

Orthogonal Range Searching for Text Indexing

Moshe Lewenstein*

Bar-Ilan University
moshe@cs.biu.ac.il

Abstract. Text indexing, the problem in which one desires to preprocess a (usually large) text for future (shorter) queries, has been researched ever since the suffix tree was invented in the early 70's. With textual data continuing to increase and with changes in the way it is accessed, new data structures and new algorithmic methods are continuously required. Therefore, text indexing is of utmost importance and is a very active research domain.

Orthogonal range searching, classically associated with the computational geometry community, is one of the tools that has increasingly become important for various text indexing applications. Initially, in the mid 90's there were a couple of results recognizing this connection. In the last few years we have seen an increase in use of this method and are reaching a deeper understanding of the range searching uses for text indexing.

In this monograph we survey some of these results.

1 Introduction

The *text indexing* problem assumes a (usually very large) text that is to be preprocessed in a fashion that will allow efficient future queries of the following type. A query is a (significantly shorter) pattern. One wants to find all text locations that match the pattern in time proportional to the *pattern length and number of occurrences*.

Two classical data structures that are most widespread amongst all the data structures solving the text indexing problem are the *suffix tree* [104] and the *suffix array* [87] (see Section 2 for definitions, time and space usage).

While text indexing for exact matches is a well studied problem, many other text indexing related problems have become of interest as the field of text indexing expands. For example, one may desire to find matches within subranges of the text [86], or to find which documents of a collection contain a searched pattern [90], or one may want our text index compressed [93].

Also, the definition of a match may vary. We may be interested in a *parameterized* match [15, 85], a *function* match [5], a *jumbled* match [4, 14, 25, 33, 89] etc.

* This paper was written while on Sabbatical in U. of Waterloo. This research was supported by the U. of Waterloo and BSF grant 2010437, a Google Research Award and GIF grant 1147/2011.

These examples are only a very few of the many different interesting ways that the field of text indexing has expanded.

New problems require more sophisticated ideas, new methods and new data structures. This indeed has happened in the realm of text indexing. New data structures have been created and known data structures from other domains have been incorporated for the use of text indexing data structures all mushrooming into an expanded, cohesive collection of text indexing methods. One of these incorporated methods is that of orthogonal range searching problems.

Orthogonal range searching refers to the preprocessing of a collection of points in d -dimensional space to allow queries on ranges defined by rectangles whose sides are aligned with the coordinate axes (orthogonal).

In the problems we consider here we assume that all input point sets are in rank space, i.e., they have coordinates on the integer grid $[n]^d = \{0, \dots, n-1\}^d$. The rank-space assumption can easily be made less restrictive, but we do not dwell on this here as the rank-space assumption works well for most of the results here.

The set of problems one typically considers in range searching are queries on the range such as emptiness, reporting (all points in the range), report any (one) point, range minimum/maximum, closest point. In general, some function on the set of points in the range.

We will consider different range searching variants in the upcoming sections and will discuss the time and space complexity of each at the appropriate place. For those interested in further reading of orthogonal range searching problems we suggest starting with [1, 28].

Another set of orthogonal range searching problems is on arrays (not point sets¹). We will lightly discuss this type of orthogonal range searching, specifically for Range Minimum Queries (RMQ).

In this monograph we take a look at some of the solutions to text indexing problems that have utilized range searching techniques. The reductions chosen are, purposely, quite straightforward with the intention of introducing the simplicity of the use of this method. Also, it took some time for the pattern matching community to adopt this technique into their repertoire. Now more sophisticated reductions are emerging and members of the community have also been contributing to better range searching solutions, reductions for hardness and more.

2 Problem Definitions and Preliminaries

Given a string S , $|S|$ is the length of S . Throughout this paper we denote $n = |S|$. An integer i is a *location* or a *position* in S if $i = 1, \dots, |S|$. The substring $S[i, \dots, j]$ of S , for any two positions $i \leq j$, is the substring of S that begins at index i and ends at index j . The *suffix* S_i of S is the substring $S[i, \dots, n]$.

¹ Nevertheless, the problems mentioned here can all be transformed to point sets using orthogonal range searching, if one so desired. This is done by adding a dimension and changing the values to be coordinates on the last dimension.

Suffix Tree. The *suffix tree* [48, 88, 101, 104] of a string S , denoted $ST(S)$, is a compact trie of all the suffixes of $S\$$ (i.e., S concatenated with a delimiter symbol $\$ \notin \Sigma$, where Σ is the alphabet set, and for all $c \in \Sigma, \$ < c$). Each of its edges is labeled with a substring of S (actually, a representation of it, e.g., the start location and its length). The “compact” property is achieved by contracting nodes having a single child. The children of every node are sorted in the lexicographical order of the substrings on the edges leading to them. Consequently, each leaf of the suffix tree represents a suffix of S , and the leaves are sorted from left to right in the lexicographical order of the suffixes that they represent. $ST(S)$ requires $O(n)$ space. The suffix tree can be prepared in $O(n + \text{Sort}(\Sigma))$, where n is the text size, Σ is the alphabet, and $\text{Sort}(Q)$ is the time required to sort the set Q [48]. For the suffix tree one can search an m -length pattern in $O(m + \text{occ})$, where occ is the number of occurrences of the pattern. If the alphabet Σ is large this potentially increases to $O(m \log |\Sigma| + \text{occ})$, as one needs to find the correct edge exiting at every node. If randomization is allowed then one can introduce hash functions at the nodes to obtain $O(m + \text{occ})$, even if the alphabet is large, without affecting the original $O(n + \text{Sort}(\Sigma))$ construction time.

Suffix Array. The *suffix array* [71, 87] of a string S , denoted $SA(S)$, is a permutation of the indices $1, \dots, n$ indicating the lexicographic ordering of the suffixes of S . For example, consider $S = \text{mississippi}$. The suffix array of S is $[11, 8, 5, 2, 1, 10, 9, 7, 4, 6, 3]$, that is $S_{11} = \text{”i”} < S_8 = \text{”ippi”} < S_5 = \text{”issippi”} < \dots < S_3 = \text{”ssissippi”}$, where $<$ denotes less-than lexicographically. The construction time of a suffix array is $O(n + \text{Sort}(\Sigma))$ [71]. The time to answer a query P of length m on the suffix array is $O(m + \log n + \text{occ})$ [87]². The $O(m + \log n)$ is required to find the range of suffixes (see Section 3.1 for details) which have P as a prefix and then since P appears as a prefix of suffix S_i it must appear at location i of the string S . Hence, with a scan of the range we can report all occurrences in additional $O(\text{occ})$ time.

Relations between the Suffix Tree and Suffix Array. Let $S = s_1 s_2 \dots s_n$ be a string. Let $SA = SA(S)$ be its suffix array and $ST = ST(S)$ its suffix tree. Consider ST ’s leaves. As these represent suffixes and they are in lexicographic ordering, ST is actually a tree over SA . In fact, one can even view ST as a search tree over SA .

Say we have a pattern P whose path from the root of ST ends on the edge entering node v in ST (the locus). Let $l(v)$ denote the leftmost leaf in the subtree of v and $r(v)$ denote the rightmost leaf in the subtree of v . Assume that i is the location of SA that corresponds to $l(v)$, i.e. the suffix $S_{SA[i]}$ is associated with $l(v)$. Likewise assume j corresponds to $r(v)$. Then the range $[i, j]$ contains all the suffixes that begin with P and it is maximal in the sense that no other suffixes begin with P . We call this range the *SA-range of P* .

Consider the previous example $S = \text{mississippi}$ with suffix array $[11, 8, 5, 2, 1, 10, 9, 7, 4, 6, 3]$. For a query pattern $P = \text{”si”}$ we have that the SA-

² This requires LCP information. Details appear in Section 3.1.

range for P is $[8, 9]$, i.e. P is a common prefix of $\{s_{SA[8]} \dots s_n, s_{SA[9]} \dots s_n\} = \{sissippi, sippi\}$.

Beforehand, we pointed out that finding the SA-range for a given P takes $O(m + \log n)$ in the suffix array. However, given the relationship between a node in the suffix tree and the SA-range in the suffix array, if we so desire, we can use the suffix tree as a search tree for the suffix array and find the SA-range in $O(m)$ time. For simplification of results, throughout this paper we assume that indeed we find SA-ranges for strings of length m in $O(m)$ time.

Moreover, one can find for $P = p_1, \dots, p_m$ all nodes in a suffix tree representing p_i, \dots, p_m for $1 \leq i \leq m$ in $O(m)$ time using suffix links. Hence, one can find all SA-ranges for p_i, \dots, p_m for $1 \leq i \leq m$ in $O(m)$ time.

3 1D Range Minimum Queries

While the rest of this paper contains results for orthogonal range searching in rank space, one cannot disregard a couple of important range searching results that are widely used in text indexing structures. The range searching we refer to is the *Range Minimum Query* (RMQ) problem on an array (not a point set). RMQ is defined as follows.

Let S be a set of linearly ordered elements whose elements can be compared (for \leq) in constant time.

<i>d</i> -Dimensional Range Minimum Query (d-RMQ)	
Input:	A d -dimensional array A over S of size $N = n_1 \cdot n_2 \cdot \dots \cdot n_d$ where n_i is the size of dimension i .
Output:	A data structure over A supporting the following queries.
Query:	Return the minimum element in a range $q = [a_1..b_1] \times [a_2..b_2] \times \dots \times [a_d..b_d]$ of A .

1-dimensional RMQ plays an important role in text indexing data structures. Hence, we give a bit of detail on results about RMQ data structure construction.

The 1-dimensional RMQ problem has been well studied. Initially, Gabow, Bentley and Tarjan [54] introduced the problem. They reduced the problem to the *Lowest Common Ancestor (LCA)* problem [61] on Cartesian Trees [103]. The *Cartesian Tree* is a binary tree defined on top of an array of n elements from a linear order. The root is the minimum element, say at location i of the array. The left subtree is recursively defined as the Cartesian tree of the sub-array of locations 1 to $i - 1$ and the right subtree is defined likewise on the sub-array from $i + 1$ to n . It is quite easy to see the connection between the RMQ problem and the Cartesian tree, which is what was utilized in [54], where the LCA problem was solved optimally in $O(n)$ time and $O(n)$ space while supporting $O(1)$ time queries. This, in turn, yielded the result of $O(n)$ preprocessing time and space for the 1D RMQ problem with answers in $O(1)$ time.

Sadakane [97] proposed a position-only solution, i.e. one that return the position of the minimum rather than the minimum itself, of $4n + o(n)$ bits space with $O(1)$ query time. Fischer and Heun [53] improved the space to $2n + o(n)$ bits and preprocessed in $O(n)$ time for subsequent $O(1)$ time queries. They also showed that the space must be of size $2n - O(\log n)$. Davoodi, Raman and Rao [42] showed how to achieve the same succinct representation in a different way with $o(n)$ working space, as opposed to the $n + o(n)$ working space in [53]. It turns out that there are two different models, the *encoding model* and the *indexing model*. The model difference was already noted in [44]. For more discussion on the modeling differences see [23]. In the encoding model we preprocess the array A to create a data structure *enc* and queries have to be answered using *enc* only, *without* access to A . In the indexing model, we create an index *idx* and are able to refer to A when answering queries. The result of Fischer and Heun [53] is the encoding model result. For the indexing model Brodal et al. [23] and Fischer and Heun [53], in parallel, showed that an index of size $O(n/g)$ bits is possible with query time $O(g)$. Brodal et al. [23] showed that this is an optimal tradeoff in the indexing model.

Range minimum queries on an array have been extended to 2D in [8, 13, 22, 23, 43, 56] and to higher dimension d in [13, 23, 31, 41].

3.1 The LCP Lemma

The Longest Common Prefix (LCP) of two strings plays a very important role in text indexing and other string matching problems. So, define as follows.

Definition 1. *Let x and y be two strings over an alphabet Σ . The longest common prefix of x and y , denoted $LCP(x, y)$, is the largest string z that is a prefix of both x and y . The length of $LCP(x, y)$ is denoted $|LCP(x, y)|$.*

The first sophisticated use of the LCP function for string matching was for string matching with errors in a paper by Landau and Vishkin [82]. An interesting and very central result to text indexing structures appears in the following lemma, which is not difficult to verify.

Lemma 1. [87] *Let $S^{[1]}, S^{[2]}, \dots, S^{[n]}$ be a sequence of n lexicographically ordered strings. Then $|LCP(S^{[i]}, S^{[j]})| = \min_{i \leq h < j} |LCP(S^{[h]}, S^{[h+1]})|$.*

This allows a data structure over the suffix array of size $O(n)$ that returns the LCP value of any two substrings in $O(1)$ time. This is done by building an RMQ data structure over the array containing the values of the LCP of lexicographically consecutive suffixes and using the lemma.

This result was implicitly³ used in [87] to reduce the $O(m \log n)$ time for a search of an m -length pattern P in a suffix array indexing a text of length n to

³ They did not actually use the RMQ data structure. Rather, since they know the path a binary search will follow, they know which interval one needs to (RMQ)query when consulting a given suffix array position (there is only one path towards it in the virtual binary search tree). So they directly store that range LCP value.

$O(m + \log n)$. The idea is as follows. Both find the SA-range for P based on a binary search of the pattern P on the suffixes of the suffix array. The $O(m \log n)$ time follows for a naive binary search because it takes $O(m)$ time to check if P is a prefix of a suffix and the $O(\log n)$ follows from the binary search.

Reducing to $O(m + \log n)$ is done as follows. The binary search is still used. Initially P is compared to the string in the center of the lexicographic ordering. This may take $O(m)$ time. However, at every stage of the binary search we maintain $|\text{LCP}(P, T_i)|$ for the suffix T_i of T with the maximal $|\text{LCP}(P, T_i)|$, over all the suffixes to which P has already been compared. When comparing P to the next suffix, say T_j , in the binary search, first $|\text{LCP}(T_i, T_j)|$ is evaluated (in constant time) if $|\text{LCP}(T_i, T_j)| \neq |\text{LCP}(P, T_i)|$ we immediately know the value of $|\text{LCP}(P, T_j)|$ - give it a moment of thought - and we can compare the character at location $|\text{LCP}(P, T_j)|+1$ of P and T_j and continue the binary search from there. Otherwise, $|\text{LCP}(T_i, T_j)| = |\text{LCP}(P, T_i)|$ in which case we continue the comparison of P and T_j (but only) from the $|\text{LCP}(P, T_j)|+1$ -th character. Hence, one can claim, in an amortized sense, that the pattern is scanned only once. So, the search time is $O(m + \log n)$.

The dynamic version of this method is much more involved but has interesting applications, see [9].

3.2 Document Retrieval

The *Document Retrieval* problem is very close to the text indexing problem. Here we are given a collection of documents D_1, \dots, D_k and desire to preprocess them in order to answer document queries Q . A *document query* asks for the set of documents where Q appears.

The *generalized suffix array* (for generalized suffix tree, see [60]) is a suffix array for a collection of texts T_1, \dots, T_k and can be viewed as the suffix array for $T_1\$1T_2\$2 \dots \$_{k-1}T_k$. However, we may remove, before finalizing the suffix array, all suffixes that start with a delimiter as they contain no interesting information. In order to solve the document retrieval problem one can build a generalized suffix array for D_1, \dots, D_k . The problem is that when one seeks a query Q one will find all the occurrences of Q in all documents, whereas we desire to know only *in which* documents Q appears and are not interested in all match locations.

A really neat trick to solve this problem was proposed by Muthukrishnan [90]. Imagine the generalized suffix array for D_1, \dots, D_k of size $n = \sum_{1 \leq i \leq k} |D_i|$ and a document retrieval query Q of length m . In $O(m + \log n)$, or even in $O(m)$ time (as discussed in the end of Section 2) it is possible to find the SA-range for Q . Now we'd like to report all documents who have a suffix in this range. So, create a *document array* for the suffix array. The *document array* for D_1, \dots, D_k will be of length n and will contain at location i the document id d if $\text{SA}[i]$ is a suffix beginning in document d . So, the former problem now becomes the problem of finding the unique id's in the SA-range of the document array.

Muthukrishnan [90] proposed a transformation to the RMQ problem in the following sense. Take the document array DA and generate, yet another, array

which we will call the *predecessor document array*. Let $\psi(i) = j$ if $j < i$, $DA(i) = DA(j)$ and for all $j < k < i$, $DA(k) \neq DA(i)$. $\psi(i) = -1$ if there is no such j . The predecessor document array has $\psi(i)$ at location i . The following observation now follows.

Lemma 2. *Let D_1, \dots, D_k be a collection of documents and let SA be their generalized suffix array. Let Q be a query and let $[i, j]$ be the SA-range of Q . There is a one-one mapping between the documents in range $[i, j]$ in the document array and the values $< i$ in range $[i, j]$ in the predecessor document array.*

Proof. Let $i_1 < i_2 < \dots < i_r$ be all locations in $[i, j]$ where document id d appears in the document array. Then the i_1 -th location of the predecessor document array will be $< i$. However, locations $i_2 < i_3 < \dots < i_r$ will contain i_1, i_2, \dots, i_{r-1} , all greater than or equal to i , in the predecessor document array. □

Hence, it is natural to consider an extended RMQ problem defined now.

Bounded RMQ	
Input:	An array A .
Output:	A data structure over A supporting the following <i>bounded RMQ</i> queries.
Query:	Given a range $[i, j]$ and a number b find all values in the range $[i, j]$ of value $< b$.

The bounded RMQ problem can be solved by recursively applying the known RMQ solution. Find an RMQ on $A[i, j]$, say it is at location r . If it is less than b then reiterate on $A[i, r - 1]$ and $A[r + 1, j]$. The preprocessing time and space are the same as those of the RMQ problem. The query time is $O(ans)$, where ans is the number of elements smaller than b .

This yields an $O(m + docc)$ solution for the document retrieval problem, where $docc$ is the number of documents in which the query Q appears.

4 Indexing with One Error

The problem of *approximate text indexing*, i.e. the text indexing problem where up to a given number of errors is allowed in a match is a much more difficult problem than text indexing. The problem is formally defined as follows.

- Input:** Text T of length n over alphabet Σ and an integer k .
- Output:** A data structure for T supporting k -error queries.
- Query:** A k -error query is a pattern $Q = q_1q_2 \dots q_m$ of length m over alphabet Σ for which we desire to find all locations i in T where Q matches with $\leq k$ errors.

We note that there are several definitions of errors. The *edit distance* allows for mismatches, insertions and deletions [84], the *Hamming distance* allows for mismatches only. For text indexing with k errors (for Hamming distance, Edit distance and more) Cole et al. [36] introduced a novel data structure which, for the Hamming distance version, uses space $(n \log^k n)$ (it is preprocessed within an $O(\log \log n)$ factor of the space complexity) and answers queries in $O(\log^k n + m + occ)$. See also [26, 100] for different space/time tradeoffs for the Hamming distance version.

Throughout the rest of this section we focus and discuss the special case of one error. Moreover, we will do so for the mismatch error, but a similar treatment will handle insertions and deletions. The reduction to range queries presented in this section was obtained in parallel by Amir et al. [10] and by Ferragina, Muthukrishnan and de Berg [51]. The goal of [51] was to show geometric data structures that solve certain methods in object oriented programming. They also used their data structure to solve the dictionary matching with one error. In [10] Amir et al. solved dictionary matching with one error and also solved the text indexing with one error. For the sake of simplicity, we will present the result of text indexing with one error from [10], but the reduction is the same for dictionary matching (see definition in [10]).

The algorithm that we will shortly describe combines a bidirectional construction of suffix trees, which had been known before. Specifically, it is similar to the data structure of [21]. However, in [21] a reduction to 2D range searching was not used.

4.1 Bidirectional Use of Suffix Arrays

For simplicity's sake we make the following assumption. *Assume that there are no exact matches of the pattern in the text.* We will relax this assumption later and show how to handle it in Section 4.3.

The Main Idea: Assume there is a pattern occurrence at text location i with a single mismatch in location $i + j - 1$. This means that $q_1 \dots q_{j-1}$ has an *exact match* at location i and $q_{j+1} \dots q_m$ has an exact match at location $i + j$.

The distance between location i and location $i + j$ is dependent on the mismatch location, and that is somewhat problematic. We therefore choose to “wrap” the pattern around the mismatch. In other words, if we stand exactly at location $i + j - 1$ of the text and look left, we see $q_{j-1} \dots q_1$. If we look right we see $q_{j+1} \dots q_m$. This leads to the following algorithm.

For the data structure supporting 1-mismatch queries construct a suffix array SA_T of text string T and a suffix array SA_{T^R} of the string T^R , where T^R is the reversed text $T^R = t_n \dots t_1$.

In order to reply to the 1-mismatch queries do as follows:




Fig. 1. A range query grid representing T and T^R . The dashed rectangle represents a mismatch at location 2 of Q and the solid rectangle represents a mismatch at location 3 of Q .

Query Reply:

For $l = 1, \dots, m$ do

1. Find the maximal SA_T -range $I_l = [i_l, j_l]$ of $q_{l+1} \dots q_m$ in SA_T , if it is non-empty.
2. Find the maximal SA_{T^R} -range $RI_l = [ri_l, rj_l]$ of $q_{l-1} \dots q_1$ in SA_{T^R} , if it is non-empty.
3. If both I_l and RI_l are non-empty, then return the intersection of SA_T and SA_{T^R} on their respective ranges.

Steps 1 and 2 of the query reply can be done for the l 's in *overall* linear time (see end of Section 2). Hence, we only need an efficient implementation of Step 3.

4.2 Set Intersection via Range Reporting

In Step 3, given SA-ranges $I_l = [i_l \dots j_l]$ and $RI_l = [ri_l \dots rj_l]$ we want to report the points in the intersection of SA_T and SA_{T^R} w.r.t. to the corresponding coordinates of the two ranges. We show that this quite straightforwardly reduces to the 2D range reporting problem.

Range Reporting in 2D (rank space)	
Input:	A point set $P = \{(x_1, y_1), \dots, (x_n, y_n)\} \subseteq [1, n] \times [1, n]$.
Output:	A data structure representing P that supports the following <i>range reporting</i> queries.
Query:	Given a range $R = [a, b] \times [c, d]$ report all points of P contained in R .

Since the arrays (the suffix array for T and the suffix array for T^R) are permutations, every number between 1 and $n + 1$ (we include suffix ϵ) appears precisely once in each array. The coordinates of every number i are (x_i, y_i) , where $x_i = SA_T^{-1}(i)$ and $y_i = SA_{T^R}^{-1}(n - i + 3)$ (the choice of $n - i + 3$ is to align the appropriate reverse suffix with suffix i — explanation: i becomes $n - i + 1$ when reversing the text. Then one needs to move over one to the mismatch location and one more to the next location). We define the point set to be $P = \{(x_2, y_2), \dots, (x_{n+1}, y_{n+1})\}$ and construct a 2D range reporting structure for it (efficiency to be discussed in a moment). It is clear that the range elements intersection corresponds precisely with a 2D query $I_l \times RI_l$.

The current best range reporting data structures in 2D are as follows:

1. **Alstrup, Brodal and Rauhe [2]:** a data structure requiring $O(n \log^\epsilon n)$ space, for any constant ϵ , that can answer queries in $O(\log \log n + k)$, where k is the number of points reported.
2. **Chan, Larsen and Pătraşcu [28]:** a data structure requiring $O(n \log \log n)$ space that can answer queries in $O(\log \log n(1 + k))$.
3. **Chan, Larsen and Pătraşcu [28]:** a data structure requiring $O(n)$ space that can answer queries in $O(\log^\epsilon n(1 + k))$. Other succinct results of interest appear in a footnote⁴.

Therefore, we have the following.

Theorem 1. *Let $T = t_1 \dots t_n$ and $Q = q_1 \dots q_m$. When no exact match exists, indexing with one error can be solved with $O(s(n))$ space such that queries can be answered in $O(qt(n, m, occ))$ time, where:*

- $s(n)$ = the space for a range reporting data structure and
- occ = the number of occurrences of Q in T with one error, and
- $qt(n, m, occ)$ = the query time for the same range reporting data structure.

Proof: Other than the range reporting data structure the space required is $O(n)$. Likewise, Steps 1 and 2 of the query response require total time $O(m)$ for $j = 1, \dots, m$. Hence, the space and time are dominated by the range reporting data structure at hand, i.e. space $O(s(n))$ and query time $qt(n, m, occ)$. □

⁴ Note that prior succinct solutions show a novel adaptation of the method of Chazelle [29] to Wavelet Trees [58], see [70], [86] and [16]. It is especially worth reading the chapter of "Application as Grids" in [91] for more results along this line.

4.3 Indexing with One Error When Exact Matches Exist

We assumed that the text contained no exact pattern occurrence in the text. In fact, the algorithm would also work for the case where there are exact pattern matches in the text, but its time complexity would suffer. Recall that the main idea of the algorithm was to pivot a pattern position and check, for every text location, whether the pattern to the left and to the right of the pivot were exact matches. However, if the pattern occurs as an exact match in the text then at that occurrence a match is announced for all m pivots. So, this means that every exact occurrence is reported m times. The worst case could end up being as bad as $O(m * occ)$ (for example if the text is a^n and the pattern is a^m then it would be $O(nm)$).

To handle the case of exact occurrences one can use the following idea. Add a third dimension to the range reporting structure representing the character in the text at the mismatch location. The desired intersection is of all suffix labels such that this character is *different from* the symbol at that respective pattern location. This leads to a specific variant of range searching.

3D 5-Sided Range Reporting (rank space)	
Input:	A point set $P = \{(x_1, y_1, z_1), \dots, (x_n, y_n, z_n)\}$ $\subseteq [1, n] \times [1, n] \times [1, n]$.
Output:	A data structure representing P that supports the following <i>range reporting</i> queries.
Query:	Given a range $R = [a, b] \times [c, d] \times [e, \infty]$ report all points of P contained in R .

3D 5-Sided Range Reporting can be solved with space $O(n \log^{O(\epsilon)} n)$ and query time $O(occ + \log \log n)$ [28].

Back to our problem. We need to update the preprocessing phase.

Preprocessing: Preprocess for 3-dimensional range queries on the matrix $[1, \dots, n] \times [1, \dots, n] \times \Sigma$. If Σ is unbounded, then use only the $O(n)$ symbols in T . The new geometric points are (x_i, y_i, z_i) , where x_i and y_i will be the same as before and $z_i = t_{i-1}$ will be the text character that needs to mismatch. This will be added in the preprocessing stage.

The only necessary modification is for Step 3 of the query reply which becomes:

3. If I_l and RI_l both exist, then return all the points in $I_l \times R_l \times \Sigma$ for which the z-coordinate is not the respective pattern mismatch symbol.

The above step can be implemented by two 3D 5-sided range queries on the three dimensional range $I_l \times R_l \times [1, a - 1]$ and $I_l \times R_l \times [a + 1, |\Sigma|]$ where a is the current pattern symbol being examined. We assume that the alphabet symbols are numbered $1, \dots, |\Sigma|$.

See Figure 2 depicting the 3D queries.




Fig. 2. The 3D 5-sided range queries on a 3D grid representing T and T^R and the appropriate "not to match" character. The dashed boxes represent a mismatch at location 2 of Q and the solid boxes represent a mismatch at location 3 of Q .

Theorem 2. Let $T = t_1 \dots t_n$ and $Q = q_1 \dots q_m$. Indexing with one error can be solved with $O(n \log^{O(\epsilon)} n)$ space and $O(occ + m \log \log n)$ query time, where occ is the number of occurrences of the pattern in the text with at most one error.

Proof: As in Theorem 1, the space and query time of the solution are dominated by the range searching structure. Hence, using the results of [28], the space is $O(n \log^{O(\epsilon)} n)$ and the query time is $O(occ + m \log \log n)$. \square

4.4 Related Material

For a succinct index for dictionary matching the best result appears in [62]. A *wildcard* character is one that matches all other symbols. When used we denote them with ϕ .

Iliopoulos and Rahman [67] consider the problem of indexing a text T to answer queries of the form $Q = Q_1 \phi^d Q_2$, where $d, |Q_1|$ and $|Q_2|$ are known during the preprocessing. This is known as gap-indexing. Bille and Gørtz [19] consider the same problem. However, they required only d to be known in advance. The solution in [19] uses a reduction to range reporting. The reduction they apply is similar to the one presented in this chapter. In fact, it is easier because the gap's location within the query pattern is known. So, one does not need to check every position of the pattern as one does in the case of mismatch. Hence, the 2D

range reporting data structure is sufficient. One does need to adapt the search for a difference of d instead of 1. This requires setting $y_i = n - i + d + 2$.

5 Compressed Full-Text Indexes

Given that texts may be very large it makes sense to compress them - and there are many methods to do so. On the other hand, these are not constructed to allow text indexing. Doing both at once has been the center of a lot of research activity over the last decade. However, reaching this stage has happened in phases. Initially, pattern matching on compressed texts was considered starting by Amir et al. [6]. By pattern matching on compressed texts we mean that a compressed text, with some predefined compressor, and a pattern are given and the goal is to find the occurrences of the pattern efficiently without decompressing the text. The second phase was text indexing while maintaining a copy of the original text and augmenting it with some sublinear data structure (usually based on a compressor) that would allow text indexing, e.g. [72]. We will call this phase, the *intermediate phase*. Finally, compressed full-text indexing was achieved, that is text indexing without the original text and with a data structure with size depending on the compressibility of the string. See Section 3 for the discussion on the encoding model vs. the indexing model. Compressed full-text indexes of note are that of Ferragina and Manzini [50], the *FM-Index*, that of Grossi and Vitter [59], the *Compressed Suffix Array* and that of Grossi, Gupta and Vitter [58], the *Wavelet Tree*. For an extensive survey on compressed full-text indexes see [93].

In this section we will present two results, that of Kärkkäinen and Ukkonen [72] and that of Claude and Navarro [34]. The latter uses a central idea of the former. However, the former is from the intermediate phase. So, it maintains a text to reference. This certainly makes the indexing easier. The latter is a compressed full-text index. The former uses the more general LZ77 and the latter uses SLPs (Straight Line Programs). Both LZ77 and SLP compressions are defined in this section.

5.1 LZ77 Compressed Indexing

The Lempel-Ziv compression schemes are among the best known and most widely used. In this section (and in Section 7) we will be interested in the variant known as LZ77 [107]. For sake of completeness we describe the LZ77 scheme here.

An Overview of the Lempel-Ziv Algorithm. Given an input string S of length n , the algorithm encodes the string in a greedy manner from left to right. At each step of the algorithm, suppose that we have already encoded $S[1, \dots, k-1]$ with $i-1$ phrases $\rho_1, \dots, \rho_{i-1}$ (phrase - to be defined shortly). We search for the location t , such that $1 \leq t \leq k-1$, for which the longest common prefix of $S_k = S[k, \dots, n]$ and the suffix S_t is maximal. Once we have found the desired location, suppose the aforementioned longest common prefix

is the substring $S[t, \dots, r]$, a *phrase*, ρ_i , will be added to the output which will include F_i the encoding of the distance to the substring (i.e., the value $k - t$), L_i the length of the substring (i.e., the value $r - t + 1$), and the next character $C_i = S[k + (r - t + 1)]$. The algorithm continues by encoding $S_{k+(r-t+1)+1} = S[k + (r - t + 1) + 1, \dots, n]$. The sequence of phrases is called the string's (LZ77) *parse* and is defined:

$$Z = \rho_1 = (F_1, L_1, C_1), \rho_2 = (F_2, L_2, C_2), \dots, \rho_{c(n)} = (F_{c(n)}, L_{c(n)}, C_{c(n)})$$

We denote with $u(i) = \sum_{j=1}^{i-1} (L_j + 1) + 1$ the start location of ρ_i in the string S .

Finally, we denote the output of the LZ77 algorithm on the input S as $LZ(S)$.

Kärkkäinen and Ukkonen Method. Farach and Thorup [47], in their paper on search in LZ77 compressed texts, noted a neat, useful observation. Say T of length n is the text to be compressed and Z is its LZ parse containing $c(n)$ phrases.

Lemma 3. [47] *Let Q be a query pattern and let j be the smallest integer such that $Q = T[j \dots j + |Q| - 1]$. Then $u(i) \in [j, j + |Q| - 1]$ for some $1 \leq i \leq c(n)$.*

In other words, the first appearance of Q in T cannot be contained in a single phrase. Kärkkäinen and Ukkonen [72] utilized this lemma to show how to augment the text with a sublinear text indexing data structure based on LZ77.

The idea is as follows. Say we search for a pattern $Q = q_1, \dots, q_m$ in the text T which has been compressed by LZ77. The occurrences of Q in T are defined differently for those that intersect with more than one phrase (which must exist if there are any matches according to Lemma 3) and those that are completely contained in a single phrase. The former are called *primary occurrences* and the latter are called *secondary occurrences*. The algorithm proposed in [72] first finds primary occurrences and then uses the primary occurrences to find secondary occurrences. Both steps use orthogonal range searching schemes.

In order to find the primary occurrences they used a bi-directional scheme based on Lemma 3. Take every phrase $\rho_i = t_{u(i)}, \dots, t_{u(i+1)-1}$ and creates its reverse $\rho_i^R = t_{u(i+1)-1}, \dots, t_{u(i)}$. Now, consider a primary occurrence and the first phrase it starts in, say ρ_i . Then there must be a j such that q_1, \dots, q_j is a suffix of ρ_i and q_{j+1}, \dots, q_m is a prefix of the suffix of T , $t_{u(i+1)}, \dots, t_n$. So, to find the primary occurrences it is sufficient to find k such that:

1. q_j, \dots, q_1 is a prefix of ρ_i^R .
2. q_{j+1}, \dots, q_m is a prefix of the suffix $t_{u(i+1)}, \dots, t_n$ of T .

Let π be the lexicographic sort of $\rho_1^R, \rho_2^R, \dots, \rho_{c(n)}^R$, i.e. $\rho_{\pi(1)}^R < \rho_{\pi(2)}^R < \dots < \rho_{\pi(c(n))}^R$. This leads to a range reporting scheme, similar to that of the previous section, of size $c(n) \times c(n)$, where the point set $P = \{(x(i), y(i)) \mid x(i) = SA_T^{-1}(u(i + 1)), y(i) = \pi^{-1}(i)\}$.

Therefore, for every $1 \leq j \leq m$ we need to (1) find the range of π for which every $\rho_{\pi(i)}^R$ within has q_j, \dots, q_1 as a prefix and (2) find the range of suffixes that have q_{j+1}, \dots, q_m as a prefix. Once we do so we revert to range reporting, as in the previous section. However, finding the ranges is not as simple as in the previous section. Recall that we want to maintain the auxiliary data in sublinear space. So, saving a suffix array for the text is not a possibility. Also, the reversed phrases need to be indexed for finding the range of relevant phrases.

To this end a sparse suffix tree was used in [73]. A *sparse suffix tree* is a compressed trie over a subset of the suffixes and is also known as a *Patricia Trie* over this suffix set. As it is a compressed trie it is of size $O(\text{subset size})$, in our case $O(c(n))$. Lately, in [18] it was shown how to construct a sparse suffix array in optimal space and near-optimal time.

The sparse suffix tree can be constructed for the suffixes $T_{u(i+1)}$ and, since we have the original text on hand, we can maintain the compressed suffix tree in $O(c(n))$ words. Navigation on this tree is the same as in a standard suffix tree. For the reversed phrases we can associate each with a prefix of the text. Reversing them gives a collection of suffixes of the reversed text. Now construct a sparse suffix tree for these reversed suffixes and this will allow finding the range in π (with the help of the existing text) as in the previous section. Hence,

Theorem 3. *Let T be a text and Z its LZ77 parse. One can construct a text indexing scheme which maintains T along with a data structure of size $O(|Z|)$ such that for a query Q we can find all its primary occurrences in $O(|Q|(|Q| + \log^\epsilon |Z|) + occ * \log^\epsilon |Z|)$, where occ is the number of primary occurrences.*

Proof. For the query Q , as described above, each pattern position is evaluated for primary occurrences. That is for each of the $|Q|$ pattern positions, we first traverse the auxiliary sparse suffix trees in $O(|Q|)$ time and once the ranges (for that pattern position is found) we perform a range query. The traversal will cost $O(|Q|^2)$ time.

Using the 2D succinct range reporting of Chan et al. [28] (see previous section) we have linear space, i.e. $O(|Z|)$, and $O(\log^\epsilon |Z|(1+k))$ query time where k is the number of points found, which is the same as the number of primary occurrences found.

Hence, over the $|Q|$ pattern positions the range querying will cost $\sum_{i=1}^{|Q|} \log^\epsilon |Z|(1 + occ_i)$ time, where occ_i is the number of primary occurrences when we split the pattern at position i . However, $\sum_{i=1}^{|Q|} \log^\epsilon |Z|(1 + occ_i) = \sum_{i=1}^{|Q|} \log^\epsilon |Z| + \sum_{i=1}^{|Q|} \log^\epsilon |Z| * occ_i = |Q| \log^\epsilon |Z| + \log^\epsilon |Z| occ$. Recalling that the traversal cost $O(|Q|^2)$ time yields the desired. □

We note that the secondary occurrences still need to be found. This is another interesting part of the paper and we refer the interested reader to [72].

See the following papers for more along the following line using Lempel Ziv compressors [12, 50, 81, 96].

5.2 SLP Text Indexing

Claude and Navarro [34] proposed a full-text indexing scheme based on *straight line programs (SLPs)*. An *SLP* is a grammar based compression for a text T . The grammar produces exactly one word T and the rules are in Chomsky Normal Form, i.e. each rule is $A \rightarrow BC$, where A, B, C are variables of the grammar or $A \rightarrow a$, where A is a variable and a is a terminal (a character of T).

The text indexing scheme that they propose follows the previous idea [72] of finding primary and secondary occurrences. However, for SLPs things are slightly different. Consider the derivation tree for the text T , that is deriving the full word T by generating from the start symbol S as the root (if $S \rightarrow AB$ then A and B will be children of the root S in the derivation tree - from here the derivation tree is applied recursively until the full T is spelled out in the left-to-right order of the leaves). Every occurrence of a pattern Q in the text T has a unique lowest variable V which produces this occurrence, but its children do not. That is if the children of V are V_1 and V_2 , i.e. there is a rule $V \rightarrow V_1V_2$, then V_1 produces x, q_1, \dots, q_j (where $x \in \Sigma^*$) and V_2 produces q_{j+1}, \dots, q_m, y (where $y \in \Sigma^*$) for some j . We say that V splits pattern Q at location j . An occurrence of Q is called a *primary occurrence* if for some V and j V splits this occurrence of Q . All other occurrences are *secondary occurrences*.

The format of the algorithm is to, once again, find the primary occurrences and then to deduce the occurrences of Q in the text therefrom. With the goal of finding the primary occurrences in mind, once again, our grid will be of size $c(n) \times c(n)$, where $c(n)$ is the size of the variable set of the grammar. Each side of the grid will have one coordinate for each variable. The range searching point set is defined per rule, $V \rightarrow V_1V_2$. The location (x, y) on the grid corresponding to V_1 (for x) on one side and V_2 (for y) on the other will have a point, V , on the grid. The ordering of the variables on either side of the grid follows from the desire to satisfy the following conditions.

1. q_j, \dots, q_1 is a prefix of V_i^R .
2. q_{j+1}, \dots, q_m is a prefix of V_l .
3. There is a rule $V \rightarrow V_iV_l$.

It is easy to see that the desired ordering, as in the LZ77 scheme, has the phrases in the x -coordinate in reverse lexicographic ordering and has the phrases in the y -coordinate in lexicographic ordering. The challenge here is to actually find the range of variables where q_j, \dots, q_1 is a prefix of V_i^R . This is because it is a full-text index and the text is not accessible any more. Nevertheless, this is doable in the SLP compression scheme using a suffix array type of search and comparing q_j, \dots, q_1 with the variable at hand. This comparison is not trivial. However, the full scheme is out of scope of this survey and we refer the reader to the full paper [34]. The result achieved is as follows:

Theorem 4. *Let T be a text of size N represented by an SLP with n variables and height h . There is a representation using $n(\log N + 3 \log n + O(\log |\Sigma| + \log h) + o(\log n))$ bits such that Q of length m can be found in $O((m(m + h) + hocc) \log n)$ query time.*

An extension of this idea to a general grammar can be found in [35], where the dependency on h was removed from the search time. There is also other work for different compressors. See [55] for one of the latest.

6 Weighted Ancestors

6.1 2-Sided Sorted Range Reporting in 2D

In this section we consider the *2-sided sorted range reporting* problem⁵ which is defined now.

2-Sided Sorted Range Reporting in 2D	
Input:	A point set $P = \{(x_1, y_1), \dots, (x_n, y_n)\} \subseteq [1, U] \times [1, U]$.
Output:	A data structure representing P that supports the following <i>2-sided sorted range reporting</i> queries.
Query:	Given a range $R = [-\infty, a] \times [-\infty, b]$ report all points of P contained in R sorted by their y -coordinate (from highest to lowest).

Note that we deviate from the assumption that the points are in rank-space. This is important for the application of this section. We now show a method to solve the 2-sided sorted range reporting. The idea is as follows.

Consider the dynamic predecessor problem in which we need to support the following operations (a) insertions/deletions of integers ($\in [1, U]$) and (b) predecessor queries. This is a classical problem and is solved with a van-Emde Boas tree [102] or with y -fast tries [105] in $O(n)$ space (n is the current number of integers) and $O(\log \log U)$ time for the operations, where $[1, U]$ is the domain of the elements.

Dietz and Raman [45] asked whether this could be made partially persistent⁶ within the same query times. In other words can one create a data structure where insertions and deletions are supported on the current version but predecessor queries can be made on any of the versions (current or previous) of the data structure. Recently, Chan [27] accomplished this by constructing a partially persistent predecessor data structure with space $O(n)$ and operations time $O(\log \log U)$. Chan's result [27] is in fact more general, showing that the first predecessor can be found (in any previous version) in $O(\log \log U)$ time but the predecessor of the predecessor (etc.) can be found in $O(1)$ time. This yields a time of $O(\log \log U + k)$ to find the k previous elements in sorted order in a chosen version of the data structure.

⁵ Results in Section 6.1 stem from wonderful research chats with Timothy Chan.

⁶ Actually Dietz and Raman [45] asked about persistency in general, which may refer to full persistence or partial persistence. We stick to partial persistence as it is sufficient for our needs.

We utilize this for the 2-sided sorted range reporting by creating a data structure for P as follows⁷. Consider the sort of the x -coordinates of P , i.e. some permutation π for which $x_{\pi(1)} < x_{\pi(2)} < \dots < x_{\pi(n)}$. Now we insert the y -coordinates into the data structure according to π . That is we insert $y_{\pi(1)}$ and then $y_{\pi(2)}$ until $y_{\pi(n)}$. Now, a 2-sided sorted range reporting query $R = [-\infty, a] \times [-\infty, b]$ is answered as follows; first use a predecessor query to find a within $x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(n)}$ - that is find i such that $x_{\pi(i)} \leq a < x_{\pi(i+1)}$. Then we go to the i -th copy of the partially persistent data structure which contains the points $(x_{\pi(1)}, y_{\pi(1)}), (x_{\pi(2)}, y_{\pi(2)}), \dots, (x_{\pi(i)}, y_{\pi(i)})$. Hence, the points are exactly the points that satisfy that their x -coordinate $\in [-\infty, a]$. Now to find the relevant points ($y \in [-\infty, b]$) in R sorted by their y -coordinate we need to apply the predecessor query. This yields an $O(\log \log U + occ)$ when using the data structure from [27]. Hence,

Theorem 5. *The 2-sided sorted range reporting problem on an n -point set over a $U \times U$ grid can be solved with $O(n)$ space and $O(\log \log U + occ)$ time.*

6.2 Weighted Ancestors to 2-Sided Range Successor in 2D

Consider the weighted ancestors problem on an edge-weighted tree introduced by Farach and Muthukrishnan [46] for the sake of obtaining a perfect-hash for substrings. An edge-weighted tree is a tree where each edge e has a weight $w(e) \in [1, U]$. Each node v is associated with a weight $w(v) = \sum_{e \in p_v} w(e)$, where p_v is the path from root-to- v . The *weighted ancestors* problem is defined as follows.

Input: An edge-weighted tree T with weight function w .

Output: A data structure supporting *weighted ancestor queries*.

Query: Given a node u and a threshold t find the ancestor v of u such that $w(v) \geq t$, but $w(p(v)) < t$, where $p(v)$ is the parent of v .

The weighted ancestor problem is a natural extension of the predecessor problem to trees. The application considered by [46] was on suffix trees. A suffix tree can be viewed as an edge-weighted tree with the edge weights denoting the length of the text with which the edge is marked. Now, say you are given indices i and j and want to find the locus of $T[i \dots j]$ in the suffix tree. This can be done by going to the leaf representing i and asking a weighted ancestor query with threshold $j - i + 1$. The answer to the query is the locus of $T[i \dots j]$.

In [46] a solution was given with $O(n \log n)$ preprocessing time, $O(n)$ space, and $O(\log \log U)$ query time. Their solution is based on a heavy path decomposition in order to linearize the input tree. Each path of the heavy path decomposition is assigned a predecessor structure. The preprocessing time of $O(n \log n)$ can be improved to $O(n)$ and this has been pointed out in [11, 79]. It should be

⁷ We point out that for our purposes, finding one successor, the results of Dietz and Raman [45] are sufficient because (a) we seek only one successor and (2) the insertions are done first and then the queries are asked.

mentioned that the authors of [46] were considering a PRAM model and hence the $O(n \log n)$ time is really $O(\log n)$ parallel time and $O(n)$ work. Lately, it was shown that if the depth of the answer is d in the tree then the query can be answered in $O(\log \log d)$ time [78].

We now present a solution for this problem using 2-sided sorted range reporting.

Consider the edge-weighted input tree T . We assume that every internal node in the tree has at least two children. Otherwise, create a dummy child with an arbitrary edge weight, say 1. Now consider the leaves l_1, \dots, l_k ordered in inorder. For every two adjacent leaves denote their lowest common ancestor with $LCA(l_i, l_{i+1})$ and $nw(i) = w(LCA(l_i, l_{i+1}))$. With one scan of the tree all these values are computable. Now generate an array of the nw values. Say we are given a weighted ancestor query, node u and threshold t . We may assume that u is a leaf. Otherwise, we simply choose a descendant leaf to represent u (the answer will be the same). Consider the node v which is the answer to the query and consider its parent $p(v)$. Since $p(v)$ has at least two children v has at least one sibling. Say, v has a sibling to its left (in inorder). Let $u = l_i$ then in the nw array the first location $j < i$ that satisfies $nw(j) < t$ is the node for which the $LCA(l_j, l_i) = p(v)$. To obtain this j we revert to 2-sided sorted range reporting. We set the points on the grid to be $P = (j, nw(j))$. The query is bounded by i in the x -coordinates and t in the y -coordinates. What we are looking for is the first answer, the element with the largest y -coordinate.

Note that once this is done it is still necessary to find v (we only obtained $p(v)$). This can be done with a predecessor structure for each node. That is, for each node $p(v)$ save the index h of the leftmost leaf l_h for each of $p(v)$'s children. A predecessor query with i will return the correct edge with child v , the weighted ancestor of u . Hence,

Theorem 6. *Let T be an n node edge-weighted tree with weights from $[1, U]$. Then using 2-sided sorted range reporting one can answer weighted ancestor queries in $O(\log \log U)$ time. The space required is $O(n)$.*

Note that for a suffix tree, the motivation in [46], the weights are from $[1, n]$. So, the query time for a suffix tree is $O(\log \log n)$.

Recall (from Section 3) the definition of an SA-range and its relation to the suffix tree. Hence, the method just described precisely finds the boundaries of the SA-range for a given suffix i (in the suffix array) and its prefix of length (threshold) t .

7 Compressed Substring Retrieval

In this section we are concerned with the *substring compression* problem. The *substring compression* problem was introduced in [37]. Some of the definitions and layout here are from [37]. The solution, specifically the reduction to *range successor queries*, is from [75].

In *substring compression* one is given a text to preprocess so that, upon request, a compressed substring is returned. The goal is to do so quickly, preferably in $O(c(s))$ time, where $c(s)$ is the size of the compressed substring. *Generalized substring compression* is the same with the following twist. The queries contain an additional context substring (or a collection of context substrings) and the answers are the substring in compressed format, where the context substring is used to make the compression more efficient.

The compressor of interest is, once again, LZ77. We use the terminology from Section 5. Some extra terminology is as follows. The string S may be encoded within the context of the string T . We denote this by $\text{LZ}(S \mid T)$. The encoded result will be equivalent to the result when LZ77 is performed on the concatenated string $T\$S$, where $\$$ is a symbol that does not appear in either S or T . However, only the portion of $\text{LZ}(T\$S)$ which represents the compression of S is output by the algorithm.

Formally, given a string S of length n , we wish to preprocess S in such a way that allows us to efficiently answer the following queries:

Substring Compression Query (SCQ(i, j)): given any two indices i and j , such that $1 \leq i \leq j \leq n$, we wish to output $\text{LZ}(S[i, \dots, j])$.

Generalized Substring Compression Query (GSCQ(i, j, α, β)): given any four indices i, j, α , and β , such that $1 \leq i \leq j \leq n$ and $1 \leq \alpha \leq \beta \leq n$, we wish to output $\text{LZ}(S[i, \dots, j] \mid S[\alpha, \dots, \beta])$.

The goal is to do answer queries quickly. Query times for both of the above query types will strongly depend on the number of phrases actually encoded. We denote these as $C(i, j)$ and $C_{\alpha, \beta}(i, j)$ for SCQ and GSCQ, respectively.

7.1 SCQ to Range Successor in 2D

Recall the definition of LZ77 from Section 5. Imagine that we have already computed the phrases for $S[i \dots k - 1]$ and desire to compute the next phrase which is a prefix of $S[k \dots j]$. In other words, we want to find the location $i \leq t \leq k - 1$ for which the longest common prefix of $S[k, \dots, j]$ and the suffix S_t is maximal. Consider the suffix S_k , which is an extension of $S[k \dots j]$. Clearly, it is sufficient to find the suffix S_t for which $|\text{LCP}(S_k, S_t)|$ is maximized (without necessarily computing the value $|\text{LCP}(S_k, S_t)|$ at this stage). Therefore we have two steps: (1) finding the location t , and (2) computing $|\text{LCP}(S_k, S_t)|$. Step (2) is easy since we assume that we have a full LCP data structure as described in Section 3. So, our goal is to solve Step (1). To do so we generalize our problem to the following.

Interval Longest Common Prefix (ILCP(k, l, r)): given k, l, r , we look for location $l \leq t \leq r$ of S for which the suffix S_t has the longest common prefix with S_k .

Clearly, for us it is sufficient to compute $\text{ILCP}(k, i, k - 1)$.

To compute the Interval Longest Common Prefix (ILCP(k, l, r)) we use a reduction to the problem of *3-sided range successor query*. That is given a 2D rank-space input on an $n \times n$ grid, a 3-sided query $R = [a, b] \times [-\infty, c]$ seeks the point in R with the largest y -coordinate. The 3-sided range successor query problem was considered under a different guise in [38] where it was called the *range next value* problem. There it was considered as an array problem for which one desires to preprocess the array to allow queries that seek the largest value on a range less than a value v . This can be translated to a grid and vice versa.

7.2 4-Sided and 3-Sided Sorted Range Reporting

The 3-sided range successor query generalizes quite nicely to the *3-sided sorted range reporting in 2D* which we define now.

3-Sided Sorted Range Reporting in 2D (rank space)	
Input:	A point set $P = \{(x_1, y_1), \dots, (x_n, y_n)\} \subseteq [1, n] \times [1, n]$.
Output:	A data structure representing P that supports the following <i>3-sided sorted range reporting</i> queries.
Query:	Given a range $R = [a, b] \times [-\infty, c]$ report all points of P contained in R sorted by their y -coordinate (from highest to lowest).

A solution for the 3-sided range successor query problem was proposed in Lenhof and Smid [83]⁸, and modified in [76] (improved query times) to work in rank space, i.e. on an $[n] \times [n]$ grid for n values with queries supported in $O(\log \log n)$ worst-case time, using $O(n \log n)$ space. However, there are now better results which solve, not only the 3-sided range successor problem, but the more general 3-sided sorted range reporting. The 3-sided sorted range reporting generalizes the 3-sided range successor query because the solutions presented can report the first location and stop.

For the same reason the 3-sided solutions work just as well for the 4-sided sorted range reporting, as they can report all points until the y -coordinates surpasses the range boundary and then stop.

The current best range 3-sided sorted range reporting data structures in 2D are as follows:

1. **Navarro and Nekrich [95]:**
 - (a): a data structure with $O(n)$ space where queries can be answered in $O(\log n(1 + k))$ where k is the number of points reported and
 - (b): a data structure with $O(n \log \log n)$ space where queries can be answered in $O(\log \log n(1 + k))$ and
 - (c): a data structure with $O(n \log^\epsilon n)$ space for any constant $\epsilon > 0$, where queries can be answered in $O(\log \log n + k)$,
2. **Crochemore et al. [38]:** a data structure that requires $O(n^{1+\epsilon})$ space for any constant $\epsilon > 0$ and can answer queries in $O(k)$ time.

⁸ Note, they called it the *Range Searching for Minimum* problem.




Fig. 3. The grid depicts P , the geometric representation of $S = abaabaabaaba$. The light point at $(8, 10)$ represents the suffix S_8 (and 10-th in the suffix array), as the substring $baab$, yet to be encoded, starts at position 8 in the string S . The grayed area of the grid represents the part of the substring which has already been encoded, that is $S[4, 7]$. When finding the start location t , we will be limited to using points found in the gray area.

7.3 The Interval Longest Common Prefix to 3-Sided Sorted Range Reporting

The reduction works as follows. Let $SA(S)$ be the suffix array of our input string S of length n . We associate each suffix S_i with its string index i and with its lexicographic index $SA^{-1}(i)$. From these two we generate a pair (x_i, y_i) , where $x_i = i$ and $y_i = SA^{-1}(i)$. We then preprocess the set $P = \{(x_i, y_i) \mid 1 \leq i \leq n\} \subseteq [1, n] \times [1, n]$ for 3-sided sorted range reporting queries. An example of the geometric representation of the scenario can be seen in Figure 3.

Computation of the ILCP. Consider the suffix S_k and the set of suffixes $\Gamma = \{S_l, \dots, S_r\}$. Since $|\text{LCP}(S_k, S_t)| = \max_{t' \in [l, r]} |\text{LCP}(S_k, S_{t'})|$, S_t is in fact the suffix lexicographically closest to S_k , out of all the suffixes of the set Γ .

We will first assume that we are searching for a suffix S_{t_1} , such that the suffix S_{t_1} is *lexicographically smaller* than S_k . The process for the case where the suffix chosen is lexicographically greater than S_k is symmetric. Therefore, once both

are found all we will need to do is to choose the best of both, i.e., the option yielding the greater $|\text{LCP}(S_k, S_t)|$ value.

Since we have assumed w.l.o.g. that S_{t_1} is lexicographically smaller than S_k , we have actually assumed that $y_{t_1} < y_k$, or equivalently, that t_1 appears to the *left* of k in the suffix array. Incorporating the lexicographical ranks of S_k and S_{t_1} into the expression, t_1 is actually the value which maximizes the expression $\max\{y_{t_1} \mid l \leq t_1 \leq r \text{ and } y_{t_1} < y_k\}$. Notice that $t_1 = x_{t_1}$.

Now consider the set $P = \{(x_i, y_i) \mid 1 \leq i \leq n\}$. Assuming that indeed $y_{t_1} < y_k$, we are interested in finding the maximal value y_{t_1} , such that $y_{t_1} < y_k$, and $l \leq x_{t_1} \leq r$. It immediately follows that the point $(x_{t_1}, y_{t_1}) \in P$ is the point in the range $[l, r] \times [-\infty, y_k - 1]$ having the maximal y -coordinate, and therefore can be obtained efficiently by obtaining the largest y -coordinate in the output of the 3-sided sorted range query. Once we have found the point (x_{t_1}, y_{t_1}) , we have t_1 , as $x_{t_1} = t_1$.

Equivalently, there exists t_2 such that S_{t_2} is the suffix lexicographically larger than S_k and closest to it. In other words, we assume $y_{t_2} > y_k$, or equivalently, that t_2 appears to the *right* of k in the suffix array. t_2 can be found using a symmetric procedure. An example of the queries performed can be seen in Figure 4.

Determining whether $t = t_1$ or $t = t_2$ is implemented by calculating both $|\text{LCP}(S_k, S_{t_1})|$ and $|\text{LCP}(S_k, S_{t_2})|$, and choosing the larger of the two. This gives us phrase ρ_i . To finish simply reiterate. Hence,




Fig. 4. Continued example from Figure 3 with string $S = abaabaabaaba$. The grid on the right-hand side depicts the 3-sided range successor query for $[l, r] \times [-\infty, y_k - 1]$, where the grid on the left-hand side depicts the $[l, r] \times [y_k + 1, \infty]$ query. In both queries the values given for the example queries are: $k = 8, l = 4$ and $r = 7$. $S[l, r]$ is the substring that has already been encoded. The query $[l, r] \times [-\infty, y_k - 1]$ outputs (4, 7) and the query on $[l, r] \times [y_k + 1, \infty]$ outputs (5, 11). S_5 is chosen since $|\text{LCPS}_8, S_5| > |\text{LCPS}_8, S_4|$.

Theorem 7. *Let T be a text of length n . We can preprocess T in $O(s(n))$ space so that we can answer substring compression queries Q in $O(|LZ(Q)|qt(n))$ time, where $s(n)$ and $qt(n)$ are the space and query times mentioned above ($occ = 1$).*

The GSCQ Problem. The *generalized substring compression* solution is more involved and uses binary searches on suffix trees applying range queries (3-sided range successor queries and emptiness queries) during the binary search. The interested reader should see [75].

Other Applications. The reduction from this section to range searching structures, i.e. the set of $P = \{(x_i, y_i) \mid 1 \leq i \leq n\} \subseteq [1, n] \times [1, n]$, defined by a suffix S_i with $x_i = i$ and $y(i) = SA^{-1}(i)$ had been considered beforehand.

This reduction was first used, to the best of our knowledge, by Ferragina [49] as part of the scheme for searching in a dynamic text indexing scheme. The reduction and point set were used also in *position restricted substring search* [86]. However, in both *range reporting* was used.

To the best of our knowledge, the first use of 3-sided range successor queries for text indexing was for *range non-overlapping indexing and successive list indexing* [76] and in parallel for position restricted substring search in [39].

See Section 9 for more range-restricted string search problems.

8 Top- k Document Retrieval

The *Top k Document Retrieval* problem is an extension of the Document Retrieval problem described in Section 3.2. The extension is to find the *top k* documents in which a given pattern Q appears, under some relevance measure. Examples of such relevance measures are (a) $tf(Q, d)$, the number of times Q occurs in document d , (b) $mind(Q, d)$, the minimum distance between two occurrences of Q and d , (c) $docrank(d)$, an arbitrary static rank assigned to document d . In general, the type of relevance measures which we shall discuss here are those that are defined by a function that assigns a numeric weight $w(S, d)$ to every substring S in document d , such that $w(S, d)$ depends only on the set of starting positions of occurrences of S in d . We call such a relevance measure a *positions based relevance measure*.

The following theorem is the culmination of the work of Hon, Shah and Vitter [65] and of Navarro and Nekrich [94].

Theorem 8. *Let D be a collection of strings (documents) of total length n , and let $w(S, d)$ be a positions based relevance measure for the documents $d \in D$. Then there exists an $O(n)$ -word space data structure that, given a string Q and an integer k reports k documents d containing Q of highest relevance, i.e. with highest $w(Q, d)$ values, in decreasing order of $w(Q, d)$, in $O(|Q| + k)$ time.*

Hon, Shah and Vitter [65] reduced this to a problem on arrays and achieved query time of $O(|Q| + k \log k)$. We will outline their idea within this section. Navarro and Nekrich [94] then showed how to adapt their solution to a 3-sided 2 dimensional range searching problem on weighted points. The solution of the range searching problem given in [94] builds upon earlier work on *top k color queries for document retrieval* [74], another interesting result. We will describe the adaptation and range searching result shortly.

8.1 Flattening the Top- k Document Retrieval Suffix Tree

Consider a generalized suffix tree ST for the document collection D . The leaves have a one-one correspondence with the locations within the documents. If leaf l is associated with location i of document d , we say that it is a d -leaf. Let l_1, l_2, \dots, l_q be the set of d -leaves. Then a node is a d -node if (a) it is a d -leaf or if (b) it is an internal node v of ST such that it is the lowest common ancestor of adjacent d -leaves l_i and l_{i+1} . Let v be a d -node. If u is the lowest common ancestor of v that is a d -node we say that u is v 's d -parent (and that v is u 's d -child). If there is no lowest common ancestor of v which is a d -node then the d -parent will be a dummy node which is the parent of the root. One can easily verify that the set of d -nodes form a tree, called a d -tree, and that an internal d -node has at least two d -children. It is also straightforward to verify that the lowest common ancestor of any two d -nodes is a d -node. Hence,

Lemma 4. *Let v be a node in the generalized suffix tree ST for the document collection D . For every document d for which the subtree of v contains a d -leaf there is exactly one d -node in the subtree of v that has a d -parent to an ancestor of v .*

Proof. 1. Every d -parent of a d -node in the subtree of v is either in the subtree or is an ancestor of v . 2. Assume, by contradiction, that there are two d -nodes x_1 and x_2 in the subtree of v each with a d -parent that is an ancestor of v . However, their lowest common ancestor, which must be a d -node, is no higher than v itself (since v is a common ancestor). Hence, their d -parents must be in v 's subtree a contradiction.

Hence, since every d -node has a d -parent, there must be exactly one d -node with a d -parent to an ancestor of v . \square

Corollary 1. *Let v be an arbitrary node in ST and let u be a descendant of v such that u is a d -node and its d -parent is an ancestor of v . Then all d -nodes u' which are descendants of v are also descendants of u .*

A node in ST may be a d -node for different d 's, say for d_{i_1}, \dots, d_{i_r} . Nevertheless, since every internal d -node has at least two d -children, the d -tree is linear in the number of d -leaves and, hence, the collection of all d -trees is linear in the size of the ST which is $O(n)$.

In light of this in [65] an array A was constructed by a pre-order traversal of the tree ST such that for each node v which is a d -node for $d \in \{d_{i_1}, \dots, d_{i_r}\}$ indexes $j + 1$ to $j + r$ are allocated in the array and contain the d_{i_1} -parent of v , \dots , the d_{i_r} -parent of v . The integer interval $[l_v, r_v]$ denotes the interval bounded by the minimal and maximal indexes in A assigned to v or its descendants. Values l_v and r_v are stored in v .

The array A was used to obtain the query result of $O(|Q| + k \log k)$ in [65]. We now show how this was used in [94].

8.2 Solving with Weighted Range Searching

Let $j + t$ be the index in A associated with d_{i_t} -node v for $d_{i_t} \in \{d_{i_1}, \dots, d_{i_r}\}$ and with its d_{i_t} -parent u_t . Let S be the string such that the locus of S is v . We generate a point $(j + t, \text{depth}(u_t))$, where depth denotes the depth of a node in the ST. The weight of the point p is $w(S, d_{i_t})$. Note that all points have different x -coordinates and are on an integer $n \times n$ grid.

It is still necessary to store a mapping from the x -coordinates of points to the document numbers. A global array of size $O(n)$ is sufficient for this task.

Queries. To answer a top- k query Q first find the locus v of Q (in $O(|Q|)$ time). Now, by Lemma 4 for each document d containing Q there is a unique d -node u which is a descendant of v with a d -parent who is an ancestor of v . By Corollary 1 $w(Q, d) = w(S, d)$, where S is the string with locus u , and $w(S, d)$ is the weight of the point corresponding to the pointer from u to its d -parent. So, there is a unique point (x, y) with $x \in [l_v, r_v]$ and $y \in [0, \text{depth}(v) - 1]$ for every document d that contains Q . Therefore, it is sufficient to report the k heaviest weight nodes in $[l_v, r_v] \times [0, \text{depth}(v) - 1]$. To do so Navarro and Nekrich [94] proposed the *three sided top- k range searching* problem.

Three sided top- k range searching	
Input:	A set of n weighted points on an $n \times n$ grid G .
Output:	A data structure over G supporting the following queries.
Query:	Given $1 \leq k, h \leq n$ and $1 \leq a \leq b \leq n$ return the k heaviest weighted points in the range $[a, b] \times [0, h]$.

In [94] a solution was given that uses $O(n)$ -word space and $O(h+k)$ query time. This result is similar to that of [74]. It is easy to note that for the application of top- k document retrieval this yields an $O(\text{depth}(v) + k)$ query time, which is $O(|Q| + k)$, an optimal solution.

8.3 External Memory Top- k Document Retrieval

Lately, a new result for top- k document retrieval for the external memory model has appeared in [98]. The result is I/O optimal and uses $O(n \log^* n)$ space.

More on research in the vicinity of top- k document retrieval can be found in [92]. See also [17] to see how to add rank functionality to a suffix tree.

9 Range Restricted String Problems

Research inspired by the problem of applying string problems limited to ranges has been of interest in the pattern matching community from around 2005. Some of the results are general. Others focus on specific applications. One such application is the *substring compression* problem that was discussed in Section 7. These problems are natural candidates for range searching solutions and indeed many of them have been solved with these exact tools.

The first three results on range restricted variants of text indexing appeared almost in parallel. The results were for *property matching* (the conference version of [7]), *substring compression* [37] and *position-restricted substring searching* [86].

Property matching is the problem of generating a text index for a text and a collection of ranges over the text. The subsequent pattern queries asks for the locations where the text appears and are fully contained in some interval. The initial definition was motivated by *weighted matching*. In *weighted matching* a text is given with probabilities on each of the text symbols and each pattern occurrence in the text has *weight* which is the multiplication of the probabilities on the text symbols associated with that occurrence. Weighted matching was reduced to property matching. In [7] a solution was given using $O(n)$ space, where n is the text size, such that queries are answered in $O(|Q| + occ_\pi)$ time, where Q is the pattern query and occ_π is the number of appearances within the interval set π . The preprocessing time was near optimal and in a combination of a couple of papers was solved in optimal time [66,69]. See also [40]. In [77] property matching was solved in the dynamic case, where intervals can be inserted and removed. Formally, π denotes the collection of intervals and the operations are:

- $\text{Insert}(s, f)$ - Insert a new interval (s, f) into π .
- $\text{Delete}(s, f)$ - Delete the interval (s, f) from π .

In [77] it was shown how to maintain a data structure under interval deletions. Queries are answered in $O(Q + occ_\pi)$ time and deletions take $O(f - s)$ time. If both insertions and deletions are allowed then the insertion/deletion time is $O(f - s + \log \log n)$, where n is the text length.

In [63] a succinct version was given that uses a compressed suffix array (CSA). The solution has a multiplicative logarithmic penalty for the query and update time.

Position-restricted substring searching is the problem where the goal is to preprocess an index to allow range-restricted queries. That is the query consists of a pattern query Q and a range described by text indices i and j . This is different from property matching because the interval is not given a-priori. On the other hand, it is one interval only. The queries considered in [86] are *position-restricted reporting* and *position-restricted counting*. Another two related queries also considered are *substring rank* and *substring select*, which are natural extensions of rank and select [24, 57, 68]. These are defined as follows.

1. **PRI-Report:** Preprocess text $T = t_1 \cdots t_n$ to answer queries $\text{Report}(Q = q_1 \cdots q_m, i, j)$, which reports all occurrences of Q in $t_i \dots t_j$.
2. **PRI-Count:** Preprocess text $T = t_1 \cdots t_n$ to answer queries $\text{Count}(Q = q_1 \cdots q_m, i, j)$, which returns the number of occurrences of Q in $t_i \dots t_j$.
3. **Substring Rank:** Preprocess text $T = t_1 \cdots t_n$ to answer queries $\text{SSR}(Q = q_1 \cdots q_m, k)$, which returns the number of occurrences of Q in $t_1 \cdots t_k$.
4. **Substring Select:** Preprocess text $T = t_1 \cdots t_n$ to answer queries $\text{SSS}(Q, k)$, which returns the k^{th} occurrence of Q in T .

We note that substring rank and position-restricted counting reduce to each other. Also, position-restricted reporting can be obtained from applying one substring-rank and $\text{occ}_{i,j}$ substring-selects, where $\text{occ}_{i,j}$ is the number of pattern occurrences in $t_i \dots t_j$. We leave it to the reader to verify the details.

Reporting: For the reporting problem Mäkinen and Navarro [86] reduced the problem to range reporting. The way to do so is to first find the SA -range of Q . Then this range and the position-restricted range i to j define a rectangle for which range reporting is used. One can use any of the data structures mentioned in Section 4. For example with $O(n \log^\epsilon n)$ space, for any constant ϵ , one can answer queries in $O(m + \log \log n + \text{occ}_{i,j})$ (we assume the alphabet is from $[1, n]$ otherwise if it is from $[1, U]$ then there is an extra additive factor of $\log \log U$). Crochemore et al. [39] noticed that a different type of reduction, namely range next value, could be more useful. The authors of [86] were more concerned with space issues. So, they also proposed a data structure which uses $n + o(n)$ space and reports in $O(m + \log n + \text{occ}_{i,j})$. The reporting time was improved by Bose et al. [20] to $O(m + \log n / \log \log n + \text{occ}_{i,j})$ with the same space constraints. Yu et al. [106] suggested a different algorithm with the same space and reporting time, but were able to report the occurrences in their original order.

Bille and Gørtz [19] went on to show that with $O(n(\log^\epsilon n + \log \log U))$ space the query time can be improved to $O(m + \text{occ}_{i,j})$, which is optimal. They also solved position-restricted reporting merged with property matching. The space and query time remain the same.

An interesting result for the reporting variant appeared in [64]. Specifically, it was shown that a succinct space polylog n time index for position-restricted substring searching is at least as hard as designing a linear space data structure for 3D range reporting in polylog n time.

Counting: In [86] the same data structure that uses $n + o(n)$ space and reports in $O(m + \log n)$ was used for counting⁹. Once again, Bose et al. [20] can improve the counting time to $O(m + \log n / \log \log n)$. Kopelowitz et al. [80] presented a counting data structure that uses $O(n(\log n / \log \log n))$ space and answers counting queries in time $O(m + \log \log |\Sigma|)$. Recently, in the upcoming journal version of [19] a similar result appears. The space is the same. However, the counting time is $O(m + \log \log n)$, which can be slightly worse.

⁹ There is another result there that assumes faster query times that is flawed. See the introduction in [80] for an explanation.

Substring Select: In [86] a solution for indexing for substring select is given. The space is $O(nK \log \sigma / \log n)$, where K is an upper bound on the size of the queried patterns. The query time is $O(m \log \sigma / \log \log n)$. This was improved in [80] to allow for any length query with $O(n \log n / \log \log n)$ space and optimal $O(m)$ query time. The proposed solution in [80] uses persistent data structure which is a basic ingredient in most of the range searching solutions.

Substring compression has been expanded on in Section 7. It was introduced in [37] and improved upon in [75]. The results are detailed in Section 7. One of the problems that is of interest in substring compression is the *ILCP* (interval longest common prefix) query (see Section 7). This inspired Amir et al. [3] to consider extensions to LCP range queries of different types.

Range Non-overlapping Indexing and Successive List Indexing [76].

In range non-overlapping indexing one wants to prepare an index so that when give a pattern query one can return a maximal set of occurrences so that the occurrences do not overlap. In successive list indexing one prepares an index to answer queries where a pattern is given along with a position i and one desires to find the first occurrence of the pattern after i . A reduction to range successor was used to solve this problem. Along with the results of sorted range reporting [95] one can solve the former with space $O(n \log^\epsilon n)$ space and $O(\log \log n + occ)$ query time. For the latter the query time is $O(\log \log n)$.

Range Successor in 2D solves several of the problems mentioned in this section. This has been discussed in Section 7 and is referred to in [38, 95, 106]. It is interesting that its generalization sorted range reporting [95] is a variant of range reporting that was considered in the community because of the unique range search problems that arise.

10 Lower Bounds on Text Indexing via Range Reporting

A novel use of range searching is its use to show lower bounds on text indexing via reductions from range reporting [32].

Theorem 9. *Let $S = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ be a set of n points in $[1, n] \times [1, n]$. We can construct a text T of length $O(n \log n)$ bits along with $O(n \log n)$ bits of auxiliary data such that we can answer range reporting queries on S with $O(\log^2 n)$ pattern match queries on T , each query is a pattern of length $O(\log n)$.*

We denote the set of x -coordinates, of S , $X = \{x_1, \dots, x_n\}$ and the set of y -coordinates, of S , $Y = \{y_1, \dots, y_n\}$.

The idea is as follows. Each point $(x, y) \in [1, n] \times [1, n]$. So, x and y both have binary representations of $\log n$ bits¹⁰. Denote with $b(w)$ the

¹⁰ We assume that n is a power of 2. Otherwise, it will be $\lfloor \log n \rfloor + 1$.

binary representation of a number w and the reverse of the binary representation with $b^R(w)$. The text T (from the theorem) constructed is $b^R(x_1)\#b(y_1)\$b^R(x_2)\#b(y_2)\$\dots\$b^R(x_n)\#b(y_n)$.

To obtain the result a collection of pattern queries on T is generated whose answers will yield an answer to the range reporting problem on the point set S . To this end, sort $b(y_1), \dots, b(y_n)$ and $b^R(x_1), \dots, b^R(x_n)$. Let π denote the ordering of the former and τ denote the ordering of the latter, i.e. $b(\pi(y_1)) < \dots < b(\pi(y_n))$, where $<$ is a lexicographic-less than, and $b^R(\tau(x_1)) < \dots < b^R(\tau(x_n))$. The n -length arrays $A = \langle b^R(\tau(x_1)), \dots, b^R(\tau(x_n)) \rangle$ and $B = \langle b(\pi(y_1)), \dots, b(\pi(y_n)) \rangle$ will be the basis of the search.

Over each of the arrays construct a binary search tree with each node representing a range of elements. Without loss of generality, consider the binary tree over B . The root represents all elements of Y . The left son is associated with one bit 0 and represents $R_0 = \{y \in Y \mid 0 \text{ is prefix of } b(y)\}$ and the right son represents $R_1 = \{y \in Y \mid 1 \text{ is prefix of } b(y)\}$ - each is a range over B - check. The left son of the left son of the root represents $R_{00} = \{y \in Y \mid 00 \text{ is prefix of } b(y)\}$, etc. In general, each node is associated with a binary string, say $b_0 \dots b_i$, formed by the walk down from the root to the node and is also associated with a range, which we call a *node-range*, $R_{b_0 \dots b_i} = \{y \in Y \mid b_0 \dots b_i \text{ is prefix of } b(y)\}$. The number of nodes in the binary tree and, hence, the number of ranges is $\leq 2n - 1$. Each range can be represented as a pair of indexes to the array. Hence, the size of the auxiliary information is $O(n)$ -words, or $O(n \log n)$ bits. We construct a complementary binary tree for A , with ranges RX .

An easy well known observation is that any range $R^{q,r} = \{y \mid 1 \leq q \leq y \leq r \leq n\}$ can be expressed as the disjoint union of at most $2 \log n$ node-ranges. The node-ranges of the disjoint union can be found by a traversal up and down the binary tree using the binary representations of q and r .

Now consider a range query on S , say $[x_{left}, x_{right}] \times [y_{bottom}, y_{top}]$. This can be seen as a query for all (x, y) such that $y \in R^{y_{bottom}, y_{top}}$ and $x \in RX^{x_{left}, x_{right}}$. By the previous observation this can be transformed into $O(\log^2 n)$ queries for all (x, y) such that x is in one of the node-ranges in the disjoint union expressing $R^{y_{bottom}, y_{top}}$ and y is in one of the node-ranges in the disjoint union expressing $RX^{x_{left}, x_{right}}$.

We show an indexing query that searches for all (x, y) such that $x \in RX_c$ and $y \in R_d$ both node-ranges, the former for the binary string c over the array A and the latter for the binary string d over the array B . We define a pattern query $Q = c^R \# d$. We query the text index with Q . Every location where Q appears corresponds to a point that is in the desired range as c^R must align with the end of an $b^R(x_i)$, which is the same as c being a prefix of $b(x_i)$ and d is a prefix of (y_i) , which is exactly the desired.

Chazelle [30] showed that in the pointer machine model an index supporting 2D range reporting in $O(\text{polylog}(n) + occ)$ query time, where occ is the number of occurrences, requires $\Omega(n(\log n / \log \log n))$ words of storage. Hence,

Theorem 10. *In the pointer machine model a text index on T of size n which returns locations of pattern occurrences in $O(\text{polylog}(n) + \text{occ})$ time requires $\Omega(n(\log n / \log \log n))$ bits.*

More on This Result and Related Work. In [32] there are also very interesting results reducing text indexing to range searching. The reduction is known as a Geometric BWT, transforming a BWT into a point representation. The reductions in both directions show that obtaining improvements in space complexity of either will imply space complexity improvements on the other.

Another two lower bounds, along the lines of the lower bound described here, are for *position restricted substring search* [64], for *forbidden patterns* [52] and for *aligned pattern matching* [99].

Acknowledgement. I wanted to thank my numerous colleagues who were kind enough to provide insightful comments on an earlier version and pointers to work that I was unaware of. These people include (in alphabetical order) Phillip Bille, Timothy Chan, Francisco Claude, Pooya Davoodi, Johannes Fischer, Travis Gagie, Roberto Grossi, Orgad Keller, Tsvi Kopelowitz, Muthu Muthukrishnan, Gonzalo Navarro, Yakov Nekrich, Rahul Shah, Sharma Thankachan, Rajeev Raman, and Oren Weimann. Special thanks to Orgad, Rahul, Sharma and Yakov for numerous Skype conversations in which I learned more than can be contained within this monologue.

References

1. Agarwal, P.K.: Range searching. In: Handbook of Discrete and Computational Geometry, pp. 575–598. CRC Press, Inc. (1997)
2. Alstrup, S., Brodal, G.S., Rauhe, T.: New data structures for orthogonal range searching. In: Proc. of Foundations of Computer Science (FOCS), pp. 198–207 (2000)
3. Amir, A., Apostolico, A., Landau, G.M., Levy, A., Lewenstein, M., Porat, E.: Range LCP. In: Asano, T., Nakano, S.-i., Okamoto, Y., Watanabe, O. (eds.) ISAAC 2011. LNCS, vol. 7074, pp. 683–692. Springer, Heidelberg (2011)
4. Amir, A., Apostolico, A., Landau, G.M., Satta, G.: Efficient text fingerprinting via Parikh mapping. Journal of Discrete Algorithms 1(5-6), 409–421 (2003)
5. Amir, A., Aumann, Y., Lewenstein, M., Porat, E.: Function matching. SIAM Journal on Computing 35(5), 1007–1022 (2006)
6. Amir, A., Benson, G., Farach, M.: Let sleeping files lie: Pattern matching in z-compressed files. Journal of Computer and System Sciences 52(2), 299–307 (1996)
7. Amir, A., Chencinski, E., Iliopoulos, C.S., Kopelowitz, T., Zhang, H.: Property matching and weighted matching. Theoretical Computer Science 395(2-3), 298–310 (2008)
8. Amir, A., Fischer, J., Lewenstein, M.: Two-dimensional range minimum queries. In: Ma, B., Zhang, K. (eds.) CPM 2007. LNCS, vol. 4580, pp. 286–294. Springer, Heidelberg (2007)

9. Amir, A., Francheschini, G., Grossi, R., Kopelowitz, T., Lewenstein, M., Lewenstein, N.: Managing unbounded-length keys in comparison-driven data structures with applications to on-line indexing. *SIAM Journal on Computing* (to appear, 2013)
10. Amir, A., Keselman, D., Landau, G.M., Lewenstein, M., Lewenstein, N., Rodeh, M.: Text indexing and dictionary matching with one error. *Journal of Algorithms* 37(2), 309–325 (2000)
11. Amir, A., Landau, G.M., Lewenstein, M., Sokol, D.: Dynamic text and static pattern matching. *ACM Transactions on Algorithms* 3(2) (2007)
12. Arroyuelo, D., Navarro, G., Sadakane, K.: Stronger Lempel-Ziv based compressed text indexing. *Algorithmica* 62(1-2), 54–101 (2012)
13. Atallah, M.J., Yuan, H.: Data structures for range minimum queries in multidimensional arrays. In: *Proc. of the Symposium on Discrete Algorithms (SODA)*, pp. 150–160 (2010)
14. Badkobeh, G., Fici, G., Kroon, S., Lipták, Z.: Binary jumbled string matching for highly run-length compressible texts. *Information Processing Letters* 113, 604–608 (2013)
15. Baker, B.S.: Parameterized pattern matching: Algorithms and applications. *Journal of Computer and System Sciences* 52(1), 28–42 (1996)
16. Barbay, J., Claude, F., Navarro, G.: Compact binary relation representations with rich functionality. *The Computing Research Repository (arXiv)*, abs/1201.3602 (2012)
17. Bialynicka-Birula, I., Grossi, R.: Rank-sensitive data structures. In: *Consens, M.P., Navarro, G. (eds.) SPIRE 2005. LNCS, vol. 3772*, pp. 79–90. Springer, Heidelberg (2005)
18. Bille, P., Fischer, J., Gørtz, I.L., Kopelowitz, T., Sach, B., Vildhøj, H.W.: Sparse suffix tree construction in small space. In: *Proc. of International Colloquium on Automata, Languages and Complexity, ICALP* (2013)
19. Bille, P., Gørtz, I.L.: Substring range reporting. In: *Giancarlo, R., Manzini, G. (eds.) CPM 2011. LNCS, vol. 6661*, pp. 299–308. Springer, Heidelberg (2011)
20. Bose, P., He, M., Maheshwari, A., Morin, P.: Succinct orthogonal range search structures on a grid with applications to text indexing. In: *Dehne, F., Gavrilova, M., Sack, J.-R., Tóth, C.D. (eds.) WADS 2009. LNCS, vol. 5664*, pp. 98–109. Springer, Heidelberg (2009)
21. Brodal, G.S., Gašieniec, L.: Approximate dictionary queries. In: *Hirschberg, D.S., Meyers, G. (eds.) CPM 1996. LNCS, vol. 1075*, pp. 65–74. Springer, Heidelberg (1996)
22. Brodal, G.S., Davoodi, P., Lewenstein, M., Raman, R., Rao, S.S.: Two dimensional range minimum queries and Fibonacci lattices. In: *Epstein, L., Ferragina, P. (eds.) ESA 2012. LNCS, vol. 7501*, pp. 217–228. Springer, Heidelberg (2012)
23. Brodal, G.S., Davoodi, P., Rao, S.S.: On space efficient two dimensional range minimum data structures. *Algorithmica* 63(4), 815–830 (2012)
24. Brodnik, A., Munro, J.I.: Membership in constant time and almost-minimum space. *SIAM Journal on Computing* 28(5), 1627–1640 (1999)
25. Butman, A., Eres, R., Landau, G.M.: Scaled and permuted string matching. *Information Processing Letters* 92(6), 293–297 (2004)
26. Chan, H.-L., Lam, T.W., Sung, W.-K., Tam, S.-L., Wong, S.-S.: A linear size index for approximate pattern matching. *Journal of Discrete Algorithms* 9(4), 358–364 (2011)
27. Chan, T.M.: Persistent predecessor search and orthogonal point location on the word ram. In: *Proc. of Symposium on Discrete Algorithms (SODA)*, pp. 1131–1145 (2011)

28. Chan, T.M., Larsen, K.G., Pătrașcu, M.: Orthogonal range searching on the RAM, revisited. In: Proc. of the Symposium on Computational Geometry, SOCG (2011)
29. Chazelle, B.: A functional approach to data structures and its use in multidimensional searching. *SIAM Journal on Computing* 17(3), 427–462 (1988)
30. Chazelle, B.: Lower bounds for orthogonal range searching: I. the reporting case. *Journal of the ACM* 37(2), 200–212 (1990)
31. Chazelle, B., Rosenberg, B.: The complexity of computing partial sums off-line. *International Journal of Computational Geometry and Applications* 1(1), 33–45 (1991)
32. Chien, Y.-F., Hon, W.-K., Shah, R., Thankachan, S.V., Vitter, J.S.: Geometric Burrows-Wheeler transform: Compressed text indexing via sparse suffixes and range searching. *Algorithmica* (to appear, 2013)
33. Cicalese, F., Fici, G., Lipták, Z.: Searching for jumbled patterns in strings. In: Holub, J., Zdárek, J. (eds.) *Proceedings of the Prague Stringology Conference (PSC)*, pp. 105–117 (2009)
34. Claude, F., Navarro, G.: Self-indexed grammar-based compression. *Fundamenta Informaticae* 111(3), 313–337 (2011)
35. Claude, F., Navarro, G.: Improved grammar-based compressed indexes. In: Calderón-Benavides, L., González-Caro, C., Chávez, E., Ziviani, N. (eds.) *SPIRE 2012. LNCS*, vol. 7608, pp. 180–192. Springer, Heidelberg (2012)
36. Cole, R., Gottlieb, L.-A., Lewenstein, M.: Dictionary matching and indexing with errors and don't cares. In: Proc. of Symposium on Theory of Computing (STOC), pp. 91–100 (2004)
37. Cormode, G., Muthukrishnan, S.: Substring compression problems. In: Proc. of Symposium on Discrete Algorithms (SODA), pp. 321–330 (2005)
38. Crochemore, M., Iliopoulos, C.S., Kubica, M., Rahman, M.S., Tischler, G., Walen, T.: Improved algorithms for the range next value problem and applications. *Theoretical Computer Science* 434, 23–34 (2012)
39. Crochemore, M., Iliopoulos, C.S., Rahman, M.S.: Finding patterns in given intervals. In: Kučera, L., Kučera, A. (eds.) *MFCS 2007. LNCS*, vol. 4708, pp. 645–656. Springer, Heidelberg (2007)
40. Crochemore, M., Kubica, M., Walen, T., Iliopoulos, C.S., Rahman, M.S.: Finding patterns in given intervals. *Fundamenta Informaticae* 101(3), 173–186 (2010)
41. Davoodi, P., Landau, G., Lewenstein, M.: Multi-dimensional range minimum queries (manuscript, 2013)
42. Davoodi, P., Raman, R., Satti, S.R.: Succinct representations of binary trees for range minimum queries. In: Gudmundsson, J., Mestre, J., Viglas, T. (eds.) *COCOON 2012. LNCS*, vol. 7434, pp. 396–407. Springer, Heidelberg (2012)
43. Demaine, E.D., Landau, G.M., Weimann, O.: On cartesian trees and range minimum queries. In: Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., Thomas, W. (eds.) *ICALP 2009, Part I. LNCS*, vol. 5555, pp. 341–353. Springer, Heidelberg (2009)
44. Demaine, E.D., López-Ortiz, A.: A linear lower bound on index size for text retrieval. *Journal of Algorithms* 48(1), 2–15 (2003)
45. Dietz, P.F., Raman, R.: Persistence, amortization and randomization. In: Proc. of Symposium on Discrete Algorithms (SODA), pp. 78–88 (1991)
46. Farach, M., Muthukrishnan, S.: Perfect hashing for strings: Formalization and algorithms. In: Hirschberg, D.S., Meyers, G. (eds.) *CPM 1996. LNCS*, vol. 1075, pp. 130–140. Springer, Heidelberg (1996)
47. Farach, M., Thorup, M.: String matching in Lempel-Ziv compressed strings. *Algorithmica* 20(4), 388–404 (1998)

48. Farach-Colton, M., Ferragina, P., Muthukrishnan, S.: On the sorting-complexity of suffix tree construction. *Journal of the ACM* 47(6), 987–1011 (2000)
49. Ferragina, P.: Dynamic text indexing under string updates. *Journal of Algorithms* 22(2), 296–328 (1997)
50. Ferragina, P., Manzini, G.: Indexing compressed text. *Journal of the ACM* 52(4), 552–581 (2005)
51. Ferragina, P., Muthukrishnan, S., de Berg, M.: Multi-method dispatching: A geometric approach with applications to string matching problems. In: *Proc. of Symposium on Theory of Computing (STOC)*, pp. 483–491 (1999)
52. Fischer, J., Gagie, T., Kopelowitz, T., Lewenstein, M., Mäkinen, V., Salmela, L., Välimäki, N.: Forbidden patterns. In: Fernández-Baca, D. (ed.) *LATIN 2012*. LNCS, vol. 7256, pp. 327–337. Springer, Heidelberg (2012)
53. Fischer, J., Heun, V.: Space-efficient preprocessing schemes for range minimum queries on static arrays. *SIAM Journal on Computing* 40(2), 465–492 (2011)
54. Gabow, H.N., Bentley, J.L., Tarjan, R.E.: Scaling and related techniques for geometry problems. In: *Proc. of the Symposium on Theory of Computing (STOC)*, pp. 135–143 (1984)
55. Gagie, T., Gawrychowski, P., Kärkkäinen, J., Nekrich, Y., Puglisi, S.J.: A faster grammar-based self-index. In: Dediu, A.-H., Martín-Vide, C. (eds.) *LATA 2012*. LNCS, vol. 7183, pp. 240–251. Springer, Heidelberg (2012)
56. Golin, M., Iacono, J., Krizanc, D., Raman, R., Rao, S.S.: Encoding 2D range maximum queries. In: Asano, T., Nakano, S.-I., Okamoto, Y., Watanabe, O. (eds.) *ISAAC 2011*. LNCS, vol. 7074, pp. 180–189. Springer, Heidelberg (2011)
57. Golynski, A., Munro, J.I., Rao, S.S.: Rank/select operations on large alphabets: a tool for text indexing. In: *Proc. of Symposium on Discrete Algorithms (SODA)*, pp. 368–373 (2006)
58. Grossi, R., Gupta, A., Vitter, J.S.: High-order entropy-compressed text indexes. In: *Proc. of Symposium on Discrete Algorithms (SODA)*, pp. 841–850 (2003)
59. Grossi, R., Vitter, J.S.: Compressed suffix arrays and suffix trees with applications to text indexing and string matching. *SIAM Journal on Computing* 35(2), 378–407 (2005)
60. Gusfield, D.: *Algorithms on strings, trees, and sequences: computer science and computational biology*. Cambridge University Press (1997)
61. Harel, D., Tarjan, R.E.: Fast algorithms for finding nearest common ancestors. *SIAM Journal on Computing* 13(2), 338–355 (1984)
62. Hon, W.-K., Ku, T.-H., Shah, R., Thankachan, S.V., Vitter, J.S.: Compressed dictionary matching with one error. In: *Proc. of the Data Compression Conference (DCC)*, pp. 113–122 (2011)
63. Hon, W.-K., Patil, M., Shah, R., Thankachan, S.V.: Compressed property suffix trees. In: *Proc. of the Data Compression Conference (DCC)*, pp. 123–132 (2011)
64. Hon, W.-K., Shah, R., Thankachan, S.V., Vitter, J.S.: On position restricted substring searching in succinct space. *Journal of Discrete Algorithms* 17, 109–114 (2012)
65. Hon, W.-K., Shah, R., Vitter, J.S.: Space-efficient framework for top-k string retrieval problems. In: *Proc. of Foundations of Computer Science (FOCS)*, pp. 713–722 (2009)
66. Iliopoulos, C.S., Rahman, M.S.: Faster index for property matching. *Information Processing Letters* 105(6), 218–223 (2008)
67. Iliopoulos, C.S., Rahman, M.S.: Indexing factors with gaps. *Algorithmica* 55(1), 60–70 (2009)

68. Jacobson, G.: Space-efficient static trees and graphs. In: FOCS, pp. 549–554 (1989)
69. Juan, M.T., Liu, J.J., Wang, Y.L.: Errata for “faster index for property matching“. Information Processing Letters 109(18), 1027–1029 (2009)
70. Kärkkäinen, J.: Repetition-Based Text Indexes. PhD thesis, University of Helsinki, Finland (1999)
71. Kärkkäinen, J., Sanders, P., Burkhardt, S.: Linear work suffix array construction. Journal of the ACM 53(6), 918–936 (2006)
72. Kärkkäinen, J., Ukkonen, E.: Lempel-Ziv parsing and sublinear-size index structures for string matching. In: Proc. 3rd South American Workshop on String Processing (WSP). International Informatics Series 4, pp. 141–155. Carleton University Press (1996)
73. Kärkkäinen, J., Ukkonen, E.: Sparse suffix trees. In: Cai, J.-Y., Wong, C.K. (eds.) COCOON 1996. LNCS, vol. 1090, pp. 219–230. Springer, Heidelberg (1996)
74. Karpinski, M., Nekrich, Y.: Top-k color queries for document retrieval. In: Proc. of Symposium on Discrete Algorithms (SODA), pp. 401–411 (2011)
75. Keller, O., Kopelowitz, T., Landau, S., Lewenstein, M.: Generalized substring compression. In: Kucherov, G., Ukkonen, E. (eds.) CPM 2009 Lille. LNCS, vol. 5577, pp. 26–38. Springer, Heidelberg (2009)
76. Keller, O., Kopelowitz, T., Lewenstein, M.: Range non-overlapping indexing and successive list indexing. In: Dehne, F., Sack, J.-R., Zeh, N. (eds.) WADS 2007. LNCS, vol. 4619, pp. 625–636. Springer, Heidelberg (2007)
77. Kopelowitz, T.: The property suffix tree with dynamic properties. In: Amir, A., Parida, L. (eds.) CPM 2010. LNCS, vol. 6129, pp. 63–75. Springer, Heidelberg (2010)
78. Kopelowitz, T., Kucherov, G., Nekrich, Y., Starikovskaya, T.A.: Cross-document pattern matching. Journal of Discrete Algorithms (to appear, 2013)
79. Kopelowitz, T., Lewenstein, M.: Dynamic weighted ancestors. In: Proc. of Symposium on Discrete Algorithms (SODA), pp. 565–574 (2007)
80. Kopelowitz, T., Lewenstein, M., Porat, E.: Persistency in suffix trees with applications to string interval problems. In: Grossi, R., Sebastiani, F., Silvestri, F. (eds.) SPIRE 2011. LNCS, vol. 7024, pp. 67–80. Springer, Heidelberg (2011)
81. Kreft, S., Navarro, G.: On compressing and indexing repetitive sequences. Theoretical Computer Science 483, 115–133 (2013)
82. Landau, G.M., Vishkin, U.: Fast string matching with k differences. Journal of Computer and System Sciences 37(1), 63–78 (1988)
83. Lenhof, H.-P., Smid, M.H.M.: Using persistent data structures for adding range restrictions to searching problems. Theoretical Informatics and Applications (ITA) 28(1), 25–49 (1994)
84. Levenshtein, V.I.: Binary codes capable of correcting deletions, insertions, and reversals. Soviet Physics Doklady 10, 707–710 (1966)
85. Lewenstein, M.: Parameterized matching. In: Encyclopedia of Algorithms (2008)
86. Mäkinen, V., Navarro, G.: Position-restricted substring searching. In: Correa, J.R., Hevia, A., Kiwi, M. (eds.) LATIN 2006. LNCS, vol. 3887, pp. 703–714. Springer, Heidelberg (2006)
87. Manber, U., Myers, E.W.: Suffix arrays: A new method for on-line string searches. SIAM Journal on Computing 22(5), 935–948 (1993)
88. McCreight, E.M.: A space-economical suffix tree construction algorithm. Journal of the ACM 23(2), 262–272 (1976)
89. Moosa, T.M., Rahman, M.S.: Indexing permutations for binary strings. Information Processing Letters 110(18-19), 795–798 (2010)

90. Muthukrishnan, S.: Efficient algorithms for document retrieval problems. In: Proc. of the Symposium on Discrete Algorithms (SODA), pp. 657–666 (2002)
91. Navarro, G.: Wavelet trees for all. In: Kärkkäinen, J., Stoye, J. (eds.) CPM 2012. LNCS, vol. 7354, pp. 2–26. Springer, Heidelberg (2012)
92. Navarro, G.: Spaces, trees and colors: The algorithmic landscape of document retrieval on sequences. The Computing Research Repository (arXiv), abs/1304.6023 (2013)
93. Navarro, G., Mäkinen, V.: Compressed full-text indexes. ACM Computing Surveys 39(1), 2 (2007)
94. Navarro, G., Nekrich, Y.: Top- k document retrieval in optimal time and linear space. In: Proc. of Symposium on Discrete Algorithms (SODA), pp. 1066–1077 (2012)
95. Nekrich, Y., Navarro, G.: Sorted range reporting. In: Fomin, F.V., Kaski, P. (eds.) SWAT 2012. LNCS, vol. 7357, pp. 271–282. Springer, Heidelberg (2012)
96. Russo, L.M.S., Oliveira, A.L.: A compressed self-index using a Ziv-Lempel dictionary. Information Retrieval 11(4), 359–388 (2008)
97. Sadakane, K.: Succinct data structures for flexible text retrieval systems. Journal of Discrete Algorithms 5(1), 12–22 (2007)
98. Shah, R., Sheng, C., Thankachan, S.V., Vitter, J.S.: On optimal top- k string retrieval. The Computing Research Repository (arXiv), abs/1207.2632 (2012)
99. Thankachan, S.V.: Compressed indexes for aligned pattern matching. In: Grossi, R., Sebastiani, F., Silvestri, F. (eds.) SPIRE 2011. LNCS, vol. 7024, pp. 410–419. Springer, Heidelberg (2011)
100. Tsur, D.: Fast index for approximate string matching. Journal of Discrete Algorithms 8(4), 339–345 (2010)
101. Ukkonen, E.: On-line construction of suffix trees. Algorithmica 14(3), 249–260 (1995)
102. van Emde Boas, P.: Preserving order in a forest in less than logarithmic time and linear space. Information Processing Letters 6(3), 80–82 (1977)
103. Vuillemin, J.: A unifying look at data structures. Communications of the ACM 23(4), 229–239 (1980)
104. Weiner, P.: Linear pattern matching algorithm. In: Proc. of the Symposium on Switching and Automata Theory, pp. 1–11 (1973)
105. Willard, D.E.: Log-logarithmic worst-case range queries are possible in space $\theta(n)$. Information Processing Letters 17(2), 81–84 (1983)
106. Yu, C.-C., Hon, W.-K., Wang, B.-F.: Improved data structures for the orthogonal range successor problem. Computational Geometry 44(3), 148–159 (2011)
107. Ziv, J., Lempel, A.: A universal algorithm for sequential data compression. IEEE Transactions on Information Theory 23(3), 337–343 (1977)