



## String indexing for top-k close consecutive occurrences

**Bille, Philip; Gørtz, Inge Li; Pedersen, Max Rishøj; Rotenberg, Eva; Steiner, Teresa Anna**

*Published in:*  
Theoretical Computer Science

*Link to article, DOI:*  
[10.1016/j.tcs.2022.06.004](https://doi.org/10.1016/j.tcs.2022.06.004)

*Publication date:*  
2022

*Document Version*  
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

*Citation (APA):*  
Bille, P., Gørtz, I. L., Pedersen, M. R., Rotenberg, E., & Steiner, T. A. (2022). String indexing for top-k close consecutive occurrences. *Theoretical Computer Science*, 927, 133-147.  
<https://doi.org/10.1016/j.tcs.2022.06.004>

---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



# String indexing for top- $k$ close consecutive occurrences <sup>☆</sup>

Philip Bille, Inge Li Gørtz, Max Rishøj Pedersen, Eva Rotenberg,  
Teresa Anna Steiner\*

DTU Compute, Richard Petersens Plads, Bygning 324, 2800 Kgs. Lyngby, Denmark



## ARTICLE INFO

### Article history:

Received 4 February 2021  
Received in revised form 6 May 2022  
Accepted 1 June 2022  
Available online 3 June 2022  
Communicated by M. Sciortino

### Keywords:

String indexing  
Pattern matching  
Consecutive occurrences

## ABSTRACT

The classic string indexing problem is to preprocess a string  $S$  into a compact data structure that supports efficient subsequent pattern matching queries, that is, given a pattern string  $P$ , report all occurrences of  $P$  within  $S$ . In this paper, we study a basic and natural extension of string indexing called the string indexing for top- $k$  close consecutive occurrences problem (SITCCO). Here, a consecutive occurrence is a pair  $(i, j)$ ,  $i < j$ , such that  $P$  occurs at positions  $i$  and  $j$  in  $S$  and there is no occurrence of  $P$  between  $i$  and  $j$ , and their distance is defined as  $j - i$ . Given a pattern  $P$  and a parameter  $k$ , the goal is to report the top- $k$  consecutive occurrences of  $P$  in  $S$  of minimal distance. The challenge is to compactly represent  $S$  while supporting queries in time close to the length of  $P$  and  $k$ . We give three time-space trade-offs for the problem. Let  $n$  be the length of  $S$ ,  $m$  the length of  $P$ , and  $\epsilon \in (0, 1]$ . Our first result achieves  $O(n \log n)$  space and optimal query time of  $O(m + k)$ . Our second and third results achieve linear space and query times either  $O(m + k^{1+\epsilon})$  or  $O(m + k \log^{1+\epsilon} n)$ . Along the way, we develop several techniques of independent interest, including a new translation of the problem into a line segment intersection problem and a new recursive clustering technique for trees.

© 2022 The Author(s). Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

## 1. Introduction

The classic string indexing problem is to preprocess a string  $S$  into a compact data structure that supports efficient subsequent pattern matching queries, that is, given a pattern string  $P$ , report all occurrences of  $P$  within  $S$ . An occurrence of  $P$  within  $S$  is an index  $i$ ,  $0 \leq i < |S|$ , such that  $P = S[i \dots i + |P| - 1]$ . In this paper, we introduce a basic extension of string indexing, where the goal is to report *consecutive occurrences* of the pattern  $P$  that occur *close* to each other in  $S$ . Here, a consecutive occurrence is a pair  $(i, j)$ ,  $i < j$ , such that  $P$  occurs at positions  $i$  and  $j$  in  $S$  and there is no occurrence of  $P$  between  $i$  and  $j$ , and close to each other means that the distance  $j - i$  between the occurrences should be small. More precisely, given a pattern  $P$  and an integer parameter  $k > 0$ , define the *top- $k$  close consecutive occurrences* of  $P$  to be the  $k$  consecutive occurrences of  $P$  in  $S$  with the smallest distances. Given a string  $S$  the *string indexing for top- $k$  close consecutive occurrences* (SITCCO) problem is to preprocess  $S$  into a data structure that supports top- $k$  close consecutive occurrences queries. The goal is to obtain a compact data structure while supporting fast queries in terms of the length of the pattern  $P$  and the number of reported occurrences  $k$ . For an example, see Fig. 1.

<sup>☆</sup> An extended abstract appeared at FSTTCS 2020 [8].

\* Corresponding author.

E-mail addresses: [phbi@dtu.dk](mailto:phbi@dtu.dk) (P. Bille), [inge@dtu.dk](mailto:inge@dtu.dk) (I.L. Gørtz), [mhrpe@dtu.dk](mailto:mhrpe@dtu.dk) (M.R. Pedersen), [erot@dtu.dk](mailto:erot@dtu.dk) (E. Rotenberg), [terst@dtu.dk](mailto:terst@dtu.dk) (T.A. Steiner).

$P = AN$   
 $S = \text{BATMAN AND ANNA SING NANANANA AND EAT BANANAS}$   
0 5 10 15 20 25 30 35 40

**Fig. 1.**  $P$  occurs at positions 4, 7, 11, 22, 24, 26, 30, 39 and 41 in  $S$ . The top 5 close consecutive occurrences are (22, 24), (24, 26), (39, 41), (4, 7), and (7, 11), with the tie between (7, 11) and (26, 30) broken arbitrarily.

Surprisingly, the `STRCCO` problem has not been studied before even though it is a natural variant of string indexing and several closely related problems have been extensively studied (see related work below).

### 1.1. Results and techniques

To state the complexity bounds, let  $n$  and  $m$  denote the lengths of  $S$  and  $P$ , respectively. An immediate approach to solve the `STRCCO` problem is to store the suffix tree of  $S$  using  $O(n)$  space. To answer a query on  $P$  with parameter  $k$ , we traverse the suffix tree to find all occurrences of  $P$ , construct the consecutive occurrences, and then sort these to output the top- $k$  close consecutive occurrences. Naively, this requires two sorts of size  $\text{occ}$ , where  $\text{occ}$  is the total number of occurrences of  $P$ , giving a query time of  $O(m + \text{occ} \log \text{occ})$ . Using more advanced data structures [12,11], the query time can be reduced to  $O(m + \text{occ})$  while still using linear space. Note that  $\text{occ}$  can be much larger than  $k$ . Alternatively, we can store at every node in the suffix tree the set of all consecutive occurrences sorted by distance using  $O(n^2)$  space. To answer a query we find the node corresponding to  $P$  and simply report the first  $k$  of the stored consecutive occurrences in optimal  $O(m + k)$  time.

To achieve better trade-offs, one might try to use a strategy similar to range minimum query (RMQ), where the ranges are subsequent ranges in the suffix array and the values are distances between pairs of suffix indexes in  $S$ . However, there are several problems with that idea: first, there are  $\Theta(|S|^2)$  possible pairs of suffix indexes within  $S$ , and it is not immediately clear how many of them can correspond to consecutive occurrences of a pattern (our arguments from Section 3 imply that this number is bounded by  $O(n \log n)$ ). Secondly, when taking the union of two ranges, the set of closest (consecutive) pairs can change completely: consider for example the string  $S = \text{A B A C A B A C D A B D A C D A B D A C}$ . While the string  $A$  has occurrences  $\{0, 2, 4, 6, 9, 12, 15, 18\}$ , the string  $AB$  has occurrences  $\{0, 4, 9, 15\}$  and  $AC$  has  $\{2, 6, 12, 18\}$ . Note that for  $P = A$ , the top-3 consecutive occurrences are (0, 2), (2, 4) and (4, 6), while for  $AB$  they are (0, 4), (4, 9) and (9, 15) and for  $AC$  they are (2, 6), (6, 12) and (12, 18). Both the pairs and the distances are completely different between  $A$  and its extensions. Thus, there is an issue of non-decomposability, which is a main challenge in this particular problem. However, in the rest of our paper we will show that we can use suffix tree decompositions and amortized arguments to bound the number of changes that can happen in the set of consecutive occurrences of substrings corresponding to positions on some paths in the suffix tree.

We obtain the following significantly improved time-space trade-offs:

**Theorem 1.** *Given a string  $S$  of length  $n$  and  $\epsilon$ ,  $0 < \epsilon \leq 1$ , we can build a data structure that can answer top- $k$  close consecutive occurrences queries using either*

- (i)  $O(n \log n)$  space and  $O(m + k)$  query time,
- (ii)  $O(\frac{n}{\epsilon})$  space and  $O(m + k^{1+\epsilon})$  query time, or
- (iii)  $O(\frac{n}{\epsilon})$  space and  $O(m + k \log^{1+\epsilon} n)$  query time.

Here,  $m$  is the length of the query pattern.

Hence, Theorem 1(i) achieves optimal query time using near-linear space. Alternatively, Theorem 1(ii) and (iii) achieve linear space, for constant  $\epsilon$ , while supporting queries in near-optimal  $O(m + k^{1+\epsilon})$  and  $O(m + k \log^{1+\epsilon} n)$  time, respectively.

To achieve Theorem 1 we develop several data structural techniques that may be of independent interest. First, we translate the problem into a line segment intersection problem on the heavy path decomposition of the suffix tree. This leads to the  $O(n \log n)$  space and optimal query time bound of Theorem 1(i). We note that Navarro and Thankachan [33] used similar techniques for a closely related problem (see related work below). To reduce space, we introduce a novel recursive clustering method on trees. The decomposition partitions the tree into a hierarchy of depth  $O(\log \log n)$  consisting of subtrees of doubly exponentially decreasing sizes. We show how to combine the decomposition with the techniques of the simple algorithm from Theorem 1(i) to obtain an  $O(n \log \log n)$  space and  $O(m + k^{1+\epsilon})$  query time solution. Then, we show how to efficiently compress the hierarchy of data structures into rank space leading to the linear space and  $O(m + k^{1+\epsilon})$  query time bound of Theorem 1(ii). Finally, we show how to use  $O(\log n)$  cluster decompositions of varying parameters together with an orthogonal range successor data structure to obtain the  $O(m + k \log^{1+\epsilon} n)$  time bound of Theorem 1(iii).

We apply these techniques to three related problems: Firstly, we address the natural “opposite” problem of reporting the  $k$  consecutive occurrences of largest distance, which can be solved using similar but not identical techniques. Secondly, we apply our framework to the related problem of reporting consecutive occurrences with distances within a specified interval,

considered by Navarro and Thankachan [33], and give an improvement for a special case. Finally, we show how this allows us to efficiently report all non-overlapping consecutive occurrences of a pattern.

## 1.2. Related work

To the best of our knowledge, the `Strcco` problem has not been studied before, even though distances between occurrences is a natural extension for string indexing and several related problems have been studied extensively.

A closely related problem was considered by Navarro and Thankachan [33], who showed how to efficiently report consecutive occurrences with distances within a specified interval. They gave an  $O(n \log n)$  space and  $O(m + \text{occ})$  time solution, where  $\text{occ}$  denotes the number of reported consecutive occurrences. We note that their result can be adapted to the `Strcco` problem to achieve the same bounds as in Theorem 1(i). However, our solution is simpler and does not rely on heavy word RAM techniques such as persistent van Emde Boas trees [14]. Our techniques can also be used to solve the problem considered by Navarro and Thankachan getting the same space and time bounds as they obtain, and we can achieve improved bounds in a special case (see Section 8).

A lot of work has been done on the related problem of string indexing for patterns under various *distance constraints*, where the goal is to report occurrences of (one or more) patterns that are within a given distance or interval of distances of each other [6,25,4,13,9,27,26]. An important difference between those works and our work is that all those solutions use time proportional to *all* pairs of occurrences with distances in the given range, in contrast to only finding *consecutive* occurrences. Note that if the goal is to find occurrences of a given maximal distance, one can find the close *consecutive* occurrences first and then construct all pairs satisfying the constraint.

Another line of related work is indexing collections of strings, called *documents*. Here the goal is to find documents containing patterns subject to various constraints. For a comprehensive overview see the survey by Navarro [30]. Several results on supporting efficient top- $k$  queries are known [29,35,22,10,21,22,31,37,20,23,32,28]. In this context the goal is to efficiently report the  $k$  documents of smallest weight. The weights can depend on the query and can be the distance between the closest pair of occurrences of a given pattern [22,31,35,22,31,28]. The problem can be solved in linear space and optimal  $O(k)$  time, in addition to finding the locus of the pattern in the suffix tree [35]. While this problem statement resembles ours, there is no direct translation from those results to our problem, since the documents are considered individually, and for a single document only the pair of occurrences with minimum distance within the document is considered.

Finally, we note that since the initial publication of the results in this paper, a subset of the authors have recently considered indexing for consecutive occurrences of two different patterns  $P_1$  and  $P_2$  [7].

## 1.3. Outline

The paper is organized as follows. In Section 2 we introduce some notation and recall results on string indexing. In Section 3 we build a simple data structure and prove Theorem 1(i). In Section 4 we recall a method for tree clustering and show how to use it to solve a simplified version of the problem. In Section 5 we introduce a recursive clustering method that allows us to use the ideas from Section 4 on the actual problem. This gives an  $O(n \log \log n)$  space and  $O(m + k^2)$  time data structure. In Section 6, we show how to reduce the space to linear while achieving the same query time, and then generalize the recursion to get Theorem 1(ii) for any  $0 < \epsilon \leq 1$ . In Section 7 we give the linear space solution with  $O(m + k \log^{1+\epsilon} n)$  query time. Finally, in Section 8 we apply our techniques to related problems.

## 2. Preliminaries

We introduce some notation and recall basic results from string indexing.

A *string*  $S$  of length  $n$  is a sequence  $S[0]S[1] \dots S[n-1]$  of characters from an alphabet  $\Sigma$ . A contiguous subsequence  $S[i, j] = S[i]S[i+1] \dots S[j-1]$  is a *substring* of  $S$ . The substrings of the form  $S[i, n]$  are the *suffixes* of  $S$ .

The *suffix tree* [38] is a compact trie of all suffixes of  $S$ , where  $\$$  is a symbol not in the alphabet, and is lexicographically smaller than any letter in the alphabet. Using perfect hashing [18], it can be stored in  $O(n)$  space and solve the string indexing problem (i.e., find and report all occurrences of a pattern  $P$ ) in  $O(m + \text{occ})$  time, where  $m$  is the length of  $P$  and  $\text{occ}$  is the number of times  $P$  occurs in  $S$ . The *suffix array* stores the suffix indices of  $S$  in lexicographic order. The suffix tree has the property that the leaves below any node represent suffixes that appear in consecutive order in the suffix array. Brodal et al. [12] show that there is a linear space data structure that allows outputting all entries within a given range of an array in sorted order using time linear in size of the output. This data structure on the suffix array together with the suffix tree can output all occurrences of a pattern *sorted by text order* in  $O(n)$  space and  $O(m + \text{occ})$  time.

For any node  $v$  in the suffix tree, we define  $\text{str}(v)$  to be the string found by concatenating all labels on the path from the root to  $v$ . The *locus* of a string  $P$ , denoted  $\text{locus}(P)$ , is the minimum depth node  $v$  such that  $P$  is a prefix of  $\text{str}(v)$ .

## 3. A simple $O(n \log n)$ space solution

In this section, we present a simple solution that solves the `Strcco` problem in  $O(n \log n)$  space and  $O(m + k)$  query time. This solution will be a key component in our more advanced structures in the following sections. We note that the

results by Navarro and Thankachan [33] for the related problem of reporting consecutive occurrences with distances within a specified interval can be modified to achieve the same complexities. However, our solution is simpler and does not rely on heavy word RAM techniques such as persistent van Emde Boas trees [14].

Let  $D(v)$  denote the set of consecutive occurrences of  $\text{str}(v)$ . Naively, if we store for each node  $v$  the set  $D(v)$  in sorted order, we can directly answer a query for the top- $k$  close consecutive occurrences of a pattern  $P$  by reporting the  $k$  smallest elements in  $D(\text{locus}(P))$ . This solves the problem in  $O(n^2)$  space and  $O(m+k)$  query time. The main idea in our simple solution is to build a heavy path decomposition of the suffix tree and compactly represent sets on the same path via a reduction to the orthogonal line segment intersection problem while maintaining optimal time queries. This is similar to the data structure by Navarro and Thankachan [33], but our reduction is different.

**Heavy path decomposition** A heavy path decomposition of a tree  $T$  is defined as follows: Starting from the root, at every node, we choose the edge to the child with the largest subtree as *heavy edge*, until we reach a leaf. Ties are broken arbitrarily. This defines a heavy path, and all edges hanging off the heavy path are *light edges*. The root of a heavy path  $h$  is called the *apex* of the path, denoted  $\text{apex}(h)$ . We then recursively decompose all subtrees hanging off the path. The heavy path decomposition has the following property:

**Lemma 2** (Sleator and Tarjan [36]). *Given a tree  $T$  of size  $n$  and a heavy path decomposition of  $T$ , any root-to-leaf path in  $T$  contains at most  $O(\log n)$  light edges.*

**Orthogonal line segment intersection** Similarly as Navarro and Thankachan [33], we are going to reduce the problem to a geometric problem on orthogonal line segment intersection. Specifically, we are going to reduce to the following problem: Let  $L$  be a set of  $n$  vertical line segments in a plane with non-negative  $x$ -coordinates. The *orthogonal line segment intersection problem* is to preprocess  $L$  to support the query:

- smallest-segments  $(y_0, k)$ : return the first  $k$  segments intersecting the horizontal line with  $y$ -coordinate  $y_0$  in left-to-right order.

We will assume that  $y_0$  is an integer, which suffices for our purpose. Let  $N$  be the maximum  $y$ -coordinate of a segment in  $L$ . The following lemma follows easily from the results on partially persistent data structures by Driscoll et al. [16].

**Lemma 3.** *We can solve the line segment intersection problem as described above in  $O(n+N)$  space and  $O(k)$  time.*

**Proof.** Consider the  $x$ -coordinates as the elements of a set  $X$  and the  $y$ -coordinate as time. The version of  $X$  at a time  $y_0$  contains exactly the  $x$ -coordinates of the line segments which intersect the horizontal line at  $y_0$ . Now, the data structure is a partially persistent sorted doubly linked list  $L$  on the elements of  $X$ . The elements are sorted in increasing order. Since we have at most  $n$  line segments, the maximum size of  $X$  as well as the maximum number of updates is  $n$ . Each update changes only  $O(1)$  pointers in the linked list. Using the node copying technique from Driscoll et al. [16] we can build a partially persistent linked list using  $O(n)$  space. To be able to find version  $y_0$  in constant time, we keep an array of size  $N$  with a pointer to the root of the version at each possible time step. For a query  $(y_0, k)$ , use the sorted linked list  $L$  to report the  $k$  smallest elements at time  $y_0$ .


If we use a linear scan to find the place to insert an element or find the element to be deleted we get a preprocessing time of  $O(n^2)$ . This can be improved to  $O(n \log n)$  by using a (non-persistent) balanced binary search tree during the preprocessing holding all elements in the current version of  $L$  together with a pointer to their node in the current version. When performing an update the binary search tree is used to find the position where the element must be inserted/deleted in  $O(\log n)$  time. After the preprocessing step the tree is discarded.  $\square$

### 3.1. Data structure

We construct a heavy path decomposition of the suffix tree  $T$  of  $S$ . Our data structure consists of a line segment data structure from Lemma 3 for each heavy path of  $T$  that compactly encodes the sets  $D(v)$  for each node  $v$  on the path.

We describe the contents of the data structure for a single heavy path  $h = v_1, \dots, v_\ell$ , where  $v_1$  is the apex of the path. Consider a consecutive occurrence  $(i, j)$  on some node on  $h$  and imagine moving down the heavy path from top to bottom. Either  $(i, j)$  is a consecutive occurrence at the apex of  $h$  or it will become a consecutive occurrence as soon as every suffix starting at an index between  $i$  and  $j$  has branched off the heavy path. Then it will stay a consecutive occurrence until either the suffix corresponding to  $i$  or the suffix corresponding to  $j$  (or both) branch off  $h$ . Thus, there exists an interval  $[d_1, d_2]$  of depths on the heavy path such that  $(i, j) \in D(v_d)$  if and only if  $d \in [d_1, d_2]$ . We say that  $(i, j)$  is *alive* in this interval.

We encode the consecutive occurrences by line segments in the plane which describe their distance and the interval in which they are alive along the heavy path. Conceptually, the  $x$ -coordinate in our coordinate system corresponds to the distance of a consecutive pair, and the  $y$ -coordinate corresponds to the depth on the heavy path. Now, for each consecutive occurrence  $(i, j)$ , we define a vertical line segment with  $x$ -coordinate set to its distance, and  $y$ -coordinate spanning the



**Fig. 2.** Line segments for a heavy path from the suffix tree for “BATMAN-AND-ANNA-SING-NANANANA-AND-EAT-BANANAS”. Here, if we have overlapping line segments, we denote by a number how many consecutive occurrences the current segment corresponds to. At depth 1, we have a line segment corresponding to pairs of consecutive occurrences of string A - there are six pairs that have a distance of 2, three pairs that have a distance of 3, two pairs that have a distance of 4, and so on. At depth 2, we encode the consecutive occurrences of string AN. Some of them are the same as for string A.

interval  $[d_1, d_2]$ , where  $[d_1, d_2]$  is the interval in which  $(i, j)$  is alive. For an example, see Fig. 2. Our data structure for  $h$  stores the above line segments in the line segment data structure from Lemma 3. For each line segment in the data structure we store a pointer to the pair of occurrences it represents. The full data structure for  $T$  consists of the line segment data structures for all of the heavy paths in  $T$ .

*Space analysis* For a given heavy path  $h$ , a leaf in the subtree of  $\text{apex}(h)$  can be in at most two consecutive occurrences in  $D(\text{apex}(h))$ . Consider a light edge  $(v_d, u)$  leaving  $h$  at depth  $d$ . Any leaf in the subtree rooted at  $u$  can be part of at most two consecutive occurrences in  $D(v_d)$ . A single leaf can thus make at most two consecutive occurrences from  $D(v_d)$  disappear in  $D(v_{d+1})$  and at most one new consecutive occurrence appear. If we consider all leaves that leave  $h$ , we therefore get at most three changes per leaf. Thus, for a given heavy path  $h$  a leaf in the subtree of  $\text{apex}(h)$  can be in at most two consecutive occurrences in  $D(\text{apex}(h))$  and can cause at most three changes of line segments in the line segment data structure for  $h$ . Since any root-to-leaf path can intersect at most  $\log n$  heavy paths, any leaf can contribute  $O(\log n)$  line segments. Overall, this means that there are at most  $O(n \log n)$  line segments in total. For a single heavy path  $h$  the line segment data structure from Lemma 3 uses linear space in the number of segments and the length of  $h$ . The sum of the lengths of the heavy paths is  $O(n)$ , since the heavy paths are disjoint. Thus the total space usage is  $O(n \log n)$ .

### 3.2. Algorithm

Given a pattern  $P$  and an integer  $k$  we can now answer a query as follows. We begin by finding  $\text{locus}(P)$  in the suffix tree. Let  $h$  be the heavy path that the locus is on and let  $d_P$  be the depth of  $\text{locus}(P)$  on  $h$ . We do a smallest-segments  $(d_P, k)$  query on the line segment data structure stored for  $h$  and report the consecutive occurrences corresponding to the returned line segments.

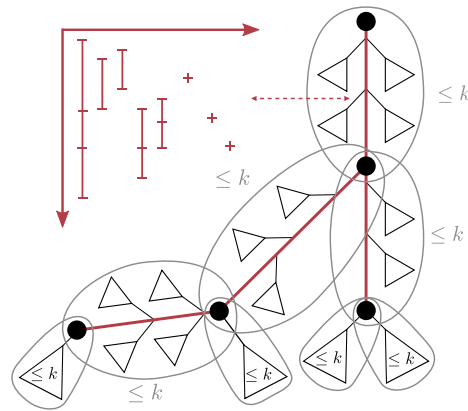
*Correctness* By definition,  $D(\text{locus}(P))$  contains the consecutive occurrences of  $P$ . Thus, every consecutive occurrence of  $P$  defines a line segment in the data structure for  $h$  and the horizontal line with  $y$ -coordinate set to  $d_P$  intersects exactly those line segments. Since we set the  $x$ -coordinate of every line segment to the distance of its consecutive occurrence, the line segments are sorted left-to-right by increasing distance. Thus, the first  $k$  line segments intersecting the horizontal line at  $y = d_P$  correspond to the top- $k$  close consecutive occurrences.

*Time analysis* The time for finding  $\text{locus}(P)$  in the suffix tree is  $O(m)$ . The time for querying the line segment data structure from Lemma 3 is  $O(k)$ , so the total time complexity is  $O(m + k)$ . This proves Theorem 1(i).

## 4. A linear space solution for fixed $k$

In this section, we present a linear space and  $O(m + k)$  time solution for the simpler problem where  $k$  is known at construction time. That is, given a string  $S$  and a positive integer  $k$ , we preprocess  $S$  into a compact data structure such that given a pattern string  $P$ , we can efficiently find the top- $k$  close consecutive occurrences of  $P$  in  $S$ . This data structure demonstrates one of the key ideas that our final result builds on.

The main idea behind the data structure is to store the line segment solution from Section 3 for some path segments of the suffix tree, such that all nodes that are not on these paths are within small subtrees. For nodes within such small



**Fig. 3.** The suffix tree is divided into clusters (gray loops) of size  $\leq k$  which are either leaf clusters, or path clusters with spines marked in red. For every spine we store a line segment data structure, also marked in red. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

subtrees we can find all consecutive occurrences without spending too much time. Specifically, we will partition the suffix tree into *clusters*, satisfying some properties. We are going to define this cluster partition next.

### 4.1. Cluster partition

For a connected subgraph  $C \subseteq T$ , a *boundary node*  $v$  is a node  $v \in C$  such that either  $v$  is the root of  $T$ , or  $v$  has an edge leaving  $C$  – that is, there exists an edge  $(v, u)$  in the tree  $T$  such that  $u \in T \setminus C$ . A *cluster* is a connected subgraph  $C$  of  $T$  with at most two boundary nodes. A cluster with one boundary node is called a *leaf cluster*. A cluster with two boundary nodes is called a *path cluster*. For a path cluster  $C$ , the two boundary nodes are connected by a unique path. We call this path the *spine* of  $C$ . A *cluster partition* is a partition of  $T$  into clusters, i.e. a set  $CP$  of clusters such that  $\bigcup_{C \in CP} V(C) = V(T)$  and  $\bigcup_{C \in CP} E(C) = E(T)$  and no two clusters in  $CP$  share any edges. Here,  $E(G)$  and  $V(G)$  denote the edge and vertex set of a (sub)graph  $G$ , respectively. We need the next lemma which follows from well-known tree decompositions [1–3,17] (see Bille and Gørtz [5] Lemma 5.1 for a direct proof).

**Lemma 4.** *Given a tree  $T$  with  $n$  nodes and a parameter  $\tau$ , there exists a cluster partition  $CP$  such that  $|CP| = O(n/\tau)$  and every  $C \in CP$  has at most  $\tau$  nodes. Furthermore, such a partition can be computed in  $O(n)$  time.*

### 4.2. Data structure

For the suffix tree of  $S$ , we build a clustering as in Lemma 4 with parameter  $\tau$  set to  $k$  to get  $O(n/k)$  clusters of size at most  $k$ . For the spine of every path cluster, we build a line segment data structure similar to the one from Section 3. The difference is that for any depth, we only maintain the line segments that correspond to the top- $k$  close consecutive occurrences for that depth. Let  $v_1, \dots, v_l$  denote the nodes on the spine, starting at the top boundary node. Note that for any consecutive occurrence that appears for the first time in  $D(v_{d+1})$  there is a consecutive occurrence in  $D(v_d)$  of smaller distance which is no longer present in  $D(v_{d+1})$ . It follows that, when moving down the spine, once a consecutive occurrence  $(i, j)$  is amongst the  $k$  closest, it will stay amongst the  $k$  closest until suffix  $i$  or  $j$  branches off the spine. Thus, there exists an interval  $[d_1, d_2]$  of consecutive depths such that  $(i, j)$  is amongst the  $k$  closest pairs in  $D(v_d)$  if and only if  $d \in [d_1, d_2]$ . For a consecutive occurrence  $(i, j)$  that is amongst the  $k$  closest for any  $v$  on the spine, we define a line segment where the  $x$ -coordinate is its distance and the  $y$ -coordinate is spanning the interval  $[d_1, d_2]$ , where  $[d_1, d_2]$  is the interval in which  $(i, j)$  is amongst the  $k$  closest pairs. For these line segments we store the data structure from Lemma 3. Again, for each line segment we store the pair of occurrences it represents. We store this data structure for the spine of each cluster and for every node that is on that spine we store a pointer to the data structure. For boundary nodes that are on multiple spines we store a pointer to any one of them. See Fig. 3 for an illustration of this structure. Additionally we store the suffix array and the sorted range reporting data structure of Brodal et al. [12] on the suffix array.

*Space analysis* We show that for every path cluster there are  $O(k)$  line segments: We still have the property that a line segment only ends if a corresponding leaf branches off the spine. In that case, it might be replaced either by a new consecutive occurrence or by a consecutive occurrence that was there before but was not amongst the  $k$  closest. Note that at any node on the spine except the boundary nodes, any subtrees branching off the spine are fully contained within the cluster, and as such have total size at most  $k$ . Between the top boundary node and the next node on the spine, we have no bound as to how many leaves can branch off – however, since we only store line segments corresponding to the top- $k$

consecutive occurrences, at most  $k$  line segments can be replaced by  $k$  other line segments. For the rest of the spine, at most  $k$  leaves can branch off in total. Every leaf that branches off can cause at most two line segments to end and two new line segments to begin. As such there can be at most  $O(k)$  line segments. As the size of the line segment data structure is linear in the number of line segments and in the length of the spine, any line segment data structure of a path cluster uses  $O(k)$  space. As both the sorted range reporting data structure and the suffix array have linear space complexity, the complete data structure occupies  $O((n/k)k + n) = O(n)$  space.

### 4.3. Algorithm

Given a pattern  $P$  we can now answer the top- $k$  query. We begin by finding  $\text{locus}(P)$  in the suffix tree. If the locus is on a spine, we query the line segment data structure for that spine. Otherwise the locus is either in a subtree hanging off a spine or in a leaf cluster. In both cases, there are at most  $k$  occurrences of our pattern  $P$ . We find all occurrences of  $P$  in text order, using the sorted range reporting data structure. This allows us to report the consecutive occurrences: Let  $i_1, \dots, i_l$  denote the leaves in text order, then the consecutive occurrences are  $(i_1, i_2), (i_2, i_3), \dots, (i_{l-1}, i_l)$ . Note that  $l \leq k$ , since the size of the subtree is at most  $k$ .

*Correctness* By construction, for any depth on a spine, the top- $k$  close consecutive occurrences of the corresponding substring will have corresponding line segments present at that depth in the line segment data structure. If the locus is on a spine, then by the arguments in Section 3, the line segment data structure will report the top- $k$  close consecutive occurrences. If the locus is not on a spine, then there are at most  $k$  occurrences of  $P$  in total, since any subtree hanging off a spine and any leaf cluster have at most  $k$  leaves. Thus, by constructing and reporting all consecutive occurrences of  $P$  we report the top- $k$  close consecutive occurrences.

*Time analysis* We find the locus in  $O(m)$  time. If we land on a spine we report in  $O(k)$  time. Otherwise, we are in a subtree of size at most  $O(k)$  and thus  $P$  has at most  $k$  occurrences. Using sorted range reporting we can find the occurrences in text order using  $O(k)$  time. The total time for a query is thus  $O(m + k)$ .

We are going to use this data structure with different parameters in Section 5. For a general parameter  $\tau$ , we have the following lemma:

**Lemma 5.** *For any positive integer  $\tau$ , there exists a cluster partition of the suffix tree and a linear space data structure with the following properties:*

1. *For any  $k \leq \tau$  and  $P$  such that  $\text{locus}(P)$  is on the spine of a cluster, we can report the top- $k$  close consecutive occurrences in  $O(m + k)$  time.*
2. *For any  $P$  such that  $\text{locus}(P)$  is not on a spine, we can report the top- $k$  close consecutive occurrences in  $O(m + \tau)$  time.*

**Proof.** We build the data structure described in this section for parameter  $\tau$  taking the role of  $k$ . In Case 1, we query the line segment data structure for the depth of  $\text{locus}(P)$  on the path and  $k$ . Since  $k \leq \tau$  this will correctly output the top- $k$  close consecutive occurrences of  $P$ . In Case 2, we have shown that we can construct the top- $\tau$  close consecutive occurrences. Using the linear time selection algorithm by Blum et al. [11] we can find the top- $k$  of those: We use the algorithm to find the consecutive occurrence of  $k^{\text{th}}$  smallest distance  $d$ ; then we traverse all the consecutive occurrences and output those of distance  $\leq d$ . If needed, we crop the output to report no more than  $k$  consecutive pairs.  $\square$

## 5. An $O(n \log \log n)$ space solution for general $k$

We now show how to leverage the solution from Lemma 5 to obtain a data structure that can answer queries for any  $k$ . The idea is to recursively cluster the suffix tree, such that we always either land on a spine with a sufficient number of consecutive occurrences stored, or in a sufficiently small subtree.

### 5.1. Data structure

Our data structure consists of the suffix tree decomposed into clusters of decreasing size, with the line segment data structure stored for every spine as before. We build it in the following way. First we build the solution from Lemma 5 with parameter  $\tau_1 = \sqrt{n}$ , resulting in clusters of size at most  $\sqrt{n}$ . For every subtree hanging off a spine and every leaf cluster, we apply the solution with parameter  $\tau_2 = \sqrt{\tau_1}$ . We keep recursively applying the solution with parameter  $\tau_i = \sqrt{\tau_{i-1}}$  until reaching a constant cluster size. For notational convenience, additionally define  $\tau_0 = n$ . See Fig. 4 for an illustration of this data structure. Again we additionally store the suffix array and the sorted range reporting data structure of Brodal et al. [12] on the suffix array.






Fig. 4. Here, we see the recursive clustering: The black clustering is the coarsest clustering and the green and blue are finer sub-clusterings.

*Space analysis* The suffix array and sorted range reporting structure occupy  $O(n)$  space. For a tree of size  $\tilde{n}$  and any  $\tau$ , the data structure from Lemma 5 uses at most  $O(\tilde{n})$  space. Since at every recursion level, we build the data structure from Lemma 5 on non-overlapping subtrees of the suffix tree, every recursion level uses at most  $O(n)$  space. As the cluster size at every level of recursion is the square root of the previous cluster size, there are at most  $O(\log \log n)$  levels. The complete data structure thus uses  $O(n \log \log n)$  space.

### 5.2. Algorithm

Given a query with pattern  $P$  and parameter  $k$ , we can now answer in the following way. As before, we begin by finding the locus of the pattern in the suffix tree. This node is now either on the spine of some cluster or in a cluster of constant size. If it is on the spine of a cluster of size  $\tau_i$ , and if  $k \leq \tau_i$ , then we query the line segment data structure for that spine, which allows us to report the top- $k$  close consecutive occurrences. Otherwise, we find all occurrences of  $P$  and construct the top- $k$  close consecutive occurrences by using linear time selection as in the proof of Lemma 5.

*Correctness* The correctness of the algorithm follows by the same arguments as previous sections.

*Time analysis* Finding the locus in the suffix tree takes  $O(m)$  time. The locus is either on the spine of a cluster, or within a cluster of constant size. In a constant sized cluster, clearly we can do all operations described above in constant time. If the locus is on the spine of a cluster with parameter  $\tau_i$ , and  $k \leq \tau_i$ , then we are in case 1 of Lemma 5 with  $\tau = \tau_i$  and can report the top- $k$  close consecutive occurrences using a total of  $O(m + k)$  time. If  $k > \tau_i$ , then we are in case 2 of Lemma 5 with  $\tau = \tau_{i-1}$ . Note that  $\tau_{i-1} = \tau_i^2 < k^2$ . Therefore, we can find the top- $k$  close consecutive occurrences in  $O(m + \tau_{i-1}) = O(m + k^2)$  time. In total, the worst case query time is then  $O(m + k^2)$ . In summary, this gives the following result:

**Lemma 6.** *Given a string  $S$  of length  $n$ , we can build a data structure that can answer top- $k$  close consecutive occurrences queries using  $O(n \log \log n)$  space and  $O(m + k^2)$  query time. Here,  $m$  is the length of the query pattern.*

## 6. A linear space solution

We now show how to reduce the space consumption of the solution presented in Section 5. Observe that in any cluster of level  $i$ , we only have  $O(\tau_i)$  objects. If we can reduce all objects within a cluster to a “universe size” of  $O(\tau_i)$  instead of  $O(n)$ , we can use  $O(\tau_i \log \tau_i)$  bits instead of  $O(\tau_i \log n)$  bits per cluster. In the following, consider a cluster  $C$  of level  $i$ .

*Reducing the line segment data structure* In the line segment data structure for cluster  $C$ , by the analysis of previous sections, there are at most  $O(\tau_i)$  line segments and  $\tau_i$  different depths on the path. Let  $c$  be a constant such that there are at most  $c\tau_i$  line segments for each cluster. We map every unique  $x$ -coordinate of a line segment to a unique element in  $\{1, \dots, c\tau_i\}$  in a way that preserves order. That is, map the minimum  $x$ -coordinate to 1, the smallest  $x$ -coordinate that is bigger than the minimum to 2, and so on. This gives us a modified line segment data structure that preserves the properties we need but is restricted to a  $c\tau_i \times \tau_i$  grid.

**Reducing the leaf pointers** For any line segment, we have to store pointers that allow us to report the corresponding pair of consecutive occurrences. Doing so naively uses  $2 \log n$  bits per line segment. In the following, we show how to reduce that to  $4 \log \tau_i$ , for a cluster  $C$  of level  $i$ . The idea is to store the offset within the suffix array range defined by the top boundary node  $r$  of  $C$ . More precisely, let  $[a_r, b_r]$  be the range in the suffix array spanning the leaves below  $r$ . Then for any leaf  $l$  in the subtree rooted at  $r$  define  $\text{off}(l) = SA^{-1}(l) - a_r$ . By the way our recursion is defined,  $C$  is fully contained in a subtree of size at most  $\tau_i^2$ , and thus  $r$  has at most  $\tau_i^2$  leaves below it. It follows that for any leaf  $l$  in the subtree of  $r$ ,  $\text{off}(l)$  is a number between in  $[0, \tau_i^2 - 1]$  and can be stored using  $2 \lceil \log \tau_i \rceil$  bits.

### 6.1. Data structure

Our data structure is now defined as follows: We have a clustering of the suffix tree as in Section 5. For every spine on level  $i$ , we store the line segment data structure reduced to a  $c\tau_i \times \tau_i$  grid. Every line segment corresponding to a pair  $(i, j)$  stores the pair  $(\text{off}(i), \text{off}(j))$  as additional information. For every node on the spine, we store a pointer to the spine data structure and to the top boundary node of the spine. Additionally, we store the suffix array and the sorted range reporting structure, as well as two integers for every node in the suffix tree, that define the range of leaves below the node in the suffix array.

**Space analysis** The suffix array and the sorted range reporting data structure use space  $O(n)$ . Storing the range in the suffix array plus at most two pointers per node use  $O(n)$  space. For a cluster  $C$  of level  $i$ , we store the line segment data structure from Lemma 3 for a  $c\tau_i \times \tau_i$  grid. Since the data structure from Lemma 3 works in the word RAM model (as do all data structures presented in this paper), we can store the data structure using  $O(\tau_i \log \tau_i)$  bits. For each of the at most  $c\tau_i$  line segments we store  $4 \log \tau_i$  bits for the encoding of the consecutive pair. Thus, we can store the data structure for cluster  $C$  using  $O(\tau_i \log \tau_i)$  bits. As in the previous section, at every recursion level, we cluster non-overlapping subtrees. The reduced cluster solution of a subtree of size  $\tilde{n}$  with parameter  $\tau_i$  uses  $O(\frac{\tilde{n}}{\tau_i} \log \tau_i) = O(\tilde{n} \log \tau_i)$  bits. The total space for all clusters of level  $i$  thus becomes  $O(n \log \tau_i)$ . Summing over all recursion levels, we get

$$\sum_{i=0}^{\lfloor \log \log n \rfloor} O(n \log \tau_i) = \sum_{i=0}^{\lfloor \log \log n \rfloor} O\left(n \log n^{(1/2^i)}\right) = \sum_{i=0}^{\lfloor \log \log n \rfloor} \frac{1}{2^i} O(n \log n) = O(n \log n) \text{ bits,}$$

that is,  $O(n)$  words.

### 6.2. Algorithm

We query the data structure as follows: If we land on a spine and  $k \leq \tau_i$ , we query the line segment data structure and get  $k$  pairs of the form  $(\text{off}(i), \text{off}(j))$ . We then use the pointer to get to the root of the spine and use the range in the suffix array to translate each encoding back to the original suffix number, using constant time per leaf. Otherwise, we proceed as described in Section 5. Since the decoding can be done in constant time per leaf, the time complexities are the same as in Section 5. We have shown the following result:

**Lemma 7.** *Given a string  $S$  of length  $n$ , we can build a data structure that can answer top- $k$  close consecutive occurrences queries using  $O(n)$  space and  $O(m + k^2)$  query time. Here,  $m$  is the length of the query pattern.*

In order to get Theorem 1(ii), we cluster according to a parameter  $\epsilon$ ,  $0 < \epsilon \leq 1$ , using the following recursion:

$$\tau_0 = n \quad \text{and} \quad \tau_i = \tau_{i-1}^{\frac{1}{1+\epsilon}}.$$

Hence, the total space in bits is now:

$$\sum_{i=0}^{\lfloor \log_{1+\epsilon} \log n \rfloor} O\left(n \log n^{1/(1+\epsilon)^i}\right) = \sum_{i=0}^{\infty} \left(\frac{1}{1+\epsilon}\right)^i O(n \log n) = \left(1 + \frac{1}{\epsilon}\right) O(n \log n),$$

that is,  $O(\frac{n}{\epsilon} \log n)$  bits, so  $O(\frac{n}{\epsilon})$  words. For the query time, there are again two cases. In the case where  $\text{locus}(P)$  is on a spine with  $k \leq \tau_i$ , we get optimal  $O(m + k)$  time, as before. For the other case, we have at most  $\tau_{i-1} = \tau_i^{1+\epsilon} < k^{1+\epsilon}$  occurrences of  $P$ , which gives us a time complexity of  $O(m + k^{1+\epsilon})$ . This concludes the proof of Theorem 1(ii).

## 7. A different tradeoff

In this section we give a solution query time  $O(m + k \log^{1+\epsilon} n)$ . The idea is to store a finer set of cluster decompositions than in the previous section and store sublinear information for each cluster decomposition. Then we use a bounded number of orthogonal range successor queries in each cluster.

*Orthogonal range successor* The *orthogonal range successor problem* is to preprocess an array  $A[0, \dots, n - 1]$  into a data structure that efficiently supports the following queries:

- $\text{RangeSuccessor}(a, b, x)$ : return the successor of  $x$  in  $A[a, \dots, b]$ , that is, the minimum  $y > x$  such that there is an  $i \in [a, b]$  with  $A[i] = y$ .
- $\text{RangePredecessor}(a, b, x)$ : return the predecessor of  $x$  in  $A[a, \dots, b]$ , that is, the maximum  $y < x$  such that there is an  $i \in [a, b]$  with  $A[i] = y$ .

Nekrich and Navarro [34] give a linear space data structure such that each range successor query takes  $O(\log^\epsilon n)$  time. We will use a range successor data structure on the suffix array to answer the following type of queries: Given an index  $i$  and the suffix array range of a pattern  $P$ , find the next position in the text after  $i$  where  $P$  occurs.

### 7.1. Data structure

We store the linear space range successor data structure from Nekrich and Navarro [34] and the sorted range reporting data structure by Brodal et al. [12] on the suffix array of  $S$ . Further, we store the suffix tree of  $S$  together with the following cluster decompositions. For each  $\kappa = 2, 4, 8, 16, \dots, 2^{\lceil \log n \rceil}$  we build the clustering decomposition of Lemma 4 of the suffix tree for cluster size  $\tau = \kappa \log n$ . For each boundary node  $v$ , we store the top- $\kappa$  close consecutive occurrences of  $\text{str}(v)$ , sorted by text position. For each node  $v$  in the suffix tree, we additionally store two bit vectors of length  $\log n$ . The first one is used to store for which  $\kappa$  node  $v$  is a boundary node, that is, the  $i$ th bit is set to 1 if  $v$  is a boundary node in the cluster decomposition with  $\kappa = 2^i$  and 0 otherwise. Similarly, the other bit vector stores for which  $\kappa$  node  $v$  is on a spine in the cluster decomposition with  $\kappa = 2^i$ .

*Space analysis* The suffix array, the range successor data structure and the sorted range reporting data structure all uses  $O(n)$  space. The suffix tree together with the bit vectors saved in the nodes also uses  $O(n)$  space. By Lemma 4, there are  $O(n/(\kappa \log n))$  boundary nodes in the cluster decomposition with cluster size  $\kappa \log n$ . For each boundary node we store  $\kappa$  values and thus the space used for a fixed  $\kappa$  is  $O(\kappa n/(\kappa \log n)) = O(n/\log n)$ . There are  $O(\log n)$  different values of  $\kappa$  and therefore the total space is  $O(n)$ .

### 7.2. Algorithm

Given a pattern  $P$  and parameter  $k$  we can now answer the top- $k$  query in the following way. As before, we begin by finding the locus of  $P$  in the suffix tree. Then we find the smallest power of two bigger than  $k$ , i.e.  $\kappa = 2^{\lceil \log k \rceil}$ , and consider the cluster decomposition defined for  $\kappa$ . There are two cases depending on whether  $\text{locus}(P)$  is on a spine in the cluster decomposition for  $\kappa$  or not.

If  $\text{locus}(P)$  is not on a spine, then as in Lemma 5 (2), there are at most  $\tau = \kappa \log n$  leaves below  $\text{locus}(P)$ , and we can construct the top- $k$  close occurrences in time  $O(\kappa \log n)$ . If  $\text{locus}(P)$  is on a spine, then we find the lower boundary node  $b$  of the cluster (note we can do that by traversing the suffix tree and checking at most  $O(\kappa \log n)$  nodes). We have the following property.

**Claim 8.** *Each of the top- $k$  close consecutive occurrences of  $P$  is either*

1. *stored at  $b$  or*
2. *includes an occurrence of  $P$  corresponding to a leaf within the cluster.*

**Proof.** Any top- $k$  consecutive occurrence of  $P$ , where both occurrences are below  $b$ , is also among the top- $k$  consecutive occurrences of  $\text{str}(b)$ . This is true because  $P$  is a prefix of  $\text{str}(b)$ , so the occurrences of  $\text{str}(b)$  is a subset of the occurrences of  $P$ . Thus any consecutive occurrence of  $P$  where both occurrences are below  $b$  is also a consecutive occurrence of  $\text{str}(b)$ . Any consecutive occurrence of  $\text{str}(b)$  that is not a consecutive occurrence of  $P$  must be split by an occurrence of  $P$  that is not below  $b$  giving rise to a least two closer consecutive occurrences of  $P$ . Thus the  $i$ -closest occurrence of  $\text{str}(b)$  must have distance at least the same as the  $i$ -closest occurrence of  $P$ . Every occurrence of  $P$  that is not below  $b$  is within the cluster.  $\square$

We find all occurrences of  $P$  that are within the cluster in text order using the sorted range reporting data structure. We can do this using two calls to the sorted range reporting data structure since the occurrences of  $P$  within the cluster correspond to two intervals in the suffix array, namely the range of  $\text{locus}(P)$  minus the range of  $b$ . For each such occurrence  $i$ , we use an orthogonal range successor query  $j = \text{RangeSuccessor}(\text{range}(\text{locus}(P)), i)$  to find the next occurrence  $j$  and then an orthogonal predecessor query  $i' = \text{RangePredecessor}(\text{range}(\text{locus}(P)), j)$  to find the last occurrence  $i'$  before  $j$ . This gives us a consecutive occurrence  $(i', j)$ . To avoid recomputing the same consecutive occurrence we skip through the list until we get to an occurrence that is after  $i'$ .

Now we have the sorted list of the top- $\kappa$  consecutive occurrences of  $\text{str}(b)$  stored at  $b$ , and a sorted list of all consecutive occurrences of  $P$  that include an occurrence corresponding to a leaf within the cluster. By Claim 8, any of the top- $k$  consecutive occurrences of  $P$  is part of one of these lists. However, some of the consecutive occurrences of  $\text{str}(b)$  might not be consecutive occurrences of  $P$ , since  $P$  might have extra occurrences inbetween. Let  $(i, j)$  be a consecutive occurrence of  $\text{str}(b)$ , and let  $i'$  be the last occurrence of  $P$  before  $j$ . If  $i' \neq i$  then  $i'$  is within the cluster and thus  $(i', j)$  is part of the consecutive occurrences we already computed. We merge the two lists, deleting all consecutive occurrences of  $\text{str}(b)$  that are not consecutive occurrences of  $P$  and all duplicates. We then find the top- $k$  of remaining consecutive occurrences, by using the linear time selection algorithm as in the proof of Lemma 5.

*Time analysis* Finding the locus in the suffix tree takes  $O(m)$  time. Finding  $b$  takes  $O(\kappa \log n)$  time, and finding the occurrences of  $P$  within the cluster using sorted range reporting also takes  $O(\kappa \log n)$  time. We make  $O(\kappa \log n)$  calls to the orthogonal range reporting data structure each using  $O(\log^\epsilon n)$  time. In total the time for this is  $O(\kappa \log^{1+\epsilon} n)$ . Merging the two lists and using the selection algorithm takes time linear in the total length of the two lists which is  $O(\kappa \log n)$ . Since  $\kappa < 2k$  the total time complexity  $O(m + k \log^{1+\epsilon} n)$ . This concludes the proof of the main result.

## 8. Extensions

Our results can be extended to a couple of related problems. In Section 8.1, we show how we can modify our data structure to solve the “opposite” problem of reporting the  $k$  consecutive occurrences of *largest* distance. The extension is quite natural, though it does require some careful analysis. In Section 8.2, we then relate the solutions from Section 8.1 and the solutions to  $\text{Srrcco}$  to the problem of finding consecutive occurrences with distances in a specified interval, considered by Navarro and Thankachan [33]. We show improved complexities for the special case where one of the interval bounds is known at indexing time. Finally, we show how to use those results to efficiently find all pairs of non-overlapping consecutive occurrences.

### 8.1. Top- $k$ far consecutive occurrences

Given a pattern  $P$  and an integer parameter  $k > 0$ , define the *top- $k$  far consecutive occurrences* of  $P$  to be the  $k$  consecutive occurrences of  $P$  in  $S$  with the largest distances. Given a string  $S$  the *string indexing for top- $k$  far consecutive occurrences problem* ( $\text{Srrfco}$ ) is to preprocess  $S$  into a data structure that supports top- $k$  far consecutive occurrences queries. The goal is to obtain a compact data structure while supporting fast queries in terms of the length of the pattern  $P$  and the number of reported occurrences  $k$ .

*Line segments and an  $O(n \log n)$  space solution* We can solve the  $\text{Srrfco}$  problem using the same strategy as for the  $\text{Srrcco}$  problem, with small modifications. We need a similar data structure from Lemma 3 to report the line segments with largest  $x$ -coordinates. As previously, assume we are given a set  $L$  of  $n$  vertical line segments. We need a data structure for the following problem:

- largest-segments  $(y_0, k)$ : return the first  $k$  segments intersecting the horizontal line with  $y$ -coordinate  $y_0$  in right-to-left order.

As before, we can assume integer coordinates and let  $N$  be the maximum  $y$ -coordinate of any line segment in  $L$ .


**Lemma 9.** *We can build a data structure that can answer largest-segment queries in  $O(n + N)$  space and  $O(k)$  time.*

**Proof.** We build the same data structure as in Lemma 3, but keep the partially persistent linked list sorted in decreasing order. The rest follows as before.  $\square$

Now, using this data structure in the solution described in Section 3, we immediately get an analogous result for  $\text{Srrfco}$ :

**Lemma 10.** *Given a string  $S$  of length  $n$ , we can build a data structure that can answer top- $k$  far consecutive occurrences queries using  $O(n \log n)$  space and  $O(m + k)$  query time.*

*Modifications to the linear space data structure* Now we extend the cluster solutions from Sections 5 and 6. We build the same recursive clusters as in Section 5. For each spine of a cluster of size  $\tau_i$ , we keep the line segments corresponding to the  $\tau_i$  consecutive occurrences of largest distance at every depth on the spine. That is, if a consecutive occurrence  $(i, j)$  is among the  $k$  farthest within  $D(v_d)$  for some  $v_d$  on the spine, define line segments for all maximal consecutive intervals  $[d_1, d_2]$  such that  $(i, j)$  is amongst the  $k$  farthest within  $D(v_d)$  for any  $d \in [d_1, d_2]$ . Again, the  $x$ -coordinate of the line segment is the distance  $j - i$ , and the  $y$ -coordinate spans  $[d_1, d_2]$ . Note that in this case, a consecutive occurrence might define more than one line segment. See Fig. 5 for an illustration of pair defining more than one line segment. We store these line segments in the data structure from Lemma 9.



**Fig. 5.** Illustration of a pair defining more than one line segment. To the left are the positions of the occurrences in  $S$ , in the middle is the spine of a cluster and to the right are the corresponding line segments. The pair  $(i, j)$  is amongst the  $k$  farthest until the occurrence  $x$  disappears, after which it is pushed out by the pair  $(a, b)$ . When  $b$  then disappears,  $(i, j)$  is again amongst the  $k$  farthest.

*Space analysis* When moving down a spine from  $v_{d-1}$  to  $v_d$ , only three different types of changes can happen to the set of the  $k$  farthest consecutive occurrences. We again denote  $D(v_d)$  to be the set of all consecutive occurrences of  $\text{str}(v_d)$ . The possible types of changes are then as follows.

- A consecutive occurrence can be removed from the  $k$  farthest because a consecutive occurrence of larger distance is added to  $D(v_d)$ . The consecutive occurrence of larger distance can only appear if an occurrence in between branched off. This leaf accounts for this change. A leaf can account for at most one such change, which triggers a line segment ending and a new line segment appearing at depth  $d$ .
- A consecutive occurrence  $(i, j)$  can disappear because either  $i$  or  $j$  branched off. Then this leaf accounts for this change. A leaf can account for at most two such changes.
- A consecutive occurrence that was present in  $D(v_{d-1})$  but not amongst the  $k$  farthest can be added to the  $k$  farthest in  $D(v_d)$ . This can only happen if a consecutive occurrence of greater distance disappeared because one or both of its occurrences branched off. Then this leaf accounts for this new line segment also, additional to the change of the disappearing consecutive occurrence(s). A leaf can account for at most two such changes.

In total, any leaf can account for at most a constant number of changes. Thus, we get the same space complexities as in Section 5.

*Algorithm* To answer a query we proceed as in Section 5. We first find  $\text{locus}(P)$ . If it is on a spine of a cluster of size  $O(\tau_i)$  and  $k < \tau_i$ , we query the line segment data structure to report the top- $k$  far consecutive occurrences. Otherwise, we find all occurrences of  $P$  in text order, construct the consecutive occurrences and use linear time selection to output the  $k$  consecutive occurrences of largest distance. This is correct by the same arguments as Section 5, and by similar arguments, achieves the same time complexities. The rank space reduction from Section 6 can be applied analogously. This gives us the following result:

**Theorem 11.** *Given a string  $S$  of length  $n$  and  $\epsilon$ ,  $0 < \epsilon \leq 1$ , we can build a data structure that can answer top- $k$  far consecutive occurrences queries using either*

- (i)  $O(n \log n)$  space and  $O(m + k)$  query time or
- (ii)  $O(\frac{n}{\epsilon})$  space and  $O(m + k^{1+\epsilon})$  query time.

Here,  $m$  is the length of the query pattern.

*Remark* We note that the construction from Section 7 does not generalize to the top- $k$  far consecutive occurrences problem, since the corresponding version of Claim 8 does not hold. A top- $k$  far consecutive occurrence of  $P$ , where both occurrences are below  $b$  is not necessarily among the top- $k$  far consecutive occurrences of  $\text{str}(b)$ . A consecutive occurrence of  $\text{str}(b)$  can be split by an occurrence of  $P$  not below  $b$ . This gives two smaller occurrences, and might cause the  $(k + 1)$ th furthest occurrence below  $b$  to be among the top- $k$  far occurrences of  $P$ . Each occurrence of  $P$  from within the cluster can split a consecutive occurrence below  $b$ , and thus we would need to store  $\Theta(\tau)$  occurrences below  $b$ , which no longer gives a linear space solution.

### 8.2. Consecutive occurrences with gaps

Given a string  $S$  the *string indexing for consecutive occurrences with gaps problem* (SICOG) is to preprocess  $S$  into a compact data structure, such that for any pattern  $P$  and a range  $[\alpha, \beta]$  we can efficiently find all consecutive occurrences of  $P$  where the distance lies within  $[\alpha, \beta]$ . The SICOG problem was considered by Navarro and Thankachan [33] and they give an  $O(n \log n)$  space and  $O(m + \text{occ})$  time solution, where  $\text{occ}$  is the number of consecutive pairs with distance in  $[\alpha, \beta]$ . Using the data structure from Section 3, we get an  $O(n \log n)$  space and  $O(m + \log n + \text{occ})$  time solution for the SICOG problem, which can be optimized using the same strategy as in [33] to achieve the same complexities. However, for a special case of the problem where either  $\alpha$  or  $\beta$  is known at indexing time we can get a similar trade-off as for the SRRCO problem: We first describe our solution for the fixed- $\alpha$  variant using the techniques from Sections 5 and 6, and then the fixed- $\beta$  variant follows by applying the same ideas combined with the data structure from Section 8.1.

**Data structure** We build the same data structure as in Section 5, with a slight modification. In the line segment data structure stored at every spine, instead of storing the  $\tau_i$  closest pairs, we store the  $\tau_i$  closest pairs that have distance  $\geq \alpha$ . This clearly occupies no more space than the solution from Section 5 and we can still apply the space optimizations of Section 6.

**Algorithm** Given  $P$  and  $\beta$ , we can now answer a query as follows. We begin by finding  $\text{locus}(P)$ . If it is in a subtree of constant size, we construct all the consecutive occurrences of  $P$  and report those that have distance within  $[\alpha, \beta]$ . If it is on a spine of a cluster of size  $\tau_i$ , we query the line segment data structure. For every consecutive occurrence we find, we check if the distance is  $\leq \beta$ . If we encounter a pair with distance  $> \beta$ , we stop reporting. If all  $\tau_i$  consecutive occurrences at  $\text{locus}(P)$  have distance  $\leq \beta$ , we then find all the consecutive occurrences of  $P$ , just as in Section 5, and scan them once to report all the consecutive occurrences with distance in  $[\alpha, \beta]$ .

For the analysis, define a *relevant pair* to be a consecutive occurrence with a distance in  $[\alpha, \beta]$ . If there are less than  $\tau_i$  relevant pairs for any locus, then they will all be stored and reported by the line segment data structure. As they are stored in order of increasing distance, once we reach a pair with distance  $> \beta$ , no further relevant pairs exist. If there are more than  $\tau_i$  relevant pairs, then we consider all occurrences of the pattern and report from those. Thus we always answer the query correctly.

If there is a consecutive occurrence among the  $\tau_i$  line segments with distance  $> \beta$ , we spend time  $O(m + \text{occ})$  finding the locus and querying the line segment data structure. Otherwise we have that  $\text{occ} \geq \tau_i$ , and thus  $\text{occ}^{1+\epsilon} \geq \tau_{i-1}$ . As before,  $P$  has at most  $\tau_{i-1}$  occurrences. Therefore, by the same arguments as in the previous section, we get the following result:

**Theorem 12.** *Given a string  $S$  of length  $n$  and  $\alpha > 0$ , we can build for any  $\epsilon$  satisfying  $0 < \epsilon \leq 1$  an  $O(\frac{n}{\epsilon})$  space data structure that can answer the following query in  $O(m + \text{occ}^{1+\epsilon})$  time: For a query pattern  $P$  and  $\beta \geq \alpha$ , report all consecutive occurrences of  $P$  in  $S$  where the distance lies in  $[\alpha, \beta]$ . Here,  $m$  is the length of the pattern and  $\text{occ}$  is the number of reported occurrences.*

By combining the same arguments with the solution for top- $k$  far consecutive occurrences, we get the following result for  $\beta$  fixed at indexing time:

**Theorem 13.** *Given a string  $S$  of length  $n$  and  $\beta > 0$ , we can build for any  $\epsilon$  satisfying  $0 < \epsilon \leq 1$  an  $O(\frac{n}{\epsilon})$  space data structure that can answer the following query in  $O(m + \text{occ}^{1+\epsilon})$  time: For a query pattern  $P$  and  $\alpha$  where  $0 < \alpha \leq \beta$ , report all consecutive occurrences of  $P$  in  $S$  where the distance lies in  $[\alpha, \beta]$ . Here,  $m$  is the length of the pattern and  $\text{occ}$  is the number of reported occurrences.*

We note that the construction from Section 7 does not generalize to the problem where  $\alpha$  is fixed, as the corresponding version of Claim 8 do not hold in this case. A consecutive occurrence below  $b$  of distance at least  $\alpha$  can be split by an occurrence of  $P$  from within the cluster introducing two new consecutive occurrences that might both have distance less than  $\alpha$ .

We can, however, get a solution to the problem when  $\alpha = 1$ . Using the data structure from Section 7, we can get all consecutive occurrences that are at most  $\beta$  apart as follows. We query the data structure for the top- $k$  close consecutive occurrences with  $k = 1, 2, 4, \dots$ , each time checking if all the top- $k$  close consecutive occurrences have distance at most  $\beta$ . As soon as we find a consecutive occurrence that has a distance more than  $\beta$  among our top- $k$  close consecutive occurrences, we stop and report all the occurrences found in this last call that have distance at most  $\beta$ . This way we ensure that  $k \leq 2 \cdot \text{occ}$ . We only find the locus once. The total query time is  $O(m + \sum_{i=0}^{\lceil \log \text{occ} \rceil} 2^i \log^{1+\epsilon} n) = O(m + \text{occ} \log^{1+\epsilon} n)$ .

**Lemma 14.** *Given a string  $S$  of length  $n$  we can build an  $O(n)$  space data structure that can answer the following query in  $O(m + \text{occ} \log^{1+\epsilon} n)$  time: For a query pattern  $P$  and  $\beta > 0$ , report all consecutive occurrences with distance at most  $\beta$ .*

**Non-overlapping consecutive occurrences** A natural and well-studied variant of string indexing is the problem of finding sets of *non-overlapping* occurrences of a pattern  $P$ . Here, a set of non-overlapping occurrences is a set of occurrences  $\{i_1, \dots, i_k\}$  of  $P$  such that the distance between any two of them is at least  $|P|$ . Several papers study the problem of finding the set

of non-overlapping occurrences of maximum size [26,15,19,24]. Note that Theorem 13 applied to  $\alpha = |P|$  solves a different variant of finding sets of non-overlapping occurrences: Namely, finding all pairs of non-overlapping *consecutive* occurrences. We call this problem the *string indexing for non-overlapping consecutive occurrences problem* (SINOCO). The SINOCO problem is inherently different from finding the maximum set of non-overlapping occurrences: For example, the maximum set of non-overlapping occurrences of the pattern  $P = \text{NANA}$  in the string  $S = \text{NANANANA}$  has size 2. However, there are no non-overlapping *consecutive* occurrences. To the best of our knowledge, the SINOCO problem has not been studied before. An immediate corollary of the results in Navarro and Thankachan [33] and Theorem 13 gives the following trade-offs for solving SINOCO:

**Corollary 15.** *Given a string  $S$  of length  $n$  and  $\epsilon$ ,  $0 < \epsilon \leq 1$ , we can build a data structure that can find all non-overlapping consecutive occurrences of a query pattern  $P$  using either*

- (i)  $O(n \log n)$  space and  $O(m + \text{occ})$  query time or
- (ii)  $O(\frac{n}{\epsilon})$  space and  $O(m + \text{occ}^{1+\epsilon})$  query time.

Here,  $m$  is the length of the query pattern and  $\text{occ}$  is the number of reported occurrences.

**Proof.** Apply the results in [33] and Theorem 13 with  $\beta = n$  and  $\alpha = |P|$ .  $\square$

## 9. Conclusion and open problems

We have introduced the natural problem of string indexing for top- $k$  close consecutive occurrences, and have given both a near-linear space solution achieving optimal query time and a linear space solution achieving a query time that is close to optimal. Using these techniques, we have given new solutions for the problem of string indexing for consecutive occurrences with gaps (SICOG). Furthermore, we have introduced the problem of finding all non-overlapping consecutive occurrences of a pattern (SINOCO) and showed that it can be reduced to a special case of SICOG.

These results open interesting new directions for further research. The most obvious open problem is to see whether it is possible to further improve the results for the main problem considered in this paper, especially, achieve linear space and optimal query time simultaneously. Secondly, it is still open whether it is possible to get an  $O(m + \text{occ})$  time and linear space solution for the special case of the SICOG problem where one of the interval endpoints is fixed, or even  $o(n \log n)$  space for the general problem. For the SINOCO problem, one might find better solutions that do not reduce it to SICOG but use additional insights about the specific structure of the problem.

## Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

## Acknowledgements

We thank anonymous reviewers of an earlier draft of this paper for their insightful comments and suggestions for improvement. Philip Bille, Inge Li Gørtz, Max Pedersen and Eva Rotenberg were partially supported by the Danish Research Council grant *Adaptive Compressed Computation* (DFF-8021-002498).

## References

- [1] Stephen Alstrup, Jacob Holm, Kristian de Lichtenberg, Mikkel Thorup, Minimizing diameters of dynamic trees, in: Proc. 24th ICALP, 1997, pp. 270–280.
- [2] Stephen Alstrup, Jacob Holm, Mikkel Thorup, Maintaining center and median in dynamic trees, in: Proc. 7th SWAT, 2000, pp. 46–56.
- [3] Stephen Alstrup, Theis Rauhe, Improved labeling scheme for ancestor queries, in: Proc. 13th SODA, 2002, pp. 947–953.
- [4] Johannes Bader, Simon Gog, Matthias Petri, Practical variable length gap pattern matching, in: Proc. 15th SEA, 2016, pp. 1–16.
- [5] Philip Bille, Inge Li Gørtz, The tree inclusion problem: in linear space and faster, ACM Trans. Algorithms 7 (3) (2011) 1–47.
- [6] Philip Bille, Inge Li Gørtz, Substring range reporting, Algorithmica 69 (2) (2014) 384–396.
- [7] Philip Bille, Inge Li Gørtz, Max Rishøj Pedersen, Teresa Anna Steiner, Gapped indexing for consecutive occurrences, in: Proc. 32nd CPM, 2021, pp. 10:1–10:19.
- [8] Philip Bille, Inge Li Gørtz, Teresa Anna Steiner, Eva Rotenberg, Max Rishøj Pedersen, String indexing for top- $k$  close consecutive occurrences, in: Proc. 40th FSTTCS, 2020.
- [9] Philip Bille, Inge Li Gørtz, Hjalte Wedel Vildhøj, Søren Vind, String indexing for patterns with wildcards, Theory Comput. Syst. 55 (1) (2014) 41–60.
- [10] Sudip Biswas, Arnab Ganguly, Rahul Shah, Sharma V. Thankachan, Ranked document retrieval for multiple patterns, Theor. Comput. Sci. 746 (2018) 98–111.
- [11] Manuel Blum, Robert W. Floyd, Vaughan Pratt, Ronald L. Rivest, Robert E. Tarjan, Time bounds for selection, J. Comput. Syst. Sci. 7 (4) (1973) 448–461.
- [12] Gerth Stølting Brodal, Rolf Fagerberg, Mark Greve, Alejandro López-Ortiz, Online sorted range reporting, in: Proc. 30th ISAAC, 2009, pp. 173–182.
- [13] Manuel Cáceres, Simon J. Puglisi, Bella Zhukova, Fast indexes for gapped pattern matching, in: Proc. 46th SOFSEM, 2020, pp. 493–504.
- [14] Timothy M. Chan, Persistent predecessor search and orthogonal point location on the word RAM, ACM Trans. Algorithms 9 (3) (2013) 1–22.
- [15] Hagai Cohen, Ely Porat, Range non-overlapping indexing, in: Proc. 20th ISAAC, 2009, pp. 1044–1053.

- [16] James R. Driscoll, Neil Sarnak, Daniel Dominic Sleator, Robert Endre Tarjan, Making data structures persistent, *J. Comput. Syst. Sci.* 38 (1) (1989) 86–124.
- [17] Greg N. Frederickson, Ambivalent data structures for dynamic 2-edge-connectivity and  $k$  smallest spanning trees, *SIAM J. Comput.* 26 (2) (1997) 484–538.
- [18] Michael L. Fredman, János Komlós, Endre Szemerédi, Storing a sparse table with  $O(1)$  worst case access time, *J. ACM* 31 (3) (1984) 538–544.
- [19] Arnab Ganguly, Rahul Shah, Sharma V. Thankachan, Succinct non-overlapping indexing, *Algorithmica* 82 (1) (2020) 107–117.
- [20] Wing-Kai Hon, Manish Patil, Rahul Shah, Sharma V. Thankachan, Jeffrey Scott Vitter, Indexes for document retrieval with relevance, in: *Space-Efficient Data Structures, Streams, and Algorithms - Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday*, 2013, pp. 351–362.
- [21] Wing-Kai Hon, Manish Patil, Rahul Shah, Shih-Bin Wu, Efficient index for retrieving top- $k$  most frequent documents, *J. Discret. Algorithms* 8 (4) (2010) 402–417.
- [22] Wing-Kai Hon, Rahul Shah, Sharma V. Thankachan, Jeffrey Scott Vitter, Space-efficient frameworks for top- $k$  string retrieval, *J. ACM* 61 (2) (2014) 1–36, Announced at 50th FOCS.
- [23] Wing-Kai Hon, Sharma V. Thankachan, Rahul Shah, Jeffrey Scott Vitter, Faster compressed top- $k$  document retrieval, in: *Proc. 23rd DCC*, 2013, pp. 341–350.
- [24] Sahar Hooshmand, Paniz Abedin, M. Oguzhan Külekci, Sharma V. Thankachan, Non-overlapping indexing - cache obliviously, in: *Proc. 29th CPM*, 2018, pp. 8:1–8:9.
- [25] Costas S. Iliopoulos, M. Sohel Rahman, Indexing factors with gaps, *Algorithmica* 55 (1) (2009) 60–70.
- [26] Orgad Keller, Tsvi Kopelowitz, Moshe Lewenstein, Range non-overlapping indexing and successive list indexing, in: *Proc. 11th WADS*, 2007, pp. 625–636.
- [27] Moshe Lewenstein, Indexing with gaps, in: *Proc. 18th SPIRE*, 2011, pp. 135–143.
- [28] J. Ian Munro, Gonzalo Navarro, Jesper Sindahl Nielsen, Rahul Shah, Sharma V. Thankachan, Top- $k$  term-proximity in succinct space, *Algorithmica* 78 (2) (2017) 379–393, Announced at 25th ISAAC.
- [29] J. Ian Munro, Gonzalo Navarro, Rahul Shah, Sharma V. Thankachan, Ranked document selection, *Theor. Comput. Sci.* 812 (2020) 149–159.
- [30] Gonzalo Navarro, Spaces, trees, and colors: the algorithmic landscape of document retrieval on sequences, *ACM Comput. Surv.* 46 (4) (2014) 1–47.
- [31] Gonzalo Navarro, Yakov Nekrich, Time-optimal top- $k$  document retrieval, *SIAM J. Comput.* 46 (1) (2017) 80–113, Announced at 23rd SODA.
- [32] Gonzalo Navarro, Sharma V. Thankachan, New space/time tradeoffs for top- $k$  document retrieval on sequences, *Theor. Comput. Sci.* 542 (2014) 83–97, Announced at 20th SPIRE.
- [33] Gonzalo Navarro, Sharma V. Thankachan, Reporting consecutive substring occurrences under bounded gap constraints, *Theor. Comput. Sci.* 638 (2016) 108–111, Announced at 26th CPM.
- [34] Yakov Nekrich, Gonzalo Navarro, Sorted range reporting, in: *Proc 13th SWAT*, 2012, pp. 271–282.
- [35] Rahul Shah, Cheng Sheng, Sharma V. Thankachan, Jeffrey Scott Vitter, Top- $k$  document retrieval in external memory, in: *Proc. 21st ESA*, 2013, pp. 803–814.
- [36] Daniel D. Sleator, Robert Endre Tarjan, A data structure for dynamic trees, *J. Comput. Syst. Sci.* 26 (3) (1983) 362–391.
- [37] Dekel Tsur, Top- $k$  document retrieval in optimal space, *Inf. Process. Lett.* 113 (12) (2013) 440–443.
- [38] Peter Weiner, Linear pattern matching algorithms, in: *Proc. 14th FOCS*, 1973, pp. 1–11.