant  $0 < \alpha < 1$ .*

Lemma 8 implies that the Ziv-Lempel trie can be stored, even using pointers, in  $O(nH_k)$  bits of space.

The pioneer work in Ziv-Lempel based indexes is due to Kärkkäinen and Ukkonen [1996a], deriving from their earlier work on sparse suffix trees [Kärkkäinen and Ukkonen 1996b]. This Ziv-Lempel based index is also the first compressed index we know of.

Before entering into the Ziv-Lempel based methods, let us review the sparse suffix tree [Kärkkäinen and Ukkonen 1996b], which is the first succinct index we know of. This is a suffix tree indexing every  $h$ -th text position. It easily finds the aligned occurrences in  $O(m)$  time. The others can start up to  $h - 1$  positions after

a sampled position. Thus we search for all the patterns of the form  $\Sigma^i P$ ,  $0 \leq i < h$ . Overall this requires  $O(\sigma^{h-1}m + occ)$  time. By choosing  $h = 1 + \epsilon \log_\sigma n$  we get  $O(mn^\epsilon + occ)$  search time and  $O((n \log n)/h) = O(n \log \sigma)$  bits of space.

## 9.2 The LZ-Index of Kärkkäinen and Ukkonen (KU-FMI)

The *LZ-Index* of Kärkkäinen and Ukkonen [1996a] (KU-LZI) uses a suffix tree that indexes only the beginnings of phrases in a Ziv-Lempel parsing of  $T$ . This parsing is a variant of the original Ziv-Lempel parsing (Definition 16). Although they use the property that  $n' \leq n/\log_\sigma n$  to show that their index is succinct, requiring  $O(n \log \sigma)$  bits of space, Lemma 8 shows that they actually require  $O(nH_k)$  bits of space. We present here the results in their definitive form [Kärkkäinen 1999].

The original parsing has the property that each new phrase has already appeared in  $T$  (or it is a new character in  $\Sigma$ ). Thus, the first occurrence of any pattern  $P$  cannot be completely inside a phrase, otherwise it would have appeared before (the exception is  $m = 1$ , which is easy to handle and we disregard here). They divide occurrences among *primary* (spanning two or more phrases) and *secondary* (completely inside a phrase). Secondary occurrences are repetitions of other primary or secondary occurrences. Assume there are  $occ_p$  primary and  $occ_s$  secondary occurrences, so that  $occ = occ_p + occ_s$ .

Primary occurrences are found as follows. For some  $1 \leq i < m$ ,  $P_{1,i}$  is the suffix of a phrase and  $P_{i+1,m}$  starts at the next phrase. To avoid reporting the same occurrence multiple times, they insist in that  $P_{1,i}$  has to be completely included in a phrase, so the partitioning  $(P_{1,i}, P_{i+1,m})$  will be unique per occurrence. The occurrences of  $P_{i+1,m}$  are found using a sparse suffix tree that indexes the phrase beginnings (and is searched only for phrase-aligned occurrences). Those of  $P_{1,i}$  within a phrase are found using *RevTrie*, which is searched for  $P_i P_{i-1} \dots P_1$  (*RevTrie* is illustrated in Fig 26 for the LZ78 parsing). For example,  $P = \text{"labar"}$  appears twice in Fig 26: in phrase 2 with partition  $(\text{"1"}, \text{"abar"})$  and in phrase 9 with partition  $(\text{"lab"}, \text{"ar"})$ .

Each of those two searches yields a *lexicographical range* in  $[1, n']$ . The sparse suffix tree yields the range  $[l_2, r_2]$  of the phrase-aligned suffixes that start with  $P_{i+1,m}$ . The trie of reverse phrases gives the range  $[l_1, r_1]$  of phrases that finish with  $P_{1,i}$ . Consider now the  $p$ -th phrase. Assume  $Z[p]$  reversed is ranked  $x_p$ -th among all reversed phrases, and that the suffix starting at phrase  $p + 1$  is ranked  $y_p$ -th among all phrase-aligned suffixes. Then we wish to report a primary occurrence (with  $P_i$  aligned at the end of  $Z[p]$ ) iff  $(x_p, y_p) \in [l_1, r_1] \times [l_2, r_2]$ . This is a two-dimensional range search problem, where we store points  $(x_p, y_p)$  and search for the range  $[l_1, r_1] \times [l_2, r_2]$ . For example, for  $P = \text{"labar"}$  and  $i = 3$ , we find range  $[l_1, r_1] = [8, 8]$  for  $\text{"ba1"}$  in *RevTrie* (as node with label  $p = 9$  is 8th in a preorder traversal of *RevTrie*; this data can be stored in each node), and range  $[l_2, r_2] = [7, 8]$  for  $\text{"ar"}$  in the sparse suffix tree (as the 7th and 8th suffixes starting phrases start with  $\text{"ar"}$ ). Then we search for  $[8, 8] \times [7, 8]$  and find point  $(7, 8)$  corresponding to  $p = 9$ .

Secondary occurrences are obtained by tracking the source of each phrase  $Z[p]$ . Given a primary occurrence  $T_{j, j+m-1}$ , we wish to find all phrases  $p$  whose source contains  $[j, j+m-1]$ . Those phrases contain secondary occurrences  $T_{j', j'+m-1}$ , which are again tracked for new copies. With some slight changes to the original

parsing, it can be ensured that no source contains another and thus source intervals can be linearly ordered. An array  $S[p]$  of phrase numbers sorted by their source interval position in  $T$ , plus a bit array  $B[1, n]$  indicating which text positions start phrase sources, permits finding each phrase copying  $[j, j + m - 1]$  in constant time:  $S[\text{rank}_1(B, j)]$  is the last phrase in  $S$  whose source starts in  $T_{1,j}$ . We traverse  $S$  backwards from that position until the source intervals finish before  $T_{j+m-1}$ .

The index space is  $O(n' \log n) = O(nH_k)$  for the sparse suffix tree and the trie of reverse phrases. Among the range search data structures considered by Kärkkäinen [1999], we take those requiring  $O(\frac{1}{\epsilon} n' \log n')$  bits of space. Array  $S$  also needs the same space, and bit array  $B$  requires  $O(n' \log n)$  bits using the techniques of Section 5.2. Thus the overall space is  $O(\frac{1}{\epsilon} nH_k)$  bits, in addition to the text.

We note that this index carries out counting and locating simultaneously. The  $m - 1$  searches in *RevTrie* (for all  $P_{1,i}$  reversed) and sparse suffix tree (for all  $P_{i+1,m}$ ) require  $O(m^2)$  time. Finding the secondary occurrences takes constant time per occurrence retrieved, so it amounts to  $O(\text{occ}_s)$ . The remaining cost is that of the two-dimensional search. Depending on the structures used, different tradeoffs are obtained [Kärkkäinen 1999]. We give the most interesting ones for us.

**Theorem 20** *The LZ-Index of Kärkkäinen and Ukkonen [1996a] (KU-LZI) offers the following space/time tradeoffs.*

Space in bits	$O(\frac{1}{\epsilon} nH_k) + O(n \log \log n / \log n) + n \log \sigma$ (for text)
Time to count	$O(m^2 + mn^\epsilon + \text{occ})$
Time to locate	free after counting
Time to display $\ell$ chars	$O(\ell)$ (text is available)
Space in bits	$O(\frac{1}{\epsilon} nH_k) + O(n \log \log n / \log n) + n \log \sigma$ (for text)
Time to count	$O(m^2 + m \log n + \frac{1}{\epsilon} \text{occ} \log^\epsilon n)$
Time to locate	free after counting
Time to display $\ell$ chars	$O(\ell)$ (text is available)
Conditions for all	$k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$ ; $0 < \epsilon < 1$ (not necessarily constant)

The first term of counting complexities can be made  $O(m^2 / \log_\sigma n)$  by letting the tries move by  $O(\log_\sigma n)$  characters in one step, yet this raises the space requirement to  $O(n \log \sigma)$  unless we use much more recent methods [Grossi et al. 2003]. By using range search data structures that appeared later [Alstrup et al. 2000], the index would require  $O(nH_k \log^\gamma n)$  bits and count in  $O(m^2 + m \log \log n + \text{occ})$  time. Finally, we point out a variant of this idea [Kärkkäinen and Sutinen 1998] that can answer queries in  $O(m)$  counting time and  $O(1)$  locating time per occurrence, for short patterns ( $m < \log_\sigma n$ ).

### 9.3 The LZ-Index of Ferragina and Manzini (FM-LZI)

The *LZ-Index* of Ferragina and Manzini [2002] (FM-LZI) is the only existing self-index taking  $O(m)$  counting time and constant time to locate each occurrence. It is based on the LZ78 parse of  $T$  (Definition 17) and requires  $O(nH_k \log^\gamma n)$  bits of space for any constant  $\gamma > 0$ .

Let us define  $T^\#$  as text  $T$  where we have inserted special characters “#” after each phrase (so  $|T^\#| = n + n'$ ). For our example text  $T = \text{"alabar a la}$

alabarda\$" we have  $T^\# = \text{"a\#l\#ab\#ar\# \#a \#la\# a\#lab\#ard\#a\#\#"}$ . We also define  $T^R$  as text  $T^\#$  read backwards,  $T^R = \text{"#\$a\#dra\#bal\#a \#al\# a\# \#ra\#ba\#l\#a"}$ . Note that position  $t$  in  $T$ , belonging to the  $p$ -th phrase, corresponds to position  $rev(t, p) = (n - t + 1) + (n' - p + 1)$  in  $T^R$ . Let  $A$  be the suffix array of  $T$  and  $A^R$  that of  $T^R$ . Finally, let  $P^R = \#P_m P_{m-1} \dots P_1$ .

The FM-LZI consists of four components: (1) the FMI of text  $T$ ; (2) the FMI of text  $T^R$ ; (3) *LZTrie*, the Ziv-Lempel trie of  $T$ ; (4) *Range*, a two-dimensional range search data structure similar to that of the KU-LZI. The first three structures require  $O(nH_k)$  bits of space, yet *Range* will dominate the space complexity. As the FMI of  $T$  is enough for counting in  $O(m)$  time, we will focus in locating the *occ* occurrences in  $O(m + occ)$  time. Occurrences of  $P$  are divided into primary and secondary as in Section 9.2 (they are called “external” and “internal” by Ferragina and Manzini [2002]).

Let us first consider secondary occurrences. Since every prefix of a phrase is also a phrase, every secondary occurrence which is not at the end of its phrase  $p$  occurs also in the phrase  $p'$  referenced by  $p$  (that is, the parent of  $p$  in the *LZTrie*). For example, in Fig. 26, pattern  $P = \text{"a"}$  occurs in phrase 10. Since it does not occur at the end of  $Z[10] = \text{"ard"}$ , it must also occur in its parent 4,  $Z[4] = \text{"ar"}$  and in turn in its parent 1,  $Z[1] = \text{"a"}$ . Let us call a trie node  $p$  *pioneer* for  $P$  if  $P$  is a suffix of  $Z[p]$ . In our example the pioneer nodes for  $P = \text{"a"}$  are 1, 7, and 8. Then, all secondary occurrences correspond to *LZTrie* nodes that descend from pioneer ones (including themselves). To find the secondary occurrences, it is sufficient to obtain the pioneer nodes and then traverse all their subtrees reporting all the text positions found (with the appropriate offsets).

Finding the pioneer nodes is easy using the FMI of  $T^R$ . It is a matter of searching for  $P^R$ , as that corresponds to occurrences of  $P$  that are phrase suffixes, or which is the same, to occurrences of  $P^\#$  in  $T^\#$ . For example, if we search for  $P^R = \text{"\#a"}$  in  $T^R$  we will find occurrences at positions 12, 15, and 31 of  $T^R$ . This corresponds to the occurrences of  $\text{"a\#"}$  in  $T^\#$ , at positions 20, 17, and 1, respectively. Aligned to the (contiguous) area of  $A^R$  corresponding to suffixes that start with “#”, we store a vector  $S$  of pointers to the corresponding *LZTrie* nodes. As the range for  $P^R$  is always contained in the area covered by  $S$ , vector  $S$  permits finding the pioneer nodes of the results of the search. As  $S$  has  $n'$  entries, it occupies  $nH_k + o(n)$  bits. Thus the  $occ_s$  secondary occurrences are reported in  $O(m + occ_s)$  time. Fig. 27 illustrates.

Let us now consider the primary occurrences. The same idea of Section 9.2 is used to find  $P_{1,i}$  at the end of a phrase and  $P_{i+1,n}$  from the next phrase. Yet, the search proceeds differently, and the FMI is shown to be more fortunate for this problem. We first search for  $P$  using the FMI of  $T$ . This single search gives us all the ranges  $[sp_{i+1}, ep_{i+1}]$  in  $A$  corresponding to the occurrences of  $P_{i+1,m}$ ,  $1 \leq i < m$  (recall Section 8.1). We now search for  $P^R$  in the FMI of  $T^R$ . After we have each range  $[sp'_i, ep'_i]$  corresponding to the occurrences of  $P_i P_{i-1} \dots P_1$  in  $A^R$ , we add character “#”, obtaining the range  $[sp_i^R, ep_i^R]$  of the occurrences of  $\#P_i P_{i-1} \dots P_1$  in  $A^R$ ,  $1 \leq i < m$ . This corresponds to occurrences of  $P_{1,i}^\#$  in  $T^\#$ . Note that the search in  $T^R$  ensures that  $P_{1,i}$  is completely contained in a phrase (and is at the end of it), while the search in  $T$  permits  $P_{i+1,m}$  span as many phrases as necessary. All this process takes  $O(m)$  time.

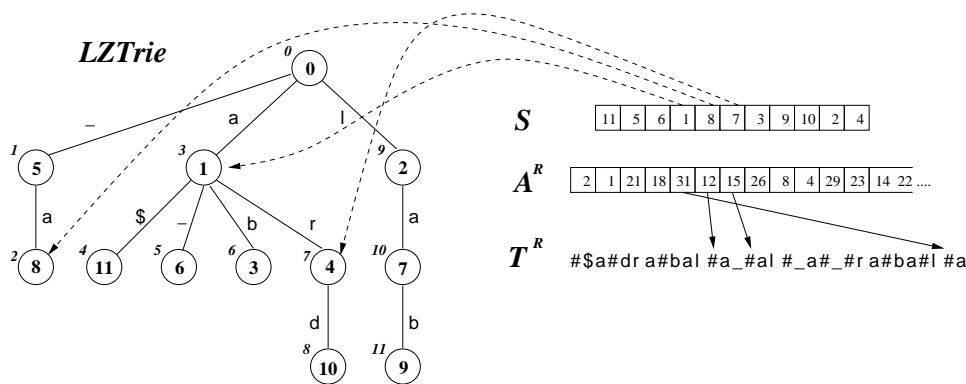


Fig. 27. Parts of the FM-LZI index. We show  $T^R$  and  $A^R$  (none of which is explicitly represented), as well as vector  $S$  and  $LZTrie$ . Only the part of  $A^R$  pointing to suffixes starting with “#” is shown in detail. This is the part  $S$  is aligned to. For legibility,  $S$  shows phrase numbers instead of pointers to  $LZTrie$ . The result of the search for “#a” is illustrated with the actual pointers from  $A^R$  to  $T^R$  and from  $S$  to  $LZTrie$ .

Consider searching for  $P = \text{"bar"}$ . There are  $m - 1 = 2$  possible partitions for  $P$ , (“b”, “ar”) and (“ba”, “r”). Those appear in  $T^\#$  as “b#ar” and “ba#r”. Using the FMI of  $T$  we get  $[sp_3, ep_3] = [20, 21]$  ( $A$  range for “r”), and  $[sp_2, ep_2] = [12, 13]$  ( $A$  range for “ar”), see Figs. 3 and 6. Using the FMI of  $T^R$  we get  $[sp_3^R, ep_3^R] = [8, 9]$  ( $A^R$  range for “#b”, corresponding to “b#” in  $T^\#$ ), and we get that  $[sp_2^R, ep_2^R]$  ( $A^R$  range for “#ab”, corresponding to “ba#” in  $T^\#$ ) is empty, see Fig. 27. Thus, we know that it is possible to find primary occurrences corresponding to partition (“b”, “ar”). The first part corresponds to  $A^R[8, 9]$  and the second to  $A[12, 13]$ . Looking at Figs. 3 and 27, we see that  $A^R[8] = 26$  in  $T^R$  matches with  $A[12] = 5$  in  $T$  and  $A^R[9] = 8$  in  $T^R$  matches with  $A[13] = 17$  in  $T$ . For example,  $t = 5$  in  $T$  matches with 26 in  $T^R$  because it belongs to the 4th phrase and thus  $t$  corresponds to  $rev(5, 4) = (21 - 5 + 1) + (11 - 4 + 1) = 25$  in  $T^R$ .

Structure *Range*, storing  $n'$  points in  $[1, n] \times [1, n]$ , is used to find those matching pairs. Let  $j = A^{-1}[t]$  be the position in  $A$  pointing to  $T_{t,n}$ , where  $t$  starts the  $p$ -th phrase in  $T$ . Similarly, let  $j^R = (A^R)^{-1}[rev(t) + 1]$  be the position in  $A^R$  pointing to  $T_{rev(t)+1,n}^R$  (which represents  $T_{1,t-1}$ ). We store pairs  $(j^R, j)$  in *Range*. A range search for  $[sp_i^R, ep_i^R] \times [sp_{i+1}, ep_{i+1}]$  retrieves all those phrase positions  $t$  such that  $P_{1,i}$  is completely included in the phrase preceding position  $t$  and  $P_{i+1,m}$  follows from  $T_t$ . Thus we report text positions  $t - i + 1$ , where each occurrence  $P$  starts. In our example, two points we would store are (8, 12) and (9, 13), corresponding to  $t = 5$  and  $t = 17$  in  $T$ . These will be retrieved by range query  $[8, 9] \times [12, 13]$ .

They use a two-dimensional data structure [Alstrup et al. 2000] that can store  $n'$  points in  $[1, n'] \times [1, n']$  using  $O(n' \log^{1+\gamma} n')$  bits for any  $\gamma > 0$ , so that they answer a query with  $res$  results in time  $O(\log \log n' + res)$ . In our case, we must query the structure once per each partition  $1 \leq i < m$ , so we pay overall  $O(m \log \log n + occ_p)$ . Note that our points are actually in  $[1, n] \times [1, n]$ . Those can be mapped to  $[1, n'] \times [1, n']$  using *rank* and *select* on bitmaps of length  $n$  with  $n'$  bits set. Using the techniques of Section 5.2, those bitmaps require  $O(n' \log n) = O(nH_k)$  bits.

The  $O(m \log \log n)$  time can be improved as follows. Basically, instead of storing only the positions  $t$  that start a phrase in *Range*, we add all positions  $[t - \log \log n + 1, t]$ . Now each cut  $(P_{1,i}, P_{i+1,m})$  would be found  $\log \log n$  times, not once. Thus we can search only for those  $i$  that are multiple of  $\log \log n$ . As we perform only  $m / \log \log n$  queries, the overall time is  $O(m + occ_p)$ . Although now we store  $n' \log \log n$  points in *Range*, its space stays  $O(nH_k \log^\gamma n)$  for any  $\gamma > 0$ . We omit some technical details to handle borders between phrases.

For patterns shorter than  $\log \log n$  we must use a different approach. Those patterns are so short that we can precompute all their primary occurrences with a four-Russians technique. There are at most  $\sigma^{\log \log n} = (\log n)^{\log \sigma}$  different short patterns, each requiring a pointer of  $\log n$  bits to its occurrence list, and the total number of primary occurrences for all short patterns is  $n'(\log \log n)^2$ , each requiring  $\log n$  bits as well. Overall, say for  $\sigma = O(\text{polylog}(n))$ , this requires  $o(nH_k \log^\gamma n + n^\beta)$  bits, for any constants  $\gamma > 0$  and  $0 < \beta < 1$ .

By using the AF-FMI rather than the original FMI, we obtain the following result. We consider the second variant of Theorem 19, as it is the most interesting for our purposes.

**Theorem 21** *The LZ-Index of Ferragina and Manzini [2002] (FM-LZI) offers the following space/time tradeoffs.*

Space in bits	$O(nH_k \log^\gamma n + n / \log^\epsilon n)$
Time to count	$O(m)$
Time to locate	$O(1)$
Time to display $\ell$ chars	$O(\ell + \log^{1+\epsilon} n)$
Condition	$\gamma > 0$ is any constant
Space in bits	$O(nH_k \log \log n + n / \log^\epsilon n)$
Time to count	$O(m)$
Time to locate	$O(\log \log n)$
Time to display $\ell$ chars	$O(\ell + \log^{1+\epsilon} n)$
Conditions for all	$\sigma = O(\text{polylog}(n));$ $k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$ ; $0 < \epsilon < 1$ is any constant

The second version is due to other results by Alstrup et al. [2000], which need  $O(n' \log n' \log \log n')$  bits of space and can search in time  $O((\log \log n)^2 + res \log \log n)$ . We can retain  $O(m)$  time by indexing  $(\log \log n)^2$  positions per phrase instead of  $\log \log n$ .

The two-dimensional range search idea has inspired other solutions to achieve constant time per occurrence on compressed suffix arrays [He et al. 2005], yet those work only for sufficiently large  $m$ .

#### 9.4 The LZ-Index of Navarro (NAV-LZI)

The *LZ-Index* of Navarro [2002, 2004] (NAV-LZI) uses essentially *LZTrie* and *RevTrie* (Definition 18), under LZ78 parsing. Despite that its complexity is not competitive, the structure is interesting because it is the only self-index not using at all the suffix array concept.



The NAV-LZI uses four structures: (1) *LZTrie*, the Ziv-Lempel trie of  $T$ , (2) *RevTrie*, storing the reverse of strings  $Z[p]$ , (3) the same two-dimensional *Range* data structure of Section 9.2, (4) *Node*, a mapping from phrase numbers to *LZTrie* nodes.

To save space, *LZTrie* and *RevTrie* tree shapes are represented using parentheses rather than full pointers. The tree shape of Fig. 26 is expressed by a preorder traversal as " $((()()())((()))())$ ", or  $Tr = 111010110010011100011000$  in bits. Munro and Raman [1997] show how a tree of  $n'$  nodes can be represented using  $2n' + o(n')$  bits so that it can be traversed in constant time per operation, by identifying each node with its open parenthesis. Some operations implemented in constant time are: first child, next sibling, parent, depth, preorder position, etc. For example, the first child of node at position  $t$  is  $t+1$  unless  $Tr[t+1] = 0$ , in which case the node is a leaf. The depth of node at position  $t$  is  $rank_1(Tr, t) - rank_0(Tr, t)$ . Other operations are more complex to implement, but the ideas derive from those in Section 5.1.

In addition to the tree space, *LZTrie* has to store the characters labeling the edges and the phrase number of each node. These are stored on arrays, in preorder traversal. In our example, the characters array is " $*a\$brd lab a$ " and the phrase numbers array is  $\{0, 1, 11, 3, 4, 10, 6, 2, 7, 9, 5, 8\}$ . Let  $pos(v)$  be the preorder position of node  $v$ ,  $pos(v) = rank_1(Tr, t)$ , where  $t$  represents node  $v$ . For example, the node  $v$  representing string  $S^v = "ar"$  corresponds to the open parenthesis at  $Tr[7]$ . It is the 5th tree node in preorder traversal,  $pos(v) = rank_1(Tr, 7) = 5$ . With position 5, we know that node  $v$  descends by character "r" from its parent, and that its phrase number is 4 (see Fig. 26). Note also that, if  $v$  is represented by  $Tr[t]$  and  $t'$  is the closing parenthesis that matches  $t$ , then  $[rank_1(Tr, t), rank_1(Tr, t')]$  is the interval of positions of all nodes that descend from  $v$  in the trie. For our example node, this interval is  $[rank_1(Tr, 7), rank_1(Tr, 10)] = [5, 6]$ , which contains  $v$  itself as well as node representing phrase 10, "ard".

*RevTrie* is a bit more complicated. Since there are *RevTrie* nodes not corresponding to any phrase (see Fig. 26), there could be more than  $n'$  nodes in *RevTrie*. To limit its size, we compress unary paths as when moving to suffix tries to suffix trees (Section 3.2). This ensures having at most  $2n'$  nodes in *RevTrie*. Clark and Munro [1996] show how a suffix tree can be represented with parentheses plus some additional information regarding the first character of each edge label and its length (like a Patricia tree [Morrison 1968]). The NAV-LZI is more extreme, as nothing apart from the tree shape (using parentheses) and the array of phrase numbers is stored. Let  $p$  be the phrase number of node  $v^R$  in *RevTrie*, representing string  $S^{v^R}$ . Then  $v = Node(p)$  is the corresponding *LZTrie* node for phrase  $p$ . Using the parent traversal operation in *LZTrie* one can obtain  $S^v$  in reverse order, and this is precisely  $S^{v^R}$ . This is the information we need to traverse *RevTrie*. Note that looking for a string  $S$  in *RevTrie* costs  $O(|S|^2 \log \sigma)$ , while the cost is  $O(|S| \log \sigma)$  in *LZTrie*. (We omit some technicalities used to obtain logarithmic rather than linear dependence on  $\sigma$ .)

Occurrences are classified according to how many phrases they span. Occurrences of type 1 are contained in a phrase, those of type 2 are contained in a pair of consecutive phrases, and those of type 3 span three phrases or more. Thus we have

$occ = occ_1 + occ_2 + occ_3$  occurrences to count and locate.

Occurrences of type 1 are found with the same idea for secondary occurrences in Section 9.3, yet with a different mechanism. We search for  $P$  reversed in  $RevTrie$ . Say we arrive at node  $v$ . Any node  $v'$  descending from  $v$ , including  $v' = v$  itself, corresponds to a phrase terminated with  $P$ . Thus  $Node(v')$  is a pioneer  $LZTrie$  node. Thus we traverse and report all subtrees of all those  $Node(v')$  nodes in  $LZTrie$ , so occurrences of type 1 are counted and located in  $O(m^2 \log \sigma + occ_1)$  time.

Occurrences of type 2 are found with the range search data structure, as in Section 9.3. This time we search for  $P_{1,i}$  reversed in  $RevTrie$ , and for  $P_{i+1,m}$  in  $LZTrie$ , obtaining two nodes  $v_{rev}$  and  $v_{lz}$ , respectively. We are interested in phrases  $p$  such that the node of phrase  $p$  descends from  $v_{rev}$  and that of  $p+1$  descends from  $v_{lz}$ . We store in  $Range$  all points  $(pos(v'), pos(v))$  for each phrase  $1 \leq p < n'$ , so that  $v' \in RevTrie$  has phrase number  $p$  and  $v \in LZTrie$  has phrase number  $p+1$ . Thus a range search for the array areas that descend from  $v_{rev}$  and  $v_{lz}$  retrieves those occurrences we want. To implement  $Range$ , a variant of a data structure [Chazelle 1988; Kärkkäinen 1999] is used which can store  $n'$  points in a grid  $[1, n'] \times [1, n']$  in  $n' \log n'(1 + o(1))$  bits, and can find the  $res$  results of a query in  $O((1 + res) \log n')$  time. Thus occurrences of type 2 cost  $O(m^3 \log \sigma + (m + occ_2) \log n)$  time.

Finally, for occurrences of type 3, the fact that every LZ78 phrase is unique is exploited. Since an occurrence of type 3 must contain a whole phrase, there cannot be more than  $m(m-1)/2 \geq occ_3$  such occurrences because this is the number of substrings of  $P$ . Essentially, the phrase equal to each substring of  $P$  is found using  $LZTrie$  and extended to a full occurrence if possible, in overall time  $O(m^2 \log \sigma + m^3)$ .

Overall, the query time is upper bounded by  $O(m^3 \log \sigma + (m + occ) \log n)$ . The space is dominated by the phrase arrays of  $LZTrie$  and  $RevTrie$ , the  $Range$  structure, and the  $Node$  mapping. Each of these requires  $n' \log n'(1 + o(1))$  bits, which makes the index space  $4nH_k(1 + o(1))$  bits.

We note that in this index counting is not easier than locating. Just like the KU-LZI, both process must be carried out simultaneously. In order to display a text substring, we first locate its phrase  $p$  and then traverse  $LZtrie$  upwards from  $Node(p)$  to obtain all its characters in reverse order. If we need more phrases, we repeat the process with  $p-1$  or  $p+1$ , and so on. This takes constant time per character, yet it displays whole phrases. The extra work to display partial phrases is  $O(n/n') = O(\log(n)/H_k)$  on average.

**Theorem 22** *The LZ-Index of Navarro [2002, 2004] (NAV-LZI) offers the following space/time tradeoff.*

Space in bits	$4nH_k(1 + o(1)) + O(n \log \log n / \log n)$
Time to count	$O(m^3 \log \sigma + (m + occ) \log n)$
Time to locate	<i>free after counting</i>
Time to display $\ell$ chars	$O(\ell + \log(n)/H_k)$ (last term is on average)
Conditions	$k \leq \alpha \log_\sigma n$ , for some constant $0 < \alpha < 1$

Note that, by using the range search data structure of Section 9.3 [Alstrup et al.

2000], the space would be  $O(nH_k \log^\gamma n)$  and the counting time would decrease to  $O(m^3 \log \sigma + m \log \log n + occ)$ .

*In practice:* The implementation of the NAV-LZI has several differences with respect to the theoretical proposal. The most important one is that *Range* structure is replaced by *RNode*, a mapping from phrase numbers to *RevTrie* nodes. Now occurrences of type 2 are found as follows: Check each phrase  $p$  in the range descending from  $v_{rev}$  and report it if  $p + 1$  descends from  $v_{lz}$ . Each such check takes constant time with *Node*. Yet, if the range of  $v_{lz}$  has less elements, do the opposite: check phrases from  $v_{lz}$  in  $v_{rev}$ , using *RNode*. *RNode* is also useful to reduce the usage of the slower *RevTrie* to a minimum: In every case where we can search for the reverse string in *LZTrie*, we avoid searching *RevTrie* as we can map results using *RNode*.

## 10. DISCUSSION

We have presented the main ideas of several compressed indexes as intuitively as possible, yet with an accuracy enough to understand the space/time tradeoffs they achieve. In many cases, these depend on several parameters in a complex way, which makes a fair comparison difficult. Just as an overview, we provide a summary in Table 1.

Table 1. Simplified space and time complexities for compressed full-text indexes. For space complexities, we present only the main term related to entropies and seek to minimize space. For time complexities, we present bounds that hold on some representative inputs, assuming for example that the alphabet is small. We refer to the theorems given earlier for accurate statements and boundary conditions.

<i>Index</i>	entropy term	time to count
GV-CSA	$nH_0$	$O(m \log^2 n)$
MAK-CSA	$2nH_k \log n$	$O(m \log n + \text{polylog}(n))$
SAD-CSA	$nH_0$	$O(m \log n)$
MN-CCSA	$nH_k \log n$	$O(m \log n)$
GGV-CSA	$nH_k$	$O(m \log \sigma + \text{polylog}(n))$
FMI	$5nH_k$	$O(m)$
HUFF-FMI	$2nH_0$	$O(mH_0)$
WT-FMI	$nH_0$	$O(m)$
RL-FMI	$nH_k \log \sigma$	$O(m)$
AF-FMI	$nH_k$	$O(m)$
KU-LZI	$O(nH_k)$	$O(m^2 + (m + occ) \log n)$
FM-LZI	$O(nH_k \log^\gamma n)$	$O(m)$
NAV-LZI	$4nH_k$	$O(m^3 \log \sigma + (m + occ) \log n)$

It is interesting at this point to discuss the most important common points in these approaches. Common points within the CSA family and within the FMI family are pretty obvious, namely they are basically different implementations of functions  $\Psi$  and  $Occ$ , respectively. There is also an obvious relation among these two families, as  $\Psi$  and  $LF$  are the inverse of each other (Lemma 4).

Which is more subtle is the relation between the different ways to achieve  $O(nH_k)$  space. Let us exclude the Ziv-Lempel based methods, as they are totally different.

For this discussion, the table of Fig. 16 is particularly enlightening. Each number  $i \in [1, n]$  can be classified according to two parameters: the list (or table column)  $T_{A[i]-1}$  and the context (or table row)  $T_{A[i], A[i+k-1]}$  where  $i$  belongs. Within a given row and column, numbers  $i$  (sharing  $T_{A[i]-1, A[i+k-1]}$ ), are classified according to  $T_{A[i+k], n}$ .

As  $A$  is sorted by  $T_{A[i], n}$  (a refinement of the  $k$ -th context order), all the  $i$  values in each table row form a contiguous subinterval of  $[1, n]$ , which advances with the row number. How do the  $i$  values within each row (corresponding to contexts  $T_{A[i], A[i+k-1]}$ ) distribute across columns, depends on  $T_{A[i]-1}$ , the characters preceding the occurrences of the context in  $T$ .

Instead of row-then-column, consider now a column-then-row order. Now  $i$  values are collected in the order given by  $T_{A[i]-1, n}$ , or renaming  $i = \Psi(j)$  (as  $\Psi$  is a permutation),  $\Psi(j)$  values are collected in the order given by  $T_{A[\Psi(j)]-1, n} = T_{A[j], n}$ . This order is of course  $j = 1, 2$ , and so on, thus we are in fact reading list  $\Psi$ . Numbers  $i$  are increasing inside each column because they are ordered by  $T_{A[i], n}$ .

The SAD-CSA structure stores  $\Psi$  in order, that is, it stores the table in column-wise order, and then row-wise inside each column. Being  $\sigma$  increasing lists, this leads to  $O(nH_0)$  space. The GGV-CSA, instead, stores  $\Psi$  in row-wise order. For this sake, it needs to record how the values inside each row distribute across columns. According to Theorem 2, it is sufficient to store that distribution information in space close to its zero-order entropy, to achieve  $nH_k$  overall space. They use the wavelet tree as a device to represent the information on which column does each row element belong to.

In a widely different view, the AF-FMI structure stores  $T^{bwt}$  context-wise (that is, row-wise in the table). For each context, it stores the characters of  $T^{bwt}$ , which are precisely  $T_i^{bwt} = A[i] - 1$ , that is, the column identifiers of the positions  $i$  lying within each context (row). They use the same wavelet tree to represent basically the same data within the same zero-order entropy space. Thus both structures are using essentially the same concept to achieve  $nH_k$  space.

The differences are due to other factors. While the GGV-CSA structure still adheres to the idea of abstract optimization of  $A$ , so that they must provide access to  $A[i]$  and use the normal binary search on  $A$ , the FMI family uses a completely different form of searching, which directly relies on  $T^{bwt}$ .

The final practical twist of the GGV-CSA is the discovery that  $\gamma$ - or  $\delta$ -encoding of consecutive  $i$  values within each cell of the table yields  $O(nH_k)$  space, independently of whether one uses row-wise or column-wise order (as there are not too many jumps across table cells). This permits a much simpler implementation of the structure, which turns out to be close to the SAD-CSA, initially believed to require  $O(nH_0)$  space.

A completely different mechanism to achieve  $O(nH_k)$  space is used in the other indexes. These rely on compressing the runs that appear in  $\Psi$ , or alternatively those that appear in  $T^{bwt}$ . The former achieve  $O(nH_k \log n)$  space by emulating the binary search on  $A$  through  $\Psi$ , whereas the latter achieve  $O(nH_k \log \sigma)$  space by emulating an FMI strategy.

Table 2 classifies the approaches that reach  $H_k$ -related space according to their approach. In one dimension, we have those based on local entropy (Theorem 2) versus those based on run lengths (Theorem 3). In the other dimension, we have

those based on  $\Psi$  or on  $Occ$ . We have included SAD-CSA as having size  $nH_k$  according to the results of Grossi et al. [2004]. We classify the FMI as using run lengths because this is the key property ensuring its  $O(nH_k)$  size, although it also uses some local entropy optimization. Remind that we left aside Ziv-Lempel methods in this discussion and in the table.

Table 2. Classifying the indexes with size related to  $H_k$ . Names are followed by a pair indicating (space,time) complexities. These are even more simplified than in Table 1.

	local entropy	run lengths
using $\Psi$	SAD-CSA ( $nH_k, m \log n$ ) GGV-CSA ( $nH_k, m \log \sigma$ )	MAK-CSA ( $2nH_k \log n, m \log n$ ) MN-CCSA ( $nH_k \log n, m \log n$ )
using $Occ$	AF-FMI ( $nH_k, m$ )	FMI ( $5nH_k, m$ ) RL-FMI ( $nH_k \log \sigma, m$ )

## 11. CONCLUSIONS

We have given a unified look at the state of the art in compressed full-text indexing. We focused on the essential ideas relating text compressibility and regularities on indexes built on it, and uncovered fundamental relations between seemingly disparate approaches. Those approaches have lead to a rich family of results, whose most important consequence is a surprising fact of text compressibility:

**Fact.** *Instead of compressing a text into a representation that does not reveal anything from the original text unless decompressed, one can obtain an equally space-efficient representation that at the same time works as a versatile index on the text.*

In other words, the indexes we have reviewed take space close to what can be obtained by the best possible compressors, both in theory and in practice. In theory, the leading term in the space complexities of the best indexes is  $nH_k$ , which is a widely accepted lower-bound estimate for text compressibility. For substring searches, the same best indexes are practically optimal, obtaining  $O(m)$  counting query time. This remarkable discovery is without any doubt one of the most important achievements ever obtained in text compression and text indexing.

However, the theory is not yet complete, as there are several open questions to be answered. A first open question is whether one can obtain  $nH_k$  space and  $O(m)$  query time on any alphabet size, and in general which is the lower bound relating these parameters. One easily notices that obtaining  $nH_k$  without any constraints is impossible as  $nH_k = 0$  for large  $k$  (in particular for  $k = n$ ). Another example is a text consisting of a permutation of  $[1, n]$ , where  $\sigma = n$ . Although  $H_k = 0$  for  $k \geq 1$ ,  $\Omega(n \log n)$  bits are necessary to represent such text. Still, there is a gap between what is obtained and what is possible. This statement also applies to the sublinear terms that might be bigger than the entropy-related part. However, recent lower bounds [Miltersen 2005] on *rank* and *select* dictionaries seem to ensure that not much progress in that direction is anymore possible.

Another open challenge is to obtain better output sensitivity in reporting queries within little space. For this goal, there are some results achieving  $O(occ+o(n))$  time for large enough  $m$  [Grossi and Vitter 2000, 2005],  $O(occ)$  time for large enough  $m$  [He et al. 2005], and even  $O(occ)$  time without any restriction on  $m$ , for small alphabets, using  $O(nH_k \log^\gamma n)$  bits of space [Ferragina and Manzini 2002]. The technique by He et al. [2005] is general and can be plugged into any of the indexes discussed before, by adding some sublinear dictionaries.

In this survey we have focused on the most basic problem, namely the exact search problem in main memory. There are many more further challenges, with regard to more complex searching, index construction and updating issues, secondary memory, and so on. A brief list of other relevant problems beyond the scope of this survey follows.

**Secondary memory:** Although their small space requirements might permit compressed indexes fit in main memory, there will always be cases where they have to operate on disk. There is not much work yet on this important issue. One of the most attractive full-text indexes for secondary memory is the String B-tree [Ferragina and Grossi 1999]. This is not, however, a succinct structure. Some proposals for succinct and compressed structures in this scenario exist [Clark and Munro 1996; Mäkinen et al. 2004]. A good survey on full-text indexes in secondary memory is by Kärkkäinen and Rao [2003].

**Construction:** Compressed indexes are usually derived from a classical one. Although it is usually simple to build a classical index and then derive its compressed version, there might not be enough space to build the classical index first. Secondary memory might be available, but many classical indexes are costly to build in secondary memory. Therefore, an important problem is how to build compressed indexes without building their classical versions first. Several papers have recently appeared on the problem of building the SAD-CSA in little space [Lam et al. 2002; Hon et al. 2003; Hon et al. 2003; Na 2005]. There is also some recent work on efficient construction of (plain) suffix arrays [Manzini and Ferragina 2004] and trees [Farach 1997]. With respect to construction of (plain) indexes in secondary memory, there is a good experimental comparison for suffix arrays [Crauser and Ferragina 2002], as well as some work on suffix trees [Farach et al. 1998; Clifford and Sergot 2003].

**Dynamism:** Most indexes considered are static, in the sense that they have to be rebuilt from scratch upon text changes. This is currently a problem even on uncompressed full-text indexes, and not much has been done. Some recent work on compressed indexes can be found [Ferragina and Manzini 2000; Hon et al. 2004; Chan et al. 2004].

**Extended functionality:** We have considered only exact string matching in this survey, yet classical full-text indexes permit much more sophisticated search tasks [Apostolico 1985; Gusfield 1997]. There has been a considerable amount of work on extending compressed suffix arrays functionalities to those of suffix trees [Grossi and Vitter 2000; Munro et al. 2001; Sadakane 2002; Sadakane 2003; Grossi et al. 2004; Kim and Park 2005; Grossi and Vitter 2005]. The idea in general is to permit the simulation of suffix tree traversals using a compressed representation of them, such as a compressed suffix array plus a parentheses

representation of the suffix tree shape [Munro and Raman 1997]. In addition, there has been some work on approximate string matching over compressed suffix arrays [Huynh et al. 2004]. Finally, it is also interesting to mention that the idea of backward searching has been used to search plain suffix arrays in  $O(m \log \sigma)$  time [Sim et al. 2003].

Overall, we believe that self-indexing is one of the most exciting research directions within text compression and text indexing, which in few years has obtained striking results and has a long way ahead, rich of challenges and possibly of new surprises.

## REFERENCES

- ABOUELHODA, M., OHLEBUSCH, E., AND KURTZ, S. 2002. Optimal exact string matching based on suffix arrays. In *Proc. 9th International Symposium on String Processing and Information Retrieval (SPIRE)*, LNCS v. 2476 (2002), pp. 31–43.
- ALSTRUP, S., BRODAL, G., AND RAHUE, T. 2000. New data structures for orthogonal range searching. In *Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS)* (2000), pp. 198–207.
- ANDERSSON, A. AND NILSSON, S. 1995. Efficient implementation of suffix trees. *Software Practice and Experience* 25, 2, 129–141.
- APOSTOLICO, A. 1985. The myriad virtues of subword trees. In *Combinatorial Algorithms on Words*, NATO ISI Series (1985), pp. 85–96. Springer-Verlag.
- ARLAZAROV, V., DINIC, E., KONROD, M., AND FARADZEV, I. 1975. On economic construction of the transitive closure of a directed graph. *Soviet Mathematics Doklady* 11, 1209–1210.
- BAEZA-YATES, R. AND RIBEIRO, B. 1999. *Modern Information Retrieval*. Addison-Wesley.
- BELL, T., CLEARY, J., AND WITTEN, I. 1990. *Text compression*. Prentice Hall.
- BENTLEY, J., SLEATOR, D., TARJAN, R., AND WEI, V. 1986. A locally adaptive compression scheme. *Communications of the ACM* 29, 4, 320–330.
- BLUMER, A., BLUMER, J., HAUSSLER, D., MCCONNELL, R., AND EHRENFEUCHT, A. 1987. Complete inverted files for efficient text retrieval and analysis. *Journal of the ACM* 34, 3, 578–595.
- BURROWS, M. AND WHEELER, D. 1994. A block sorting lossless data compression algorithm. Technical Report 124, Digital Equipment Corporation.
- CHAN, H.-L., HON, W.-K., AND LAM, T.-W. 2004. Compressed index for a dynamic collection of texts. In *Proc. 15th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 3109 (2004), pp. 445–456.
- CHAZELLE, B. 1988. A functional approach to data structures and its use in multidimensional searching. *SIAM Journal on Computing* 17, 3, 427–462.
- CLARK, D. 1996. *Compact Pat Trees*. Ph. D. thesis, University of Waterloo, Canada.
- CLARK, D. AND MUNRO, I. 1996. Efficient suffix trees on secondary storage. In *Proc. 7th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (1996), pp. 383–391.
- CLIFFORD, R. AND SERGOT, M. 2003. Distributed and paged suffix trees for large genetic databases. In *Proc. 14th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 2676 (2003), pp. 70–82.
- COLUSSI, L. AND DE COL, A. 1996. A time and space efficient data structure for string searching on large texts. *Information Processing Letters* 58, 5, 217–222.
- COVER, T. AND THOMAS, J. 1991. *Elements of Information Theory*. Wiley.
- CRAUSER, A. AND FERRAGINA, P. 2002. A theoretical and experimental study on the construction of suffix arrays in external memory. *Algorithmica* 32, 1, 1–35.
- CROCHEMORE, M. AND VÉRIN, R. 1997. Direct construction of compact directed acyclic word graphs. In *Proc. 8th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 1264 (1997), pp. 116–129.

- ELIAS, P. 1975. Universal codeword sets and representation of the integers. *IEEE Transactions on Information Theory* 21, 2, 194–203.
- FARACH, M. 1997. Optimal suffix tree construction with large alphabets. In *Proc. 38th IEEE Symposium on Foundations of Computer Science (FOCS)* (1997), pp. 137–143.
- FARACH, M., FERRAGINA, P., AND MUTHUKRISHNAN, S. 1998. Overcoming the memory bottleneck in suffix tree construction. In *Proc. 39th IEEE Symposium on Foundations of Computer Science (FOCS)* (1998), pp. 174–185.
- FERRAGINA, P. AND GROSSI, R. 1999. The string B-tree: A new data structure for string search in external memory and its applications. *Journal of the ACM* 46, 2, 236–280.
- FERRAGINA, P. AND MANZINI, G. 2000. Opportunistic data structures with applications. In *Proc. 41st IEEE Symposium on Foundations of Computer Science (FOCS)* (2000), pp. 390–398.
- FERRAGINA, P. AND MANZINI, G. 2001. An experimental study of an opportunistic index. In *Proc. 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2001), pp. 269–278.
- FERRAGINA, P. AND MANZINI, G. 2002. On compressing and indexing data. Technical Report TR-02-01, Dipartimento di Informatica, University of Pisa, Italy. To appear in *Journal of the ACM*.
- FERRAGINA, P. AND MANZINI, G. 2004. Compression boosting in optimal linear time using the Burrows-Wheeler transform. In *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2004), pp. 655–663.
- FERRAGINA, P., MANZINI, G., MÄKINEN, V., AND NAVARRO, G. 2004a. An alphabet-friendly FM-index. In *Proc. 11th International Symposium on String Processing and Information Retrieval (SPIRE)*, LNCS v. 3246 (2004), pp. 150–160.
- FERRAGINA, P., MANZINI, G., MÄKINEN, V., AND NAVARRO, G. 2004b. Compressed representation of sequences and full-text indexes. Technical Report 2004-05 (Dec.), Technische Fakultät, Universität Bielefeld, Germany.
- FERRAGINA, P., MANZINI, G., MÄKINEN, V., AND NAVARRO, G. 2004c. Succinct representation of sequences. Technical Report TR/DCC-2004-5 (Aug.), Department of Computer Science, University of Chile, Chile.
- FREDKIN, E. 1960. Trie memory. *Communications of the ACM* 3, 490–500.
- GIEGERICH, R., KURTZ, S., AND STOYE, J. 1999. Efficient implementation of lazy suffix trees. In *Proc. 3rd Workshop on Algorithm Engineering (WAE)*, LNCS v. 1668 (1999), pp. 30–42.
- GONNET, G., BAEZA-YATES, R., AND SNIDER, T. 1992. *Information Retrieval: Data Structures and Algorithms*, Chapter 3: New indices for text: Pat trees and Pat arrays, pp. 66–82. Prentice-Hall.
- GONZÁLEZ, R., GRABOWSKI, S., MÄKINEN, V., AND NAVARRO, G. 2005. Practical implementation of rank and select queries. In *Poster Proceedings Volume of 4th Workshop on Efficient and Experimental Algorithms (WEA'05)* (Greece, 2005), pp. 27–38. CTI Press and Ellinika Grammata.
- GRABOWSKI, S., MÄKINEN, V., AND NAVARRO, G. 2004. First Huffman, then Burrows-Wheeler: an alphabet-independent FM-index. In *Proc. 11th International Symposium on String Processing and Information Retrieval (SPIRE)*, LNCS v. 3246 (2004), pp. 210–211. Short paper. Full version as Technical Report TR/DCC-2004-4, Department of Computer Science, University of Chile, July 2004.
- GROSSI, R., GUPTA, A., AND VITTER, J. 2003. High-order entropy-compressed text indexes. In *Proc. 14th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2003), pp. 841–850.
- GROSSI, R., GUPTA, A., AND VITTER, J. 2004. When indexing equals compression: Experiments with compressing suffix arrays and applications. In *Proc. 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2004), pp. 636–645.
- GROSSI, R. AND VITTER, J. 2000. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. In *Proc. 32nd ACM Symposium on Theory of*



- Computing (STOC)* (2000), pp. 397–406.
- GROSSI, R. AND VITTER, J. 2005. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. *SIAM Journal on Computing*. To appear.
- GUSFIELD, D. 1997. *Algorithms on Strings, Trees and Sequences: Computer Science and Computational Biology*. Cambridge University Press.
- HE, M., MUNRO, I., AND RAO, S. 2005. A categorization theorem on suffix arrays with applications to space efficient text indexes. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2005), pp. 23–32.
- HON, W.-K., LAM, T.-W., SADAKANE, K., AND SUNG, W.-K. 2003. Constructing compressed suffix arrays with large alphabets. In *Proc. 14th Annual International Symposium on Algorithms and Computation (ISAAC)* (2003), pp. 240–249.
- HON, W.-K., LAM, T.-W., SADAKANE, K., SUNG, W.-K., AND YU, S.-M. 2004. Compressed index for dynamic text. In *Proc. 14th IEEE Data Compression Conference (DCC)* (2004), pp. 102–111.
- HON, W.-K., SADAKANE, K., AND SUNG, W.-K. 2003. Breaking a time-and-space barrier in constructing full-text indices. In *Proc. 44th IEEE Symposium on Foundations of Computer Science (FOCS)* (2003), pp. 251–260.
- HUFFMAN, D. 1952. A method for the construction of minimum-redundancy codes. *Proceedings of the I.R.E.* 40, 9, 1090–1101.
- HUYNH, T., HON, W.-K., LAM, T.-W., AND SUNG, W.-K. 2004. Approximate string matching using compressed suffix arrays. In *Proc. 15th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 3109 (2004), pp. 434–444.
- IRVING, R. 1995. Suffix binary search trees. Technical Report TR-1995-7 (April), Computer Science Department, University of Glasgow, UK.
- JACOBSON, G. 1989. Space-efficient static trees and graphs. In *Proc. 30th IEEE Symposium on Foundations of Computer Science (FOCS)* (1989), pp. 549–554.
- KÄRKKÄINEN, J. 1995. Suffix cactus: a cross between suffix tree and suffix array. In *Proc. 6th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 937 (1995), pp. 191–204.
- KÄRKKÄINEN, J. 1999. *Repetition-Based Text Indexing*. Ph. D. thesis, Department of Computer Science, University of Helsinki, Finland.
- KÄRKKÄINEN, J. AND RAO, S. 2003. *Algorithms for Memory Hierarchies*, Chapter 7: Full-text indexes in external memory, pp. 149–170. LNCS v. 2625. Springer.
- KÄRKKÄINEN, J. AND SANDERS, P. 2003. Simple linear work suffix array construction. In *Proc. 30th International Colloquium on Automata, Languages and Programming (ICALP)*, LNCS v. 2719 (2003), pp. 943–955.
- KÄRKKÄINEN, J. AND SUTINEN, E. 1998. Lempel-Ziv index for  $q$ -grams. *Algorithmica* 21, 1, 137–154.
- KÄRKKÄINEN, J. AND UKKONEN, E. 1996a. Lempel-Ziv parsing and sublinear-size index structures for string matching. In *Proc. 3rd South American Workshop on String Processing (WSP)* (1996), pp. 141–155. Carleton University Press.
- KÄRKKÄINEN, J. AND UKKONEN, E. 1996b. Sparse suffix trees. In *Proc. 2nd Annual International Conference on Computing and Combinatorics (COCOON)*, LNCS v. 1090 (1996), pp. 219–230.
- KIM, D. AND PARK, H. 2005. A new compressed suffix tree supporting fast search and its construction algorithm using optimal working space. In *Proc. 16th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 3537 (2005), pp. 33–44.
- KIM, D., SIM, J., PARK, H., AND PARK, K. 2003. Linear-time construction of suffix arrays. In *Proc. 14th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 2676 (2003), pp. 186–199.
- KNUTH, D. 1973. *The Art of Computer Programming*, Volume 3: Sorting and Searching. Addison-Wesley.
- KO, P. AND ALURU, S. 2003. Space efficient linear time construction of suffix arrays. In *Proc. 14th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 2676

- (2003), pp. 200–210.
- KOSARAJU, R. AND MANZINI, G. 1999. Compression of low entropy strings with Lempel-Ziv algorithms. *SIAM Journal on Computing* 29, 3, 893–911.
- KURTZ, S. 1998. Reducing the space requirements of suffix trees. Report 98-03, Technische Fakultät, Universität Bielefeld, Germany.
- LAM, T.-W., SADAKANE, K., SUNG, W.-K., AND YIU, S.-M. 2002. A space and time efficient algorithm for constructing compressed suffix arrays. In *Proc. 8th Annual International Conference on Computing and Combinatorics (COCOON)* (2002), pp. 401–410.
- LARSSON, N. AND SADAKANE, K. 1999. Faster suffix sorting. Technical Report LU-CS-TR:99-214, Department of Computer Science, Lund University, Sweden.
- LEMPEL, A. AND ZIV, J. 1976. On the complexity of finite sequences. *IEEE Transactions on Information Theory* 22, 1, 75–81.
- MÄKINEN, V. 2000. Compact suffix array. In *Proc. 11th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 1848 (2000), pp. 305–319.
- MÄKINEN, V. 2003. Compact suffix array — a space-efficient full-text index. *Fundamenta Informaticae* 56, 1–2, 191–210.
- MÄKINEN, V. AND NAVARRO, G. 2004a. Compressed compact suffix arrays. In *Proc. 15th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 3109 (2004), pp. 420–433.
- MÄKINEN, V. AND NAVARRO, G. 2004b. New search algorithms and time/space tradeoffs for succinct suffix arrays. Technical Report C-2004-20 (April), University of Helsinki, Finland.
- MÄKINEN, V. AND NAVARRO, G. 2004c. Run-length FM-index. In *Proc. DIMACS Workshop: “The Burrows-Wheeler Transform: Ten Years Later”* (Aug. 2004), pp. 17–19.
- MÄKINEN, V. AND NAVARRO, G. 2005a. Succinct suffix arrays based on run-length encoding. In *Proc. 16th Annual Symposium on Combinatorial Pattern Matching (CPM)*, LNCS v. 3537 (2005), pp. 45–56.
- MÄKINEN, V. AND NAVARRO, G. 2005b. Succinct suffix arrays based on run-length encoding. Technical Report TR/DCC-2005-4 (March), Department of Computer Science, University of Chile, Chile.
- MÄKINEN, V., NAVARRO, G., AND SADAKANE, K. 2004. Advantages of backward searching — efficient secondary memory and distributed implementation of compressed suffix arrays. In *Proc. 15th Annual International Symposium on Algorithms and Computation (ISAAC)*, LNCS v. 3341 (2004), pp. 681–692.
- MANBER, U. AND MYERS, G. 1993. Suffix arrays: a new method for on-line string searches. *SIAM Journal on Computing* 22, 5, 935–948.
- MANZINI, G. 2001. An analysis of the Burrows-Wheeler transform. *Journal of the ACM* 48, 3, 407–430.
- MANZINI, G. AND FERRAGINA, P. 2004. Engineering a lightweight suffix array construction algorithm. *Algorithmica* 40, 1, 33–50.
- MCCREIGHT, E. 1976. A space-economical suffix tree construction algorithm. *Journal of the ACM* 23, 2, 262–272.
- MILTERSEN, P. 2005. Lower bounds on the size of selection and rank indexes. In *Proc. 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2005), pp. 11–12.
- MORRISON, D. 1968. PATRICIA — practical algorithm to retrieve information coded in alphanumeric. *Journal of the ACM* 15, 4, 514–534.
- MUNRO, I. 1996. Tables. In *Proc. 16th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, LNCS v. 1180 (1996), pp. 37–42.
- MUNRO, I. AND RAMAN, V. 1997. Succinct representation of balanced parentheses, static trees and planar graphs. In *Proc. 38th IEEE Symposium on Foundations of Computer Science (FOCS)* (1997), pp. 118–126.
- MUNRO, I., RAMAN, V., AND RAO, S. 2001. Space efficient suffix trees. *Journal of Algorithms* 39, 2, 205–222.
- NA, J. 2005. Linear-time construction of compressed suffix arrays using  $o(n \log n)$ -bit working space for large alphabets. In *Proc. 16th Annual Symposium on Combinatorial Pattern*

- Matching (CPM)*, LNCS v. 3537 (2005), pp. 57–67.
- NAVARRO, G. 2002. Indexing text using the ziv-lempel trie. In *Proc. 9th International Symposium on String Processing and Information Retrieval (SPIRE)*, LNCS v. 2476 (2002), pp. 325–336.
- NAVARRO, G. 2004. Indexing text using the Ziv-Lempel trie. *Journal of Discrete Algorithms* 2, 1, 87–114.
- PAGH, R. 1999. Low redundancy in dictionaries with  $O(1)$  worst case lookup time. In *Proc. 26th International Colloquium on Automata, Languages and Programming (ICALP)* (1999), pp. 595–604.
- RAMAN, R. 1996. Priority queues: small, monotone and trans-dichotomous. In *Proc. 4th European Symposium on Algorithms (ESA)*, LNCS v. 1136 (1996), pp. 121–137.
- RAMAN, R., RAMAN, V., AND RAO, S. 2002. Succinct indexable dictionaries with applications to encoding  $k$ -ary trees and multisets. In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2002), pp. 233–242.
- RAO, S. 2002. Time-space trade-offs for compressed suffix arrays. *Information Processing Letters* 82, 6, 307–311.
- SADAKANE, K. 2000. Compressed text databases with efficient query algorithms based on the compressed suffix array. In *Proc. 11th International Symposium on Algorithms and Computation (ISAAC)*, LNCS v. 1969 (2000), pp. 410–421.
- SADAKANE, K. 2002. Succinct representations of  $lcp$  information and improvements in the compressed suffix arrays. In *Proc. 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)* (2002), pp. 225–232.
- SADAKANE, K. 2003. New text indexing functionalities of the compressed suffix arrays. *Journal of Algorithms* 48, 2, 294–313.
- SCHÜRMAN, K. AND STOYE, J. 2005. An incomplex algorithm for fast suffix array construction. In *Proc. 7th Workshop on Algorithm Engineering and Experiments and 2nd Workshop on Analytic Algorithmics and Combinatorics (ALENEX/ANALCO)* (2005). SIAM Press. To appear.
- SEDGEWICK, R. AND FLAJOLET, P. 1996. *An Introduction to the Analysis of Algorithms*. Addison-Wesley.
- SIM, J., KIM, D., PARK, H., AND PARK, K. 2003. Linear-time search in suffix arrays. In *Proc. 14th Australasian Workshop on Combinatorial Algorithms (AWOCA)* (2003), pp. 139–146.
- UKKONEN, E. 1995. On-line construction of suffix trees. *Algorithmica* 14, 3, 249–260.
- WEINER, P. 1973. Linear pattern matching algorithm. In *Proc. 14th Annual IEEE Symposium on Switching and Automata Theory* (1973), pp. 1–11.
- WITTEN, I., MOFFAT, A., AND BELL, T. 1999. *Managing Gigabytes* (second ed.). Morgan Kaufmann Publishers.
- ZIV, J. AND LEMPEL, A. 1977. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory* 23, 3, 337–343.
- ZIV, J. AND LEMPEL, A. 1978. Compression of individual sequences via variable length coding. *IEEE Transactions on Information Theory* 24, 5, 530–536.