

Dynamic Data Structures for Document Collections and Graphs

J. Ian Munro
Cheriton School of Computer
Science
University of Waterloo
imunro@uwaterloo.ca.

Yakov Nekrich
Cheriton School of Computer
Science
University of Waterloo
ynekrich@uwaterloo.ca

Jeffrey Scott Vitter
Department of Electrical
Engineering & Computer
Science
University of Kansas
jsv@ku.edu.

ABSTRACT

In the dynamic indexing problem, we must maintain a changing collection of text documents so that we can efficiently support insertions, deletions, and pattern matching queries. We are especially interested in developing efficient data structures that store and query the documents in compressed form. All previous compressed solutions to this problem rely on answering rank and select queries on a dynamic sequence of symbols. Because of the lower bound in [Fredman and Saks, 1989], answering rank queries presents a bottleneck in compressed dynamic indexing. In this paper we show how this lower bound can be circumvented using our new framework. We demonstrate that the gap between static and dynamic variants of the indexing problem can be almost closed. Our method is based on a novel framework for adding dynamism to static compressed data structures. Our framework also applies more generally to dynamizing other problems. We show, for example, how our framework can be applied to develop compressed representations of dynamic graphs and binary relations.

Categories and Subject Descriptors

F2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems—*Sorting and searching*; H.3.1 [Information Storage and Retrieval]: Content Analysis and Indexing—*Indexing methods*

Keywords

Compressed Data Structures; Text Indexes; Graph Data Structures

1. INTRODUCTION

Motivated by the preponderance of massive data sets (so-called “big data”), it is becoming increasingly useful to store data in compressed form and moreover to manipulate and query the data while in compressed form. For that reason,

such compressed data structures have been developed in the context of text indexing, graph representations, XML indexes, labeled trees, and many other applications. In this paper we describe a general framework to convert known static compressed data structures into dynamic compressed data structures. We show how this framework can be used to obtain significant improvements for two important dynamic problems: maintaining a dynamic graph and storing a dynamic collection of documents. We expect that our general framework will find further applications.

In the indexing problem, we keep a text or a collection of texts in a data structure, so that, given a query pattern, we can list all occurrences of the query pattern in the texts. This problem is one of the most fundamental in the area of string algorithms. Data structures that use $O(n \log n)$ bits of space can answer pattern matching queries in optimal time $O(|P| + \text{occ})$, where $|P|$ denotes the length of the query pattern P and occ is the number of occurrences of P . Because of the large volumes of data stored in text data bases and document collections, we are especially interested in data structures that store the text or texts in compressed form and at the same time can answer pattern matching queries efficiently. Compressed indexing problem was extensively studied in the static scenario and during the last two decades significant progress has been achieved; we refer to a survey [30] for an overview of previous results in this area.

In the dynamic indexing problem, also known as the library management problem, we maintain a collection of documents (texts) in a data structure under insertions and deletions of texts. It is not difficult to keep a dynamic collection of texts in $O(n)$ words (i.e., $O(n \log n)$ bits) and support pattern matching queries at the same time. For instance, we can maintain suffixes of all texts in a suffix tree; when a new text is added or deleted, we add all suffixes of the new text to the suffix tree (respectively, remove all suffixes of the deleted text from the suffix tree). We refer the reader to the full version of this paper [29] for a more detailed description of the $O(n \log n)$ -bit solution. The problem of keeping a dynamic document collection in compressed form is however more challenging. Compressed data structures for the library management problem were considered in a number of papers [9, 25, 8, 25, 23, 26, 15, 24, 16, 20, 32, 31]. In spite of previous work, the query times of previously described dynamic data structures significantly exceed the query times of the best static indexes. In this paper we show that the gap between the static and the dynamic variants

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
PODS'15, May 31–June 4, 2015, Melbourne, Victoria, Australia.
Copyright © 2015 ACM 978-1-4503-2757-2/15/05 ...\$15.00.
<http://dx.doi.org/10.1145/2745754.2745778>.

of the compressed indexing problem can be closed or almost closed. Furthermore we show that our approach can be applied to the succinct representation of dynamic graphs and binary relations that supports basic adjacency and neighbor queries. Again our technique significantly reduces the gap between static and dynamic variants of this problem.

These problems arise often in database applications. For example, reporting or counting occurrences of a string in a dynamic collection of documents is an important operation in text databases and web browsers. Similar tasks also arise in data analytics. Suppose that we keep a search log and want to find out how many times URLs containing a certain substring were accessed. Finally the indexing problem is closely related to the problem of substring occurrence estimation [33]. The latter problem is used in solutions of the substring selectivity estimation problem [10, 21, 22]; we refer to [33] for a more extensive description. Compressed storage schemes for such problems help us save space and boost general performance because a larger portion of data can reside in the fast memory. Graph representation of data is gaining importance in the database community. For instance, the set of subject-predicate-object RDF triples can be represented as a graph or as two binary relations [12]. Our compressed representation applied to an RDF graph enables us to support basic reporting and counting queries on triples. An example of such a query is given x , to enumerate all the triples in which x occurs as a subject. Another example is, given x and p , to enumerate all triples in which x occurs as a subject and p occurs as a predicate.

Previous Results. Static Case.

We will denote by $|T|$ the number of symbols in a sequence T or in a collection of sequences; $T[i]$ denotes the i -th element in a sequence T and $T[i..j] = T[i]T[i+1] \dots T[j]$. Suffix trees and suffix arrays are two handbook data structures for the indexing problem. Suffix array keeps (references to) all suffixes $T[i..n]$ of a text T in lexicographic order. Using a suffix array, we can find the range of suffixes starting with a query string P in $t_{\text{range}} = O(|P| + \log n)$ time; once this range is found, we can locate each occurrence of P in T in $t_{\text{locate}} = O(1)$ time. A suffix tree is a compact trie that contains references to all suffixes $T[i..n]$ of a text T . Using a suffix trie, we can find the range of suffixes starting with a query string P in $t_{\text{range}} = O(|P|)$ time; once this range is found, we can locate every occurrence of P in T in $t_{\text{locate}} = O(1)$ time. A large number of compressed indexing data structures are described in the literature; we refer to [30] for a survey. These data structures follow the same two-step procedure for answering a query: first, the range of suffixes that start with P is found in $O(t_{\text{range}})$ time, then we locate each occurrence of P in T in $O(t_{\text{locate}})$ time. Thus we report all occ occurrences of P in $O(t_{\text{range}} + \text{occ} \cdot t_{\text{locate}})$ time. We can also extract any substring $T[i..i + \ell]$ of T in $O(t_{\text{extract}})$ time. Data structures supporting queries on a text T can be extended to answer queries on a collection \mathcal{C} of texts: it suffices to append a unique symbol $\$i$ at the end of every text T_i from \mathcal{C} and keep the concatenation of all T_i in the data structure.

We list the currently best and selected previous results for static text indexes with asymptotically optimal space usage in Table 1. All listed data structures can achieve different space-time trade-offs that depend on parameter s : an index typically needs about $nH_k + o(n \log \sigma) + O(n \log n/s)$ bits

and t_{locate} is proportional to s . Henceforth H_k denotes the k -th order empirical entropy and σ denotes the alphabet size¹. We assume that $k \leq \alpha \log_\sigma n - 1$ for a constant $0 < \alpha < 1$. H_k is the lower bound on the average space usage of any statistical compression method that encodes each symbol using the context of k previous symbols [27]. The currently fastest such index of Belazzougui and Navarro [6] reports all occurrences of P in $O(|P| + s \cdot \text{occ})$ time and extracts a substring of length ℓ in $O(s + \ell)$ time. Thus their query time depends only on the parameter s and the length of P . Some recently described indices [3, 6] achieve space usage $nH_k + o(nH_k) + o(n)$ or $nH_k + o(nH_k) + O(n)$ instead of $nH_k + o(n \log \sigma)$.

If we are interested in obtaining faster data structures and can use $\Theta(n \log \sigma)$ bits of space, then better trade-offs between space usage and time are possible [18, 19]. For the sake of space, we describe only one such result. The data structure of Grossi and Vitter [19] uses $O(n \log \sigma)$ bits and reports occurrences of a pattern in $O(|P|/\log_\sigma n + \log^\epsilon n + \text{occ} \log^\epsilon n)$ time; see Table 3. We remark that the fastest data structure in Table 1 needs $\Omega(n \log^{1-\epsilon} n)$ space to obtain the same time for t_{locate} as in [19]. If a data structure from Table 1 uses $O(n \log \sigma)$ space, then $t_{\text{locate}} = \Omega(\log_\sigma n)$.

Dynamic Document Collections.

In the dynamic indexing problem, we maintain a collection of documents (strings) under insertions and deletions. An insertion adds a new document to the collection, a deletion removes a document from the collection. For any query substring P , we must return all occurrences of P in all documents. When a query is answered, relative positions of occurrences are reported. To be precise, we must report all pairs (doc, off) , such that P occurs in a document doc at position off . We remark that relative positions of P (with respect to document boundaries) are reported. Hence an insertion or a deletion of a document does not change positions of P in other documents. Indexes for dynamic collections of strings were also studied extensively [9, 25, 8, 23, 26, 15, 24, 16, 20, 32, 31]. The fastest previously known result for the case of large alphabets is described in [31]. Their data structure, that builds on a long line of previous work, uses $nH_k + o(n \log \sigma) + O(n \log n/s) + O(\rho \log n)$ bits of space, where ρ is the number of documents; queries are answered in $O(|P| \log n / \log \log n + \text{occ} \cdot s \cdot \log n / \log \log n)$ time and updates are supported in $O(\log n + |T_u| \log n / \log \log n)$ amortized time, where T is the document inserted into or deleted from the collection. See Table 2 for some other previous results.

An important component of previous dynamic solutions is a data structure supporting rank and select queries: a sequence S over an alphabet $\Sigma = \{1, \dots, \sigma\}$ is kept in a data structure so that the i -th occurrence of a symbol $a \in \Sigma$ and the number of times a symbol a occurs in $S[1..i]$ for any $1 \leq i \leq n$ can be computed. Thus progress in dynamic indexing was closely related to progress in dynamic data structures for rank and select queries. In [31] the authors obtain a

¹Let S be an arbitrary string over an alphabet $\Sigma = \{1, \dots, \sigma\}$. A *context* $s_i \in \Sigma^k$ is an arbitrary string of length k . Let $n_{s_i, a}$ be the number of times the symbol a is preceded by a context s_i in S and $n_{s_i} = \sum_{a \in \Sigma} n_{s_i, a}$. Then $H_k = - \sum_{s_i \in \Sigma^k} \sum_{a \in \Sigma} n_{s_i, a} \log \frac{n_{s_i, a}}{n_{s_i}}$ is the k -th order empirical entropy of S .

Ref.	Space ($+O(n\frac{\log n}{s})$)	t_{range}	t_{locate}	t_{extract}	σ
[18]	$nH_k + o(n \log \sigma)$	$O(P \log \sigma + \log^4 n)$	$O(s \log \sigma)$	$O((s + \ell) \log \sigma)$	$\log^{\text{const}} n$
[34]	$nH_k + o(n \log \sigma)$	$O(P \log n)$	$O(s)$	$O(s + \ell)$	
[13]	$nH_k + o(n \log \sigma)$	$O(P \frac{\log \sigma}{\log \log n})$	$O(s \frac{\log \sigma}{\log \log n})$	$O((s + \ell) \frac{\log \sigma}{\log \log n})$	
[5]	$nH_k + o(n \log \sigma)$	$O(P \log \log \sigma)$	$O(s \log \log \sigma)$	$O((s + \ell) \log \log \sigma)$	
[3]	$nH_k + o(nH_k) + o(n)$	$O(P \frac{\log \sigma}{\log \log n})$	$O(s \frac{\log \sigma}{\log \log n})$	$O((s + \ell) \frac{\log \sigma}{\log \log n})$	
[3]	$nH_k + o(nH_k) + o(n)$	$O(P \log \log \sigma)$	$O(s \log \log \sigma)$	$O((s + \ell) \log \log \sigma)$	
[3]	$nH_k + o(nH_k) + o(n)$	$O(P)$	$O(s)$	$O(s + \ell)$	
[6]	$nH_k + o(nH_k) + O(n)$	$O(P)$	$O(s)$	$O(s + \ell)$	

Table 1: Asymptotically optimal space data structures for static indexing. Occurrences of a string P can be found in $O(t_{\text{range}} + t_{\text{locate}} \cdot \text{occ})$ time. A substring $T[i..i + \ell]$ of T can be extracted in $O(t_{\text{extract}})$ time. Results are valid for any $k \leq \alpha \log_{\sigma} n - 1$ and $0 < \alpha < 1$.

Ref.	Space ($+O(n\frac{\log n}{s}) + \rho \log n$)	t_{range}	t_{locate}	t_{extract}	Insert/ Delete	σ
[8]	$O(n)$	$O(P \log n)$	$O(\log^2 n)$	$O((\log n + \ell) \log n)$	$O(T_u \log n)$	const
[25]	$nH_k + o(n \log \sigma)$	$O(P \log n \log \sigma)$	$O(s \log n \log \sigma)$	$((s + \ell) \log n \log \sigma)$	$O(T_u \log n \log \sigma)$	
[31]	$nH_k + o(n \log \sigma)$	$O(P \log n)$	$O(s \log n)$	$O((s + \ell) \log n)$	$O(T_u \log n)$	
[31]	$nH_k + o(n \log \sigma)$	$O(P \frac{\log n}{\log \log n})$	$O(s \log n / \log \log n)$	$O((s + \ell) \frac{\log n}{\log \log n})$	$O(\log n + T_u \frac{\log n}{\log \log n})^{\text{A}}$	
Our	$nH_k + o(n \log \sigma)$	$O(P \log \log n)$	$O(s)$	$O(s + \ell)$	$O(T_u \log^{1+\epsilon} n)$	$\log^{\text{const}} n$
Our	$nH_k + o(n \log \sigma)$	$O(P \log \log n \log \log \sigma)$	$O(s \log \log \sigma)$	$O((s + \ell) \log \log \sigma)$	$O(T_u \log^{\epsilon} n) /$ $O(T_u (\log^{\epsilon} n + s))$	
Our	$nH_k + o(n \log \sigma)$	$O(P \log \log n)$	$O(s)$	$O(s + \ell)$	$O(T_u \log^{\epsilon} n)^{\text{R}} /$ $O(T_u (\log^{\epsilon} n + s))^{\text{R}}$	

Table 2: Asymptotically optimal space data structures for dynamic indexing. The same notation as in Table 1 is used. Randomized update procedures that achieve specified cost in expectation are marked with **R**. Amortized update costs are marked with **A**. T_u denotes the document that is inserted into (resp. deleted from) the data structure during an update operation. In previous papers on dynamic indexing only the cases of $s = \log n$ or $s = \log_{\sigma} n \log \log n$ was considered, but extension to an arbitrary value of s is straightforward.

dynamic data structure that supports rank and select in $O(\log n / \log \log n)$ time. By the lower bound of Fredman and Saks [14], this query time is optimal in the dynamic scenario. It was assumed that the solution of the library management problem described in [31] achieves query time that is close to optimal.

Our Results.

In this paper we show that the lower bound on dynamic rank-select problem can be circumvented and describe data structures that need significantly less time to answer queries. Our results close or almost close the gap between static and dynamic indexing. If the alphabet size $\sigma = \log^{O(1)} n$, we can obtain an $(nH_k + o(n \log \sigma) + O(n\frac{\log n}{s}))$ -bit data structure that answers queries in $O(|P| \log \log n + \text{occ} \cdot s)$ time; updates are supported in $O(|T_u| \log^{1+\epsilon} n)$ time, where T_u denotes the document that is inserted into or deleted from the index. Our second data structure supports updates in $O(|T_u| \log^{\epsilon} n)$ expected time and answers queries in $O(|P| \log \log n + \text{occ} \cdot s)$ time for an arbitrarily large alphabet². If the update procedure is deterministic, then queries are answered in $O((|P| \log \log n + \text{occ} \cdot s) \log \log \sigma)$ time and

updates are supported in $O(|T_u| \log^{\epsilon} n)$ worst-case time. See Table 2. If $O(n \log \sigma)$ bits of space are available, then our dynamic data structure matches the currently fastest static result of Grossi and Vitter [19]. We can report all occurrences of a pattern P in $O(|P| / \log_{\sigma} n + \log^{\epsilon} n + \text{occ} \cdot \log^{\epsilon} n)$ time. This is the first compressed dynamic data structure that achieves $t_{\text{range}} = o(|P|)$ if $\sigma = n^{o(1)}$. Compared to the fastest previous data structure that needs the same space, we achieve $O(\log n \log_{\sigma} n)$ factor improvement in query time. A variant of this data structure with deterministic update procedure answers queries in $O(|P| (\log \log n)^2 / \log_{\sigma} n + \log n + \text{occ} \cdot \log^{\epsilon} n)$ time. See Table 3.

Our data structures can also count occurrences of a pattern P in $O(t_{\text{count}})$ time. For previously described indexes $t_{\text{count}} = t_{\text{range}}$. In our case, $t_{\text{count}} = t_{\text{range}} + \log n / \log \log n$ or $t_{\text{count}} = (t_{\text{range}} + \log n / \log \log n) \log \log n$. Times needed to answer a counting query are listed in Table 4. However, if our data structures support counting queries, then update times grow slightly, as shown in Table 4.

All of the above mentioned results are obtained as corollaries of two general transformations. Using these transformations, that work for a very broad class of indexes, we can immediately turn almost any static data structure with good pre-processing time into an index for a dynamic collection of texts. The query time either remains the same or increases by a very small multiplicative factor. Our method can be applied to other problems where both compressed representation and dynamism are desirable.

²Dynamic indexes also need $O(\rho \log n)$ bits to navigate between documents, where ρ is the number of documents. Since $\rho \log n$ is usually negligible in comparison to n , we ignore this additive term, except for Tables 2 and 3, to simplify the description.

Binary Relations and Graphs.

One important area where our techniques can also be used is compact representation of directed graphs and binary relations. Let $R \subseteq L \times O$ be a binary relation between labels from a set L and objects from a set O . Barbay et al. [5] describe a compact representation of a static binary relation R (i.e., the set of object-label pairs) that consists of a sequence S_R and a bit sequence B_R . S_R contains the list of labels related to different objects and is ordered by object. That is, S_R lists all labels related to an object o_1 , then all labels related to an object o_2 , etc. The binary sequence B_R contains unary-encoded numbers of labels related to objects o_1, o_2, \dots . Barbay et al [5] showed how S_R and B_R can be used to support basic queries on binary relations, such as listing or counting all labels related to an object, listing or counting all objects related to a label, and telling whether a label and an object are related. Their method reduces queries on a binary relation R to rank, select, and access queries on S_R and B_R . Another data structure that stores a static binary relation and uses the same technique is described in [2]. Static compact data structures described in [5, 2] support queries in $O(\log \log \sigma_l)$ time per reported datum, where σ_l is the number of labels. For instance, we can report labels related to an object (resp. objects related to a label) in $O((k+1) \log \log \sigma_l)$ time, where k is the number of reported items; we can tell whether an object and a label are related in $O(\log \log \sigma_l)$ time. In [31], the authors describe a dynamization of this approach that relies on dynamic data structures answering rank and select queries on dynamic strings S_R and B_R . Again the lower bound on dynamic rank queries sets the limit on the efficiency of this approach. Since we need $\Omega(\log n / \log \log n)$ time to answer rank queries, the data structure of Navarro and Nekrich [31] needs $O(\log n / \log \log n)$ time per reported item, where n is the number of object-label pairs in R . Updates are supported in $O(\log n / \log \log n)$ amortized time and the space usage is $nH + \sigma_l \log \sigma_l + t \log t + O(n + \sigma_l \log \sigma_l)$ where n is the number of pairs, H is the zero-order entropy of the string S_R , σ_l is the number of labels and t is the number of objects. In [31] the authors also show that we can answer basic adjacency and neighbor queries on a directed graph by regarding a graph as a binary relation between nodes. Again reporting and counting out-going and in-going neighbors of a node can be performed in $O(\log n)$ time per delivered datum.

In this paper we show how our method for dynamizing compressed data structures can be applied to binary relations and graphs. Our data structure supports reporting labels related to an object or reporting objects related to a label in $O(\log \log n \cdot \log \log \sigma_l)$ time per reported datum. We support counting queries in $O(\log n)$ time and updates in $O(\log^\epsilon n)$ worst-case time. The same query times are also achieved for the dynamic graph representation. The space usage of our data structures is dominated by nH where n is the number of pairs in a binary relation or the number of edges in a graph and H is the zero-order entropy of the string S_R defined above. Thus the space usage of our data structure matches that of [31] up to lower-order factors. At the same time we show that reporting queries in a dynamic graph can be supported without dynamic rank and select queries.

Overview.

The main idea of our approach can be described as follows. The input data is distributed among several data structures. We maintain a fraction of the data in an uncompressed data structure that supports both insertions and deletions. We bound the number of elements stored in uncompressed form so that the total space usage of the uncompressed data structure is affordable. Remaining data is kept in several compressed data structures that do not support updates. New elements (respectively new documents) are always inserted into the uncompressed data structure. Deletions from the static data structures are implemented by the lazy deletions mechanism: when a deletion takes place, then the deleted element (respectively the document) is marked as deleted. We keep positions of marked elements in a data structure, so that all elements in a query range that are not marked as deleted can be reported in $O(1)$ time per element. When a static structure contains too much obsolete data (because a certain fraction of its size is marked as deleted), then this data structure is purged: we create a new instance of this data structure that does not contain deleted elements. If the uncompressed data structure becomes too big, we move its content into a (new) compressed data structure. Organization of compressed data structures is inspired by the logarithmic method, introduced by Bentley and Saxe [7]: the size of compressed data structures increases geometrically. We show that re-building procedures can be scheduled in such way that only a small fraction of data is kept in uncompressed form at any given time. Since the bulk of data is kept in static data structures, our approach can be viewed as a general framework that transforms static compressed data structures into dynamic ones.

In Section 2 we describe Transformation 1; Transformation 1, based on the approach outlined above, can be used to turn a static indexing data structure into a data structure for dynamic collection of documents with amortized update cost. The query costs of the obtained dynamic data structure are the same as in the underlying static data structure. In Section 3 we describe Transformation 2 that turns a static indexing data structure into a dynamic data structure with worst-case update costs. We use more sophisticated division into sub-collections and slightly different re-building procedures in Transformation 2. In Section 4 we describe how to obtain new solutions of the dynamic indexing problem using our static-to-dynamic transformations. Finally Section 5 contains our data structures for dynamic graphs and binary relations.

2. DYNAMIC DOCUMENT COLLECTIONS

In this section we show how a static compressed index \mathcal{I}_s can be transformed into a dynamic index \mathcal{I}_d . \mathcal{C} will denote a collection of texts T_1, \dots, T_p . We say that an index \mathcal{I}_s is $(u(n), w(n))$ -constructible if there is an algorithm that uses $O(n \cdot w(n))$ additional workspace and constructs \mathcal{I}_s in $O(n \cdot u(n))$ time. Henceforth we make the following important assumptions about the static index \mathcal{I}_s . \mathcal{I}_s needs at most $|S|\phi(S)$ bits of space for any symbol sequence S and the function $\phi(\cdot)$ is monotonous: if any sequence S is a concatenation of S_1 and S_2 , then $|S|\phi(S) \geq |S_1|\phi(S_1) + |S_2|\phi(S_2)$. We also assume that \mathcal{I}_s reports occurrences of a substring in \mathcal{C} using the two-step method described in the introduction: first we identify the range $[a, b]$ in the suffix array, such that all suffixes that start with P are in $[a, b]$; then we find the

Ref.	t_{range}	t_{locate}	t_{extract}	Update	σ
[19]	$O(P /\log_{\sigma} n + \log^{\varepsilon} n)$	$O(\log^{\varepsilon} n)$	$O(\ell/\log_{\sigma} n)$	static	const
[8]	$O(P \log n)$	$O(\log^2 n)$	$O((\log n + \ell) \log n)$	$O(T_u \log n)$	
[31]	$O(P \log n)$	$O(\log n \log_{\sigma} n)$	$O((\log_{\sigma} n + \ell) \log n)$	$O(T_u \log n)$	
Our	$O(P /\log_{\sigma} n + \log^{\varepsilon} n)$	$O(\log^{\varepsilon} n)$	$O(\ell/\log_{\sigma} n)$	$O(T_u \log^{\varepsilon} n)^{\mathbf{R}}$	
Our	$O(P (\log \log n)^2/\log_{\sigma} n + \log n)$	$O(\log^{\varepsilon} n)$	$O(\ell/\log_{\sigma} n)$	$O(T_u \log^{\varepsilon} n)$	

Table 3: $O(n \log \sigma)$ -bit indexes. Dynamic data structures need additional $\rho \log n$ bits. Randomized update costs are marked with **R**.

Space	Counting	Updates	σ
$nH_k + o(n \log \sigma)$	$O(P \log \log n + \log n)$	$O(T_u \log n)$	$\log^{\text{const}} n$
$nH_k + o(n \log \sigma)$	$O(P \log \log \sigma \log \log n + \log n)$	$O(T_u \log n)$	
$nH_k + o(n \log \sigma)$	$O(P \log \log n + \log n)$	$O(T_u \log n)^{\mathbf{R}}$	
$O(n \log \sigma)$	$O(P /\log_{\sigma} n + \log n/\log \log n)$	$O(T_u \log n)^{\mathbf{R}}$	
$O(n \log \sigma)$	$O(P (\log \log n)^2/\log_{\sigma} n + \log n)$	$O(T_u \log n)$	

Table 4: Costs of counting queries for our data structures. Randomized update costs are marked with **R**. The first three rows correspond to the last three rows in Table 2, the last two rows correspond to the last two rows in Table 3.

positions of suffixes from $[a, b]$ in the document(s). These operations will be called range-finding and locating. Moreover the rank of any suffix $T_i[l..]$ in the suffix array can be found in time $O(t_{\text{SA}})$. The class of indexes that satisfy these conditions includes all indexes that are based on compressed suffix arrays or the Burrows-Wheeler transform. Thus the best currently known static indexes can be used in Transformation 1 and the following transformations described in this paper.

Our result can be stated as follows.

TRANSFORMATION 1. *Suppose that there exists a static $(u(n), w(n))$ -constructible index \mathcal{I}_s that uses $|S|\phi(S)$ space for any document collection S . Then there exists a dynamic index \mathcal{I}_d that uses $|S|\phi(S) + O(|S|(\frac{\log \sigma}{\tau} + w(n) + \frac{\log \tau}{\tau}))$ space for a parameter $\tau = O(\log n / \log \log n)$; \mathcal{I}_d supports insertions and deletions of documents in $O(u(n) \log^{\varepsilon} n)$ time per symbol and $O(u(n) \cdot \tau + t_{\text{SA}} + \log^{\varepsilon} n)$ time per symbol respectively. Update times are amortized. The asymptotic costs of range-finding, extracting, and locating are the same in \mathcal{I}_s and \mathcal{I}_d .*

We start by showing how to turn a static index into a semi-dynamic deletion-only index using $O((n/\tau) \log \tau)$ additional bits. Then we will show how to turn a semi-dynamic index into a fully-dynamic one.

Supporting Document Deletions.

We keep a bit array B whose entries correspond to positions in the suffix array SA of \mathcal{C} . $B[j] = 0$ if $SA[j]$ is a suffix of some text T_f , such that T_f was already deleted from \mathcal{C} and $B[j] = 1$ otherwise. We keep a data structure V that supports the following operations on B : $\text{zero}(j)$ sets the j -th bit in B to 0; $\text{report}(j_1, j_2)$ reports all 1-bits in $B[j_1..j_2]$. V is implemented using Lemma 3, so that $\text{zero}(i)$ is supported in $O(\log^{\varepsilon} n)$ time and $\text{report}(j_1, j_2)$ is answered in $O(k)$ time, where k is the number of output bit positions. If B contains at most n/τ zeros, then B and V need only $O((n \log \tau)/\tau)$ bits. Lemma 3 is proved in Section A.1.

When a document T_f is deleted, we identify the positions of T_f 's suffixes in SA and set the corresponding bits in B




Figure 1: Sub-collections \mathcal{C}_i for dynamizing a deletion-only index. A data structure for \mathcal{C}_0 is fully-dynamic and stores documents in uncompressed form.

to 0. When the number of symbols in deleted documents equals (n/τ) , we re-build the index for \mathcal{C} without deleted documents in $O(n \cdot u(n))$ time. The total amortized cost of deleting a document is $O(u(n)\tau + t_{\text{SA}} + \log^{\varepsilon} n)$ per symbol. To report occurrences of some string P in \mathcal{C} , we identify the range $[s..e]$ such that all suffixes in $SA[s..e]$ start with P in $O(t_{\text{range}})$ time. Using V , we enumerate all j , such that $s \leq j \leq e$ and $B[j] = 1$. For every such j , we compute $SA[j]$ in $O(t_{\text{locate}})$ time.

Fully-Dynamic Index.

We split \mathcal{C} into a constant number of sub-collections $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_r$ such that $|\mathcal{C}_i| \leq \max_i$ for all i . The maximum size of the i -th sub-collection, \max_i , increases geometrically: $\max_0 = 2n/\log^2 n$ and $\max_i = 2(n/\log^2 n) \log^{\varepsilon-i} n$ for a constant $\varepsilon > 0$; see Fig. 1. There is no lower bound on the number of symbols in a sub-collection \mathcal{C}_i ; for instance, any \mathcal{C}_i can be empty. Our main idea is to store \mathcal{C}_0 in uncompressed form and \mathcal{C}_i for $i \geq 1$ in semi-static deletion-only data structures. Insertions into \mathcal{C}_i for $i \geq 1$ are supported by re-building the semi-static index of \mathcal{C}_i . We also re-build all sub-collections when the total number of elements is increased by a constant factor (global re-build).

We store the document collection \mathcal{C}_0 in uncompressed form. Suffixes of all documents in \mathcal{C}_0 are kept in an (uncompressed) suffix tree \mathcal{D}_0 . We can insert a new text T into \mathcal{D}_0 or delete T from \mathcal{D}_0 in $O(|T|)$ time. Using \mathcal{D}_0 , all occurrences of a pattern P in \mathcal{C}_0 can be reported in $O(|P| + \text{occ})$ time. Since $|\mathcal{C}_0| \leq 2n/\log^2 n$, we need $O(n/\log n)$ bits to store \mathcal{C}_0 . For completeness we describe the data structure for \mathcal{C}_0 in the full version of this paper [29].

Every \mathcal{C}_i for $i \geq 1$ is kept in a semi-dynamic data structure described in the first part of this section. Let $\text{size}(i)$ denote the total length of all undeleted texts in \mathcal{C}_i . When a new document T must be inserted, we find the first (smallest) collection \mathcal{C}_j such that $\sum_{i=0}^j \text{size}(i) + |T| \leq \max_j$ where $\max_j = 2(n/\log^2 n) \log^{\varepsilon^j} n$. That is, we find the first sub-collection \mathcal{C}_j that can accommodate the new text T and all preceding subcollections without exceeding the size limit. If $j = 0$, we insert the new text into \mathcal{C}_0 . Otherwise, if $j \geq 1$, we discard the old indexes for all \mathcal{C}_i where $0 \leq i \leq j$, set $\mathcal{C}_j = (\cup_{i=0}^j \mathcal{C}_i) \cup T$ and construct a new semi-static index for \mathcal{C}_j . If $\sum_{i=0}^j \text{size}(i) + |T| > \max_j$ for all j , we start a global re-build procedure: all undeleted texts from old sub-collections are moved to the new sub-collection \mathcal{C}_r and parameters \max_i are re-calculated; new sub-collections \mathcal{C}_i for $0 \leq i < r$ are initially empty after the global re-build.

We start a global re-build procedure when the total number of elements is at least doubled. Hence, the amortized cost of a global re-build is $O(u(n))$. The amortized cost of re-building sub-collections can be analyzed as follows. When a sub-collection \mathcal{C}_j is re-built, we insert all symbols from subcollections \mathcal{C}_i , $0 \leq i < j$ and the new text T into \mathcal{C}_j . Our insertion procedure guarantees that $\sum_{i=1}^{j-1} \text{size}(i) + |T| > \max_{j-1}$. We need $O(\max_j \cdot u(n))$ time to construct a new index for \mathcal{C}_j . The cost of re-building \mathcal{C}_j can be distributed among the new text symbols inserted into \mathcal{C}_j . Since $\max_{j-1} = \max_j / \log^\varepsilon n$, the amortized cost of inserting a new symbol into \mathcal{C}_j is $O(u(n) \cdot \log^\varepsilon n)$. Every text is moved at most once to any subcollection \mathcal{C}_j for any j such that $1 \leq j \leq \lceil 2/\varepsilon \rceil$. Hence the total amortized cost of an insertion is $O((1/\varepsilon)u(n) \cdot \log^\varepsilon n)$ per symbol.

A query is answered by querying all non-empty sub-collections \mathcal{C}_i for $i = 0, 1, \dots, r$. Since $r = O(1)$, query times are the same as in the underlying static index. Splitting a collection into sub-collection does not increase the space usage because the function $\phi(\cdot)$ is monotonous. We need $O((n/\tau) \log \tau)$ bits to keep data structures V for all \mathcal{C}_i . Another $O(nw(n))$ bits are needed for global and local re-builds. Finally we need $O((n/\tau) \log \sigma)$ bits to store the symbols from deleted documents. Since there are no more than $O(n/\tau)$ deleted symbols, we use $O((n/\tau) \log \sigma) + o(n \log \sigma)$ additional bits to store them; a more detailed analysis is given in the full version [29]. Hence, the total space overhead of our dynamic index is $O(n(w(n) + (\log \sigma + \log \tau)/\tau))$.

3. WORST-CASE UPDATES

In this section we will prove the following result.

TRANSFORMATION 2. *Suppose that there exists a static $(u(n), w(n))$ -constructible index \mathcal{I}_s that uses $|S|\phi(S)$ space for any document collection S . Then there exists a dynamic index \mathcal{I}_d that uses $|S|\phi(S) + O(|S| \frac{\log \sigma + \log \tau + w(n)}{\log \sigma})$ space for any parameter $\tau = O(\log n / \log \log n)$; \mathcal{I}_d supports insertions and deletions of documents in $O(u(n) \log^\varepsilon n)$ time per symbol and $O(u(n) \cdot (\log^\varepsilon n + \tau \log \tau) + t_{\text{SA}})$ time per symbol*

respectively. The asymptotic costs of range-finding increases by $O(\tau)$; the costs of extracting and locating are the same in \mathcal{I}_s and \mathcal{I}_d .

We use the index of Transformation 1 as the starting point. First we give an overview of our data structure and show how queries can be answered. Then we describe the procedures for text insertions and deletions.

Overview.

The main idea of supporting updates in worst-case is to maintain several copies of the same sub-collection. An old copy of \mathcal{C}_j is locked while a new updated version of \mathcal{C}_{j+1} that includes \mathcal{C}_j is created in the background. When a new version of \mathcal{C}_{j+1} is finished, we discard an old locked sub-collection. When a new document \mathcal{T} must be inserted, we insert it into \mathcal{C}_0 if $|\mathcal{C}_0| + |T| \leq \max_0$. Otherwise we look for the smallest $j \geq 0$, such that \mathcal{C}_{j+1} can accommodate both \mathcal{C}_j and T ; then we move both T and all documents from \mathcal{C}_j into \mathcal{C}_{j+1} ³. If the new document T is large, $|T| \geq \max_j / 2$, we can afford to re-build \mathcal{C}_{j+1} immediately after the insertion of T . If the size of T is smaller than $\max_j / 2$, re-building of \mathcal{C}_{j+1} is postponed. For every following update, we spend $O(\log^\varepsilon n \cdot u(n))$ time per symbol on creating the new version of \mathcal{C}_{j+1} . The old versions of \mathcal{C}_j , \mathcal{C}_{j+1} are retained until the new version is completed. If the number of symbols that are marked as deleted in \mathcal{C}_j exceeds $\max_j / 2$, we employ the same procedure for moving \mathcal{C}_j to \mathcal{C}_{j+1} : \mathcal{C}_j is locked and we start the process of constructing a new version \mathcal{C}_{j+1} that contains all undeleted documents from \mathcal{C}_j .

The disadvantage of delayed re-building is that we must keep two copies of every document in $\mathcal{C}_j \cup \mathcal{C}_{j+1}$ until new \mathcal{C}_{j+1} is completed. In order to reduce the space usage, we keep only a fraction of all documents in sub-collections \mathcal{C}_i . All \mathcal{C}_i for $0 \leq i \leq r$ will contain $O(n/\tau)$ symbols, where τ is the parameter determining the trade-off between space overhead and query time. The remaining documents are kept in top sub-collections $\mathcal{T}_1, \dots, \mathcal{T}_g$ where $g \leq 2\tau$. Top sub-collections are constructed using the same delayed approach. But once \mathcal{T}_i is finished, no new documents are inserted into \mathcal{T}_i . We may have to re-build a top collection or merge it with another \mathcal{T}_j when the fraction of deleted symbols exceeds a threshold value $1/\tau$. We employ the same rebuilding-in-the-background approach. However, we will show that the background procedures for maintaining \mathcal{T}_i can be scheduled in such a way that only one \mathcal{T}_j is re-built at any given moment. Hence, the total space overhead due to re-building and storage of deleted elements is bounded by an additive term $O(n(\log \sigma + w(n))/\tau)$.

Data Structures.

We split a document collection \mathcal{C} into subcollections $\mathcal{C}_0, \mathcal{C}_1, \dots, \mathcal{C}_r$, $\mathcal{L}_1, \dots, \mathcal{L}_r$ and top subcollections $\mathcal{T}_1, \dots, \mathcal{T}_g$ where $g = O(\tau)$. We will also use auxiliary collections $\mathcal{N}_1, \dots, \mathcal{N}_{r+1}$ and temporary collections $\text{Temp}_1, \dots, \text{Temp}_r$. Temp_i are also used to answer queries but each non-empty

³Please note the difference between Transformations 1 and 2. In Transformation 1 we look for the sub-collection \mathcal{C}_j that can accommodate the new document and all smaller sub-collections $\mathcal{C}_0, \dots, \mathcal{C}_{j-1}$. In Transformation 2 we look for the sub-collection \mathcal{C}_{j+1} that can accommodate that can accommodate the new document and the preceding sub-collection \mathcal{C}_j . We made this change in order to avoid some technical complications caused by delayed re-building.

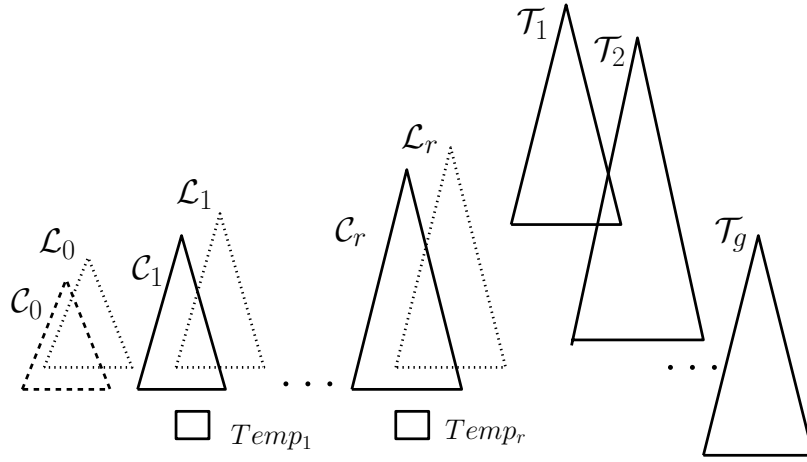


Figure 2: Dynamization with worst-case update guarantees. Only main sub-collections used for answering queries are shown. \mathcal{L}'_r and auxiliary collections \mathcal{N}_i are not shown.

$Temp_i$ contains exactly one document; $Temp_i$ are used as temporary storage for new document that are not yet inserted into “big” collections. The sizes of sub-collections can be defined as a function of parameter n_f such that $n_f = \Theta(n)$; the value of n_f changes when n becomes too large or too small. Let $\max_i = 2(n_f/\log^2 n) \log^{3\epsilon} n$. We maintain the invariant $|\mathcal{C}_i| \leq \max_i$ for all i , $0 \leq i \leq r$, but r is chosen in such way that $n_f/\log^{2-r\epsilon} n_f = n_f/\tau$. Every \mathcal{T}_i contains $\Omega(n_f/\tau)$ symbols. If \mathcal{T}_i contains more than one text, then its size is at most $4n_f/\tau$; otherwise \mathcal{T}_i can be arbitrarily large. When a collection \mathcal{C}_j is merged with \mathcal{C}_{j+1} , the process of re-building \mathcal{C}_j can be distributed among a number of future updates (insertions and deletions of documents). During this time \mathcal{C}_j is locked: we set $\mathcal{L}_j = \mathcal{C}_j$ and initialize a new empty sub-collection \mathcal{C}_j . When a new subcollection $\mathcal{N}_{j+1} = \mathcal{C}_{j+1} \cup \mathcal{C}_j$ is completed, we set $\mathcal{C}_{j+1} = \mathcal{N}_{j+1}$ and discard old \mathcal{C}_{j+1} and \mathcal{L}_j . A query is answered by querying all non-empty \mathcal{C}_i , \mathcal{L}_i , $Temp_i$, and \mathcal{T}_i . Therefore the cost of answering a range-finding query grows by $O(\tau)$. The costs of locating and extracting are the same as in the static index. We show main sub-collections used by our method on Figure 2.

Insertions.

When a document T is inserted, we consider all j , such that $0 \leq j \leq r$ and the data structure \mathcal{L}_j is not empty. For every such j , we spend $O(|T| \log^\epsilon n \cdot u(n))$ units of time on constructing \mathcal{N}_{j+1} . If \mathcal{N}_{j+1} for some $0 \leq j \leq r-1$ is completed, we set $\mathcal{C}_{j+1} = \mathcal{N}_{j+1}$ and $\mathcal{N}_{j+1} = Temp_{j+1} = \emptyset$; if \mathcal{N}_{r+1} is completed, we set $\mathcal{T}_{g+1} = \mathcal{N}_{r+1}$, increment the number of top collections g , and set $\mathcal{N}_{r+1} = Temp_{r+1} = \emptyset$. Then we look for a sub-collection that can accommodate the new document T . If $|T| \geq n/\tau$, we create the index for a new sub-collection \mathcal{T}_i that contains a single document T . If $|T| < n/\tau$, we look for the smallest index j , such that $|\mathcal{C}_{j+1}| + |\mathcal{C}_j| + |T| \leq \max_{j+1}$. That is, \mathcal{C}_{j+1} can accommodate both the preceding sub-collection \mathcal{C}_j and T . If $|T| \geq \max_j/2$, we set $\mathcal{C}_{j+1} = \mathcal{C}_j \cup \mathcal{C}_{j+1} \cup T$ and create an index for the new \mathcal{C}_{j+1} in $O(|\mathcal{C}_{j+1}| \cdot u(n)) = O(|T| \log^\epsilon n \cdot u(n))$ time. If $|T| < \max_j/2$, the collection \mathcal{C}_j is locked. We set $\mathcal{L}_j = \mathcal{C}_j$, $\mathcal{C}_j = \emptyset$ and initiate the process of creating $\mathcal{N}_{j+1} = \mathcal{C}_j \cup \mathcal{C}_{j+1} \cup T$. The cost of creating the new index for \mathcal{N}_{j+1} will

be distributed among the next \max_j update operations. We also create a temporary static index $Temp_{j+1}$ for the text T in $O(|T|u(n))$ time. This procedure is illustrated on Fig. 3. If the index j is not found and $|\mathcal{C}_i| + |\mathcal{C}_{i+1}| + |T| > \max_{i+1}$ for all i , $0 \leq i < r$, we lock \mathcal{C}_r (that is, set $\mathcal{L}_r = \mathcal{C}_r$ and $\mathcal{C}_r = \emptyset$) and initiate the process of constructing $\mathcal{N}_{r+1} = \mathcal{L}_r \cup T$. We also create a temporary index $Temp_{r+1}$ for the document T in $O(|T|u(n))$ time.

Deletions.

Indexes for sub-collections \mathcal{C}_i , $1 \leq i \leq r$, and \mathcal{T}_j , $1 \leq j \leq g$, support lazy deletions in the same way as in Section 2: when a document is deleted from a sub-collection, we simply mark the positions of suffixes in the suffix array as deleted and set the corresponding bits in the bit vector B to 0. Augmenting an index so that lazy deletions are supported is done in exactly the same way as in Section 2.

We will need one additional sub-collection \mathcal{L}'_r to support deletions. If a sub-collection \mathcal{C}_j for $1 \leq j \leq r-1$ contains $\max_j/2$ deleted elements, we start the process of re-building \mathcal{C}_j and merging it with \mathcal{C}_{j+1} . This procedure is the same as in the case of insertions. We lock \mathcal{C}_j by setting $\mathcal{L}_j = \mathcal{C}_j$ and $\mathcal{C}_j = \emptyset$. The data structure $\mathcal{N}_{j+1} = \mathcal{C}_{j+1} \cup \mathcal{L}_j$ will be re-built during the following $\max_j/2$ updates. If a sub-collection \mathcal{C}_r contains $\max_r/2$ deleted symbols, we set $\mathcal{L}'_r = \mathcal{C}_r$ and $\mathcal{C}_r = \emptyset$. The sub-collection \mathcal{L}'_r will be merged with the next sub-collection \mathcal{T}_i to be re-built.

If a collection \mathcal{T}_i contains a single document and this document is deleted, then \mathcal{T}_i is discarded. We also bound the number of deleted symbols in any \mathcal{T}_i by n_f/τ . This is achieved by running the following background process. After each series of $n_f/(2\tau \log \tau)$ symbol deletions, we identify \mathcal{T}_j that contains the largest number of deleted symbols. During the next $n_f/(2\tau \log \tau)$ symbol deletions we build the new index for \mathcal{T}_j without the deleted symbols. At the same time we remove the deleted symbols from \mathcal{L}'_r if \mathcal{L}'_r exists. If \mathcal{L}'_r exists and contains at least $n_f/2\tau$ undeleted symbols, we create an index for a new sub-collection \mathcal{T}'_{g+1} and increment the number g of top collections. If \mathcal{L}'_r exists, but contains less than $n_f/2$ undeleted symbols, we merge \mathcal{L}'_r with the largest \mathcal{T}_j that contains more than one document and split the result if necessary: if the number of undeleted symbols in $\mathcal{L}'_r \cup \mathcal{T}_j$ does




Figure 3: Suppose that C_{j+1} is the first sub-collection that can accommodate both C_j and a new document T_n . If C_j must be rebuilt in the background, we “rename” C_j to L_j and initialize another (initially empty) C_j . New document T_n is put into a separate collection $Temp_{j+1}$ (a). A background process creates a new collection N_{j+1} that contains all documents from L_j , C_{j+1} and $Temp_{j+1}$ (b). When N_{j+1} is finished, we discard C_{j+1} , L_j and $Temp_{j+1}$, and set $C_{j+1} = N_{j+1}$ (c). Our procedure guarantees that N_{j+1} is completed before the new sub-collection C_j must be re-built again.

not exceed $2n_f/\tau$, we construct an index for $T_j \cup L'_r$ without deleted symbols; otherwise, we split $T_j \cup L'_r$ into two parts T_j^1, T_j^2 and create indexes for the new sub-collections. Our method guarantees us that the number of deleted elements in any collection \mathcal{T}_i does not exceed $O(n_f/\tau)$ as follows from a Theorem of Dietz and Sleator [11].

LEMMA 1 ([11], THEOREM 5). *Suppose that x_1, \dots, x_g are variables that are initially zero. Suppose that the following two steps are iterated: (i) we add a non-negative real value a_i to each x_i such that $\sum a_i = 1$ (ii) set the largest x_i to 0. Then at any time $x_i \leq 1 + h_{g-1}$ for all $i, 1 \leq i \leq g$, where h_i denotes the i -th harmonic number.*

Let m_i be the number of deleted elements in the i -th top collection \mathcal{T}_i and $\delta = n_f/(2\tau \log \tau)$. We define $x_i = m_i/\delta$. We consider the working of our algorithm during the period when the value of n_f is fixed. Hence, δ is also fixed and the number of variables x_i is $O(\tau)$ (some x_i can correspond to empty collections). Every iteration of the background process sets the largest x_i to 0. During each iteration $\sum x_i$ increases by 1. Hence, the values of x_i can be bounded from above by the result of Lemma 1: $x_i \leq 1 + h_{2\tau}$ for all i at all times. Hence $m_i = O((n_f/2\tau \log \tau) \log \tau) = O(n_f/\tau)$ for all i because $h_i = O(\log i)$. Thus the fraction of deleted symbols in each \mathcal{T}_i is $O(1/\tau)$.

It is easy to show that the sub-collections that we use are sufficient for our algorithm. When a sub-collection \mathcal{L}_j is initialized, C_j is empty. The situation when C_j cannot accommodate a new document T_n and a preceding sub-collection C_{j-1} can happen only after $\max_j - \sum_{t=1}^{j-1} \max_t$ new symbol insertions. Since we spend $O(\log^\varepsilon n \cdot u(n))$ time for constructing N_{j+1} with each new symbol insertion, we can choose constants in such a way that construction of N_{j+1} is finished (and \mathcal{L}_j is discarded) after $\max_j/2 < \max_j - \sum_{t=1}^{j-1} \max_t$ symbol insertions. The situation when C_j contains $\max_j/2$ deleted symbols can happen after at least \max_j new symbol updates ($\max_j/2$ insertions and $\max_j/2$ deletions). Hence, the collection \mathcal{L}_j is discarded before C_j has to be locked again. In our description of update procedures we assumed that the parameter n_f is fixed. We can maintain the invariant $n_f = \Theta(n)$ using standard methods; for completeness we will provide a description in the full version [29].

The space overhead caused by storing copies of deleted elements is bounded by $O(n/\tau)$: all C_i contain $O(n/\tau)$ symbols and at most every second symbol in each C_i is from a deleted document; the fraction of deleted symbols in each \mathcal{T}_i does not exceed $O(1/\tau)$. By the same argument, at any moment of time at most $O(n/\tau)$ symbols are in sub-collections that are re-built. Hence re-building procedures running in the background need $O(nw(n)/\tau)$ bits of space. Since each \mathcal{T}_i contains at most $O(|\mathcal{T}_i|/\tau)$ deleted symbols, we can store the data structure V , which enables us to identify undeleted elements in any range of the suffix array and is implemented as described in Lemma 3, using $O(|\mathcal{T}_i| \log \tau/\tau)$ bits. Data structures V for all \mathcal{T}_i need $O(n \log \tau/\tau)$ bits. Hence, the total space overhead of \mathcal{I}_d compared to \mathcal{I}_s is $O(n \frac{w(n) + \log \tau + \log \sigma}{\tau})$ bits.

Counting Occurrences.

Our dynamic indexes can be easily extended so that pattern counting queries are supported.

THEOREM 1. *We can augment the indexes \mathcal{I}_d of Transformations 1- 2 with $O((n \log \tau)/\tau)$ additional bits so that all occurrences of a pattern can be counted in $O(t_{\text{count}})$ time, where $t_{\text{count}} = (t_{\text{range}} + \log n/\log \log n)(r + \tau)$ and τ is defined as in the proofs of respective Transformations. If counting is supported, update times are increased by $O(\log n/\log \log n)$ additive term per symbol.*

PROOF. Every semi-dynamic index for a sub-collection C_i (respectively \mathcal{T}_i) already keeps a vector B that enables us to identify the suffixes of already deleted documents in the suffix array. We also store each B in a data structure of Navarro and Sadakane [32] that supports rank queries in $O(\log n/\log \log n)$ time and updates in $O(\log n/\log \log n)$ time. If B contains $O(|B|/\tau)$ zero values, then the structure of [32] needs $O((|B|/\tau) \log \tau)$ bits. Using this data structure, we can count the number of 1's in any portion $B[a..b]$ of B in the same time. To answer a counting query, we first answer a range-finding query in every sub-collection. For every non-empty range that we found, we count the number of 1's in that range. Finally, we sum the answers for all sub-collections. Since a range-finding query returns the range of all suffixes that start with a query pattern and each

1 in V corresponds to a suffix of an undeleted document, our procedure is correct. \square

4. DYNAMIC INDEXES

To obtain our results on dynamic document collections, we only need to plug some currently known static indexes into Transformations 1 and 2. We prove the statements about constructibility of static indexes in the full version of this paper [29].

The static index of Belazzougui and Navarro, described in [6], is $(\log^\varepsilon n, \log \sigma)$ -constructible. Their index achieves $t_{\text{range}} = O(|P|)$, $t_{\text{extract}} = O(s + \ell)$, and $t_{\text{SA}} = t_{\text{locate}} = O(s)$ for arbitrarily large alphabets; it needs $nH_k + O(n \frac{\log n}{s}) + O(n \frac{H_k}{\log \log n}) + O(n)$ bits. We apply Transformation 2 with $\tau = \log \log n$. The construction algorithm for this index relies on randomized algorithm for constructing an mmphf functions [6]; therefore the update procedures of our dynamic data structure also rely on randomization in this case. The resulting dynamic index uses $nH_k + O(n \frac{\log n}{s}) + O(n \frac{\log \sigma}{\log \log n}) + O(n)$ bits. This index achieves $t_{\text{range}} = O(|P| \log \log n)$, $t_{\text{extract}} = O(s + \ell)$, $t_{\text{locate}} = O(s)$. Insertions and deletions are supported in $O(|T| \log^\varepsilon n)$ time and $O(|T|(\log^\varepsilon n + s))$ expected time respectively. If counting queries are also supported, then $t_{\text{count}} = O(|P| \log \log n + \log n)$ and updates take $O(|T| \log n)$ expected time.

The index of Barbay et al. [3, 2] is also $(\log^\varepsilon n, \log \sigma)$ -constructible and uses $nH_k + O(n \frac{\log n}{s}) + o(n \log \sigma)$ bits. If the alphabet size $\sigma = \log^{O(1)} n$, this index achieves $t_{\text{range}} = O(|P|)$, $t_{\text{extract}} = O(s + \ell)$, and $t_{\text{locate}} = O(s)$; it uses $nH_k + O(n \frac{\log n}{s}) + o(n \log \sigma)$ bits. If we set $\tau = \log \log n$ and apply Transformation 2, we obtain a dynamic data structure with $t_{\text{range}} = O(|P| \log \log n)$, $t_{\text{extract}} = O(s + \ell)$, and $t_{\text{SA}} = t_{\text{locate}} = O(s)$. For an arbitrary alphabet size σ , the index of Barbay et al. [3, 2] achieves $t_{\text{range}} = O(|P| \log \log \sigma)$, $t_{\text{extract}} = O((s + \ell) \log \log \sigma)$, and $t_{\text{SA}} = t_{\text{locate}} = O(s \log \log \sigma)$. Again we set $\tau = \log \log n$ and apply Transformation 2. We obtain a dynamic index that has query costs $t_{\text{range}} = O(|P| \log \log \sigma \log \log n)$, $t_{\text{extract}} = O((s + \ell) \log \log \sigma)$, and $t_{\text{SA}} = t_{\text{locate}} = O(s \log \log \sigma)$. Insertions and deletions are supported in $O(|T| \log^\varepsilon n)$ time and $O(|T|(\log^\varepsilon n + s))$ time respectively. If counting queries are also supported, then $t_{\text{count}} = O(|P| \log \log n \log \log \sigma + \log n)$ (resp. $t_{\text{count}} = O(|P| \log \log n + \log n)$ if $\sigma = \log^{O(1)} n$) and updates take $O(|T| \log n)$ time.

The index of Grossi and Vitter [19] is $(\log^\varepsilon n, \log \sigma)$ -constructible. It achieves $t_{\text{locate}} = O(\log^\varepsilon n)$, $t_{\text{range}} = O(|P|/\log_\sigma n + \log^\varepsilon n)$ and $t_{\text{extract}} = O(\ell/\log_\sigma n)$. We apply Transformation 2 with $\tau = 1/\delta$ for a constant δ . The resulting dynamic index uses $O(n \log \sigma (1 + 1/\delta)) = O(n \log \sigma)$ bits and has the following query costs: $t_{\text{locate}} = O(\log^\varepsilon n)$, $t_{\text{range}} = O(|P|/\log_\sigma n + \log^\varepsilon n)$, $t_{\text{extract}} = O(\ell/\log_\sigma n)$. As shown in the full version [29], in this case the data structure for uncompressed sequence \mathcal{C}_0 relies on hashing. Therefore the update procedure is randomized. Updates are supported in $O(|T| \log^{2\varepsilon} n)$ expected time, but we can replace ε with $\varepsilon/2$ in our construction and reduce the update time to $O(|T| \log^\varepsilon n)$. If counting queries are also supported, then $t_{\text{count}} = O(|P|/\log_\sigma n + \log n/\log \log n)$ and updates take $O(|T| \log n)$ expected time. If we want to support updates using a deterministic procedure, then the cost of searching in \mathcal{C}_0 grows to $O(|P|(\log \log n)^2/\log_\sigma n + \log n)$. In

this case $t_{\text{range}} = t_{\text{count}} = O(|P|(\log \log n)^2/\log_\sigma n + \log n)$, $t_{\text{locate}} = O(\log^\varepsilon n)$, and $t_{\text{extract}} = O(\ell/\log_\sigma n)$.

5. DYNAMIC GRAPHS AND BINARY RELATIONS

Let R denote a binary relation between t objects and σ_l labels. In this section we denote by n the cardinality of R , i.e., the number of object-label pairs. We will assume that objects and labels are integers from intervals $[1, \sigma_l]$ and $[1, t]$ respectively. Barbay et al. [4] showed how a static relation R can be represented by a string S . A dynamization of their approach based on dynamic data structures for rank and select queries is described in [31].

Let M be a matrix that represents a binary relation R ; columns of R correspond to objects and rows correspond to matrices. The string S is obtained by traversing M columnwise (i.e., objectwise) and writing the labels. An additional bit string N encodes the numbers of labels related to objects: $N = 1^{n_1}01^{n_2}0\dots 1^{n_t}$, where n_i is the number of labels related to the i -th object. Using rank, select, and access queries on N and S , we can enumerate objects related to a label, enumerate labels related to an object, and decide whether an object and a label are related.

Deletion-Only Data Structure.

We keep R in S and N described above; S and N are stored in static data structures. If a pair (e, l) is deleted from R , we find the element of S that encodes this pair and mark it as deleted. We record marked elements (i.e. pairs that are deleted but are still stored in the data structure) in a bit vector D : $D[i] = 0$ if and only if the pair $S[i]$ is marked as deleted. We maintain the data structure of Lemma 3 on D . Moreover we keep D in a data structure described in [17]; this data structure enables us to count the number of 1-bits in any range of D . For each label a we also keep a data structure D_a . D_a is obtained by traversing the a -th row of M : if $M[a, j] \neq 0$, then we append 0 to D_a if (a, j) is marked as deleted; if $M[a, j] \neq 0$ and (a, j) is not marked as deleted, we append 1 to D_a . For each D_a we also maintain data structures for reporting and counting 1-bits described above. Finally we record indices of deleted labels and objects in two further bit sequences. The static data structures on S and N are implemented as in [2], so that rank and select queries are answered in $O(\log \log \sigma_l)$ time and any $S[i]$ or $N[i]$ can be retrieved in constant time.

If we need to list labels related to an object i , we first find the part of S that contains these labels. Let $l = \text{rank}_1(\text{select}_0(i - 1, N), N)$ and $r = \text{rank}_1(\text{select}_0(i, N), N)$. We list all elements of $S[l..r]$ that are not marked as deleted by enumerating all 1-bits in $D[l..r]$. Then we access and report $S[i_1], S[i_2], \dots, S[i_f]$, where i_1, i_2, \dots, i_f are positions of 1-bits in $D[l..r]$. In order to list objects related to a label a , we find positions of 1-bits in D_a . Then we access and report $\text{select}_a(j_1, S), \text{select}_a(j_2, S), \dots$, where j_1, j_2, \dots denote positions of 1-bits in D_a . In order to determine whether an object i and a label a are related, we compute $d = \text{rank}_a(r, S) - \text{rank}_a(l, S)$, where l and r are as defined above. If $d = 0$, then the object i and the label a are not related. If $d = 1$, we compute $j = \text{select}_a(\text{rank}_a(r, S), S)$; i and a are related if and only if $D[j] = 1$.

When (e, l) is deleted, we find the position j of (e, l) in S and set $D[j] = 0$; j can be found with a constant num-

ber of rank and select queries. We also set $D_a[j'] = 0$ for $j' = \text{rank}_a(S, j)$. When an empty label or an empty object is removed, we simply record this fact by adding it to a compact list of empty labels (resp. empty objects). When the number of pairs that are marked as deleted exceeds n/τ , we start the process of re-building the data structure. The cost of re-building is distributed among the following updates; we will give a more detailed description in the exposition of the fully-dynamic data structure.

Fully-Dynamic Data Structure.

We split a binary relation R , regarded as a set of object-label pairs, into subsets and keep these subsets in data structures $\mathbb{C}_0, \mathbb{C}_1, \dots, \mathbb{C}_r, \mathbb{L}_1, \dots, \mathbb{L}_r$, and $\mathbb{T}_1, \dots, \mathbb{T}_g$ for $g = \Theta(\tau)$. We set the parameter $\tau = \log \log n$. Only \mathbb{C}_0 is stored in a fully-dynamic data structure, but we can afford to keep \mathbb{C}_0 in $O(\log n)$ bits per item because it contains only a small fraction of pairs. All other pairs are stored in deletion-only data structures described above. Distribution of pairs among subsets and procedures for re-building deletion-only data structures are the same as in Section 3. To simplify a description, we will not distinguish between a subset and a data structure that stores it.

\mathbb{C}_0 contains at most $\max_0 = 2n/\log^2 n$ pairs. Each structure \mathbb{C}_i for $r \geq i \geq 1$ contains at most $\max_i = 2n/\log^{2-i\epsilon} n$ pairs. Every \mathbb{T}_i contains at most $2n/\tau$ pairs. Data structure \mathbb{C}_0 contains object-label pairs in uncompressed form and uses $O(\log n)$ bits per pair. For every object i that occurs in \mathbb{C}_0 we keep a list L_i that contains all labels that occur in pairs $(i, \cdot) \in \mathbb{C}_0$; for each label a that occurs in \mathbb{C}_0 we keep a list of objects that occur in pairs $(\cdot, a) \in \mathbb{C}_0$. Using these lists we can enumerate all objects related to a label or labels related to an object in \mathbb{C}_0 in $O(1)$ time per datum. If we augment lists L_i with predecessor data structures described in [1], we can also find out whether an object i and a label a are related in $O((\log \log \sigma_l)^2)$ time.

All pairs in $\mathbb{C}_1, \mathbb{L}_1, \dots, \mathbb{C}_r, \mathbb{L}_r$, and $\mathbb{T}_1, \mathbb{T}_\tau$ are kept in deletion-only data structures described above. A new object-label pair (i, a) is inserted into \mathbb{C}_0 if \mathbb{C}_0 contains less than \max_0 pairs. Otherwise we look for the smallest j , $0 \leq j < r$, such that $|\mathbb{C}_{j+1}| + |\mathbb{C}_j| + 1 \leq \max_{j+1}$. We lock \mathbb{C}_j by setting $\mathbb{L}_j = \mathbb{C}_j$, $\mathbb{C}_j = \emptyset$ and initiate the process of creating $\mathbb{N}_{j+1} = \mathbb{C}_j \cup \mathbb{C}_{j+1} \cup \{(i, a)\}$. If $|\mathbb{C}_{i+1}| + |\mathbb{C}_i| + 1 \leq \max_{i+1}$ for all $i < r$, we lock \mathbb{C}_r and start the process of constructing $\mathbb{N}_{j+1} = \mathbb{C}_r \cup \{(i, a)\}$. The cost of creating \mathbb{N}_j is distributed among the next \max_j updates in the same way as in Section 3. We observe that data structures Temp_i are not needed now because each update inserts only one element (pair) into the relation R . We guarantee that each structure \mathbb{C}_i for some $1 \leq i \leq r$ contains at most $\max_i/2$ pairs marked as deleted and \mathbb{T}_i for $1 \leq i \leq r$ contains an $O(1/\tau)$ fraction of deleted pairs. Procedures for re-building data structures that contain too many pairs marked as deleted are the same as in Section 3.

Our fully-dynamic data structure must support insertions and deletions of new objects and labels. An object that is not related to any label or a label that is not related to any object can be removed from a data structure. This means that both the number of labels σ_l and the number of objects t can change dynamically. Removing and inserting labels implies changing the alphabets of strings S that are used in deletion-only data structures. Following [31] we store two global tables, NS and SN ; SN maps labels to

integers bounded by $O(\sigma_l)$ (global label alphabet) and NS maps integers back to labels. We also keep bitmaps GC_i and GT_i, GL_i , and GN_i for all subsets \mathbb{C}_i, L_i, N_i , and \mathbb{T}_i . $GC_i[j] = 1$ if the label that is assigned to integer j occurs in \mathbb{C}_i and $GC_i[j] = 0$ otherwise; GT_i, GL_i , and GN_i keep the same information for subsets \mathbb{T}_i, L_i , and N_i . Using these bit sequences we can map the symbol of a label in the global alphabet to the symbol of the same label in the effective alphabet⁴ used in one of subsets. When a label a is deleted, we mark $SN[a]$ as free. When a new label a' is inserted, we set $SN[a']$ to a free slot in SN (a list of free slots is maintained). When some subset, say \mathbb{C}_i is re-built, we also re-build the bit sequence GC_i .

In order to list objects related to a label a , we first report all objects that are related to $SN[a]$ and stored in \mathbb{C}_0 . Then we visit all subsets \mathbb{C}_i, L_i , and \mathbb{T}_i and report all objects related to $\text{rank}_1(SN[a], GC_i)$, $\text{rank}_1(SN[a], GL_i)$, and $\text{rank}_1(SN[a], GT_i)$ respectively. We remark that a global symbol of a label can be mapped to a wrong symbol in the local effective alphabet. This can happen if some label a' is removed and its slot in $SN[]$ is assigned to another label a but the bitmap of say GC_i is not yet re-built. In this case $\text{rank}_1(SN[a], GC_i)$ will map a to the symbol for the wrong label a' . But a' can be removed only if all object-label pairs containing a' are deleted; hence, all pairs (i, a') in \mathbb{C}_i are marked as deleted and the query to \mathbb{C}_i will correctly report nothing. We can report labels related to an object and tell whether a certain object is related to a certain label using a similar procedure. We visit $O(\log \log n)$ data structures in order to answer a query. In all data structures except for \mathbb{C}_0 , we spend $O(\log \log \sigma_l)$ time per reported datum. An existential query on \mathbb{C}_0 takes $O((\log \log \sigma_l)^2)$ time; all other queries on \mathbb{C}_0 take $O(1)$ time per reported datum. Hence all queries are answered in $O(\log \log n \log \log \sigma_l)$ time per reported datum. A counting query takes $O(\log n / \log \log n)$ time in each subset. Hence, we can count objects related to a label or labels related to an object in $O(\log n)$ time.

All bit sequences D and D_a in all subsets use $O((n/\tau) \log \tau)$ bits. Every string S stored in a deletion-only data structure needs $|S|H_0(S) + o(|S| \log \sigma_l)$ bits. Hence all strings S use at most $nH + o(n \log \sigma_l)$ bits, where $H = \sum_{1 \leq a \leq \sigma_l} \frac{n_i}{n} \log \frac{n}{n_i}$. Bit sequences GC_i, GL_i , and GT_i use $O(\sigma_l \tau) = o(n \log \sigma_l)$ bits. Now we consider the space usage of bit sequences N stored in deletion-only data structures. Let m_i denote the number of pairs in a data structure \mathbb{T}_i . N consists of m_i 1's and t 0's. If $m_i > t$, then the bit sequence N stored as a part of \mathbb{T}_i uses $m_i \log \frac{m_i+t}{m_i} = O(m_i)$ bits. If $t \geq m_i$, N uses $O(m_i \log \tau)$ bits because $m_i = \Theta(n/\tau)$. Hence all N stored in all \mathbb{T}_i use $O(n \log \tau)$ bits. In our data structure we set $\tau = \log \log n$. If $\sigma_l = \Omega(\log^{1/4} n)$, $O(n \log \tau) = o(n \log \sigma_l)$. Otherwise $t = \Omega(n/\log n)$ because $n \leq t \cdot \sigma_l$; if $t = \Omega(n/\log n)$, $O(n \log \tau) = o(t \log t)$. Data structures that are re-built at any moment of time contain $O(n/\tau)$ elements and use $O(\frac{n}{\tau} \log \sigma_l) = o(n \log \sigma_l)$ bits. Extra space that we need to store elements marked as deleted is bounded by $o(n \log \sigma_l)$; this can be shown in the same way as in Section 3.

THEOREM 2. *A dynamic binary relation that consists of n pairs relating t objects to σ_l labels can be stored in $nH + o(n \log \sigma_l) + o(t \log t) + O(t + n + \sigma_l \log n)$ bits where $H =$*

⁴An effective alphabet of a sequence S contains only symbols that occur in S at least once.

$\sum_{1 \leq a \leq \sigma_1} \frac{n_a}{n} \log \frac{n}{n_a}$ and n_a is the number of objects related to a label a . We can determine whether an object and a label are related in $O(\log \log \sigma_1 \log \log n)$ time and report all objects related to a label (resp. all labels related to an object) in $O((k+1) \log \log \sigma_1 \log \log n)$ time, where k is the number of reported items. We can count objects related to a label or labels related to an object in $O(\log n)$ time. Updates are supported in $O(\log^\epsilon n)$ time.

Directed graph is a frequently studied instance of a binary relation. In this case both the set of labels and the set of objects are identical with the set of graph nodes. There is an edge from a node u to a node v if the object u is related to the label v .

THEOREM 3. *A dynamic directed graph that consists of σ_1 nodes and $n \geq \sigma_1$ edges can be stored in $nH + o(n \log \sigma_1) + O(n + \sigma_1 \log n)$ bits where $H = \sum_{1 \leq a \leq \sigma_1} \frac{n_a}{n} \log \frac{n}{n_a}$ and n_a is the number of outgoing edges from node a . We can determine if there is an edge from one node to another one in $O(\log \log \sigma_1 \log \log n)$ time and report all neighbors (resp. reverse neighbors) of a node in $O((k+1) \log \log \sigma_1 \log \log n)$ time, where k is the number of reported nodes. We can count neighbors or reverse neighbors of a node in $O(\log n)$ time. Updates are supported in $O(\log^\epsilon n)$ time.*

6. CONCLUSIONS

In this paper we described a general framework for transforming static compressed indexes into dynamic ones. We showed that, using our framework, we can achieve the same or almost the same space and time complexity for dynamic indexes as was previously obtained by static indexes. Our framework is applicable to a broad range of static indexes that includes a vast majority of currently known results in this area. Thus, using our techniques, we can easily modify almost any compressed static index, so that insertions and deletions of documents are supported. It will likely be possible to apply our framework to static indexes that will be obtained in the future. Our approach also significantly reduces the cost of basic queries in compact representations of dynamic graphs and binary relations. We expect that our ideas can be applied to the design of other compressed data structures.

Acknowledgments.

The authors wish to thank Djamel Belazzougui for clarifying the construction time of the static index in [6] and Gonzalo Navarro for explaining some technical details of dynamic indexes used in [26].

7. REFERENCES

- [1] A. Andersson and M. Thorup. Dynamic ordered sets with exponential search trees. *J. ACM*, 54(3):13, 2007.
- [2] J. Barbay, F. Claude, T. Gagie, G. Navarro, and Y. Nekrich. Efficient fully-compressed sequence representations. *Algorithmica*, 69(1):232–268, 2014.
- [3] J. Barbay, T. Gagie, G. Navarro, and Y. Nekrich. Alphabet partitioning for compressed rank/select and applications. In *Proc. 21st ISAAC*, pages 315–326 (part II), 2010.
- [4] J. Barbay, A. Golynski, J. I. Munro, and S. S. Rao. Adaptive searching in succinctly encoded binary relations and tree-structured documents. *Theoretical Computer Science*, 387(3):284–297, 2007.
- [5] J. Barbay, M. He, J. I. Munro, and S. S. Rao. Succinct indexes for strings, binary relations and multi-labeled trees. *ACM Transactions on Algorithms*, 7(4):article 52, 2011.
- [6] D. Belazzougui and G. Navarro. Alphabet-independent compressed text indexing. In *Proc. 19th Annual European Symposium on Algorithms (ESA 2011)*, pages 748–759, 2011.
- [7] J. L. Bentley and J. B. Saxe. Decomposable searching problems i: Static-to-dynamic transformation. *J. Algorithms*, 1(4):301–358, 1980.
- [8] H. Chan, W.-K. Hon, T.-H. Lam, and K. Sadakane. Compressed indexes for dynamic text collections. *ACM Transactions on Algorithms*, 3(2):article 21, 2007.
- [9] H.-L. Chan, W.-K. Hon, and T.-W. Lam. Compressed index for a dynamic collection of texts. In *Proc. 15th CPM*, LNCS 3109, pages 445–456, 2004.
- [10] S. Chaudhuri, V. Ganti, and L. Gravano. Selectivity estimation for string predicates: Overcoming the underestimation problem. In *Proc. 20th International Conference on Data Engineering, (ICDE 2004)*, pages 227–238, 2004.
- [11] P. F. Dietz and D. D. Sleator. Two algorithms for maintaining order in a list. In *Proc. 19th Annual ACM Symposium on Theory of Computing (STOC 1987)*, pages 365–372, 1987.
- [12] J. D. Fernández, M. A. Martínez-Prieto, and C. Gutierrez. Compact representation of large RDF data sets for publishing and exchange. In *Proc. 9th International Semantic Web Conference, (ISWC 2010)*, pages 193–208, 2010.
- [13] P. Ferragina, G. Manzini, V. Mäkinen, and G. Navarro. Compressed representations of sequences and full-text indexes. *ACM Transactions on Algorithms*, 3(2):article 20, 2007.
- [14] M. Fredman and M. Saks. The cell probe complexity of dynamic data structures. In *Proc. 21st STOC*, pages 345–354, 1989.
- [15] R. González and G. Navarro. Improved dynamic rank-select entropy-bound structures. In *Proc. 8th LATIN*, LNCS 4957, pages 374–386, 2008.
- [16] R. González and G. Navarro. Rank/select on dynamic compressed sequences and applications. *Theoretical Computer Science*, 410:4414–4422, 2009.
- [17] R. González and G. Navarro. Rank/select on dynamic compressed sequences and applications. *Theor. Comput. Sci.*, 410(43):4414–4422, 2009.
- [18] R. Grossi, A. Gupta, and J. S. Vitter. High-order entropy-compressed text indexes. In *Proc. 14th SODA*, pages 841–850, 2003.
- [19] R. Grossi and J. S. Vitter. Compressed suffix arrays and suffix trees with applications to text indexing and string matching. *SIAM J. Comput.*, 35(2):378–407, 2005.
- [20] M. He and J. I. Munro. Succinct representations of dynamic strings. In *Proc. 17th SPIRE*, pages 334–346, 2010.

- [21] H. V. Jagadish, R. T. Ng, and D. Srivastava. Substring selectivity estimation. In *Proc. 18th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS 1999)*, pages 249–260, 1999.
- [22] P. Krishnan, J. S. Vitter, and B. R. Iyer. Estimating alphanumeric selectivity in the presence of wildcards. In *Proceedings of the 1996 ACM SIGMOD International Conference on Management of Data (SIGMOD1996)*, pages 282–293, 1996.
- [23] S. Lee and K. Park. Dynamic rank-select structures with applications to run-length encoded texts. In *Proc. 18th CPM*, LNCS 4580, pages 95–106, 2007.
- [24] S. Lee and K. Park. Dynamic rank/select structures with applications to run-length encoded texts. *Theoretical Computer Science*, 410(43):4402–4413, 2009.
- [25] V. Mäkinen and G. Navarro. Dynamic entropy-compressed sequences and full-text indexes. In *Proc. 17th CPM*, LNCS 4009, pages 307–318, 2006.
- [26] V. Mäkinen and G. Navarro. Dynamic entropy-compressed sequences and full-text indexes. *ACM Transactions on Algorithms*, 4(3):article 32, 2008.
- [27] G. Manzini. An analysis of the burrows-wheeler transform. *J. ACM*, 48(3):407–430, 2001.
- [28] C. W. Mortensen, R. Pagh, and M. Patrascu. On dynamic range reporting in one dimension. In *Proc. 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 104–111, 2005.
- [29] J. I. Munro, Y. Nekrich, and J. S. Vitter. Dynamic data structures for document collections and graphs. *CoRR*, abs/1503.05977, 2015.
- [30] G. Navarro and V. Mäkinen. Compressed full-text indexes. *ACM Comput. Surv.*, 39(1), 2007.
- [31] G. Navarro and Y. Nekrich. Optimal dynamic sequence representations. In *Proc. 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2013)*, pages 865–876, 2013.
- [32] G. Navarro and K. Sadakane. Fully-functional static and dynamic succinct trees. *ACM Transactions on Algorithms*, 10(3):article 16, 2014.
- [33] A. Orlandi and R. Venturini. Space-efficient substring occurrence estimation. In *Proc. 30th ACM SIGMOD-SIGACT-SIGART Symposium on Principles of Database Systems (PODS 2011)*, pages 95–106, 2011.
- [34] K. Sadakane. New text indexing functionalities of the compressed suffix arrays. *J. Algorithms*, 48(2):294–313, 2003.

APPENDIX

A.1. REPORTING 1-BITS IN A BIT VECTOR

We show how to store a bit vector with a small number of zeros in small space, so that all 1-values in an arbitrary range can be reported in optimal time. This result is used by our method that transforms a static index into an index that supports deletions. We start by describing an $O(n)$ -bit data structure. Then we show how space usage can be reduced to $O((n \log \tau)/\tau)$

LEMMA 2. *There exists an $O(n)$ -bit data structure that supports the following operations on a bit vector B of size n : (i) $zero(i)$ sets $B[i] = 0$ (ii) $report(s, e)$ enumerates all j such that $s \leq j \leq e$ and $B[j] = 1$. Operation $zero(i)$ is supported in $O(\log^\epsilon n)$ time and a query $report(s, e)$ is answered in $O(k)$ time, where k is the number of output bit positions.*

PROOF. We divide the vector B into words $W_1, \dots, W_{\lceil n/\log n \rceil}$ of $\log n$ bits. We say that a word W_t is non-empty if at least one bit in W_t is set to 1. We store the indices of all non-empty words in a data structure that supports range reporting queries in $O(k)$ time, where k is the number of reported elements, and updates in $O(\log^\epsilon n)$ time [28]. For every word W_i we can find the rightmost bit set to 1 before the given position p or determine that there is no bit set to 1 to the right of p in $O(1)$ time. This can be done by consulting a universal look-up table of size $o(n)$ bits. To report positions of all 1-bits in $B[s..e]$, we find all non-empty words whose indices are in the range $\lceil s/\log n \rceil, \lfloor e/\log n \rfloor$. For every such word, we output the positions of all 1-bits. Finally, we also examine the words $W_{\lfloor s/\log n \rfloor}$ and $W_{\lceil e/\log n \rceil}$ and report positions of 1-bits in these two words that are in $B[s..e]$. The total query time is $O(k)$. Operation $zero(i)$ is implemented by setting the bit $i - \lfloor i/\log n \rfloor \log n$ in the word $W_{\lfloor i/\log n \rfloor}$ to 0. If $W_{\lfloor i/\log n \rfloor}$ becomes empty, we remove $\lfloor i/\log n \rfloor$ from the range reporting data structure. \square

LEMMA 3. *Let B be a bit vector of size n with at most $O(\frac{n}{\tau})$ zero values for $\tau = O(\log n / \log \log n)$. B can be stored in $O(n \frac{\log \tau}{\tau})$ -bit data structure that supports the following operations on B : (i) $zero(i)$ sets $B[i] = 0$ (ii) $report(s, e)$ enumerates all j such that $s \leq j \leq e$ and $B[j] = 1$. Operation $zero(i)$ is supported in $O(\log^\epsilon n)$ time and a query $report(s, e)$ is answered in $O(k)$ time, where k is the number of output bit positions.*

PROOF. We divide B into words W_i of τ bits. Indices of non-empty words are stored in the the data structure B' , implemented as in Lemma 2. Every word W_i is represented as follows: we store the number of zeros in W_i using $O(\log \tau)$ bits. A subword with f zeros, where $0 \leq f \leq \tau$, is encoded using $f(\log \tau)$ bits by specifying positions of 0-bits. For every word W_i we can find the rightmost bit set to 1 before the given position p or determine that there is no bit set to 1 to the right of p in $O(1)$ time. This can be done by consulting a universal look-up table of size $o(n)$ bits. Query processing is very similar to Lemma 2. To report 1-bits in $B[s..e]$, we find all non-empty words whose indices are in the range $\lceil s/\tau \rceil, \lfloor e/\tau \rfloor$. For every such word, we output the positions of all 1-bits. Finally, we also examine the words $W_{\lfloor s/\tau \rfloor}$ and $W_{\lceil e/\tau \rceil}$ and report positions of 1-bits in these two words that are in $B[s..e]$.

Operation $zero(i)$ is implemented by setting the corresponding bit in some word W_i to 0 and changing the word encoding. If W_i becomes empty, the i -th bit in B' is set to 0. Issues related to memory management can be resolved as in [31].

We need $O(n/\tau)$ bits to store the data structure B' for non-empty words. Let n_f denote the number of words

with f zero values. All words W_i need $\sum_{f=1}^{\tau} n_f \cdot f \cdot \log \tau = \log \tau \sum n_f \cdot f = O((n/\tau) \log \tau)$ because $\sum n_f \cdot f = O(n/\tau)$. \square