

Binary Jumbled Pattern Matching on Trees and Tree-Like Structures

Travis Gagie · Danny Hermelin ·
Gad M. Landau · Oren Weimann

Received: 12 September 2013 / Accepted: 11 November 2014 / Published online: 27 November 2014
© Springer Science+Business Media New York 2014

Abstract Binary jumbled pattern matching asks to preprocess a binary string S in order to answer queries (i, j) which ask for a substring of S that is of length i and has exactly j 1-bits. This problem naturally generalizes to vertex-labeled trees and graphs by replacing “substring” with “connected subgraph”. In this paper, we give an $O(n^2 / \log^2 n)$ -time solution for trees, matching the currently best bound for (the simpler problem of) strings. We also give an $O(g^{2/3} n^{4/3} / (\log n)^{4/3})$ -time solution for strings that are compressed by a context-free grammar of size g in Chomsky normal form. This solution improves the known bounds when the string is compressible under many popular compression schemes. Finally, we prove that on graphs the problem is fixed-parameter tractable with respect to the treewidth w of the graph, even for a

Gad M. Landau: Supported in part by the National Science Foundation (NSF) Grant 0904246, the Israel Science Foundation (ISF) Grant 347/09, and the United States-Israel Binational Science Foundation (BSF) Grant 2008217. Oren Weimann: Supported in part by the Israel Science Foundation Grant 794/13.

Preliminary version of this paper appeared in the 21st Annual European Symposium on Algorithms (ESA 2013).

T. Gagie
University of Helsinki, Helsinki, Finland
e-mail: travis.gagie@aalto.fi

D. Hermelin
Ben-Gurion University, Beersheba, Israel
e-mail: hermelin@bgu.ac.il

G. M. Landau · O. Weimann (✉)
University of Haifa, Haifa, Israel
e-mail: oren@cs.haifa.ac.il

G. M. Landau
e-mail: landau@cs.haifa.ac.il

constant number of different vertex-labels, thus improving the previous best $n^{O(w)}$ algorithm.

Keywords Pattern matching · Tree pattern matching · Permutation pattern matching · Grammar compression · Graph motifs

1 Introduction

Jumbled pattern matching is an important variant of classical pattern matching with several applications in computational biology, ranging from alignment [5] and SNP discovery [7], to the interpretation of mass spectrometry data [10] and metabolic network analysis [26]. In the most basic case of strings, the problem asks to determine whether a given pattern P can be rearranged so that it appears in a given text T . That is, whether T contains a substring of length $|P|$ where each letter of the alphabet occurs the same number of times as in P . Using a straightforward sliding window algorithm, such a jumbled occurrence can be found optimally in $O(n)$ time on a text of length n . While jumbled pattern matching has a simple efficient solution, its *indexing* problem is much more challenging. In the indexing problem, we preprocess a given text T so that on queries P we can determine quickly whether T has a jumbled occurrence of P . Very little is known about this problem besides the trivial naive solution.

Most of the interesting results on indexing for jumbled pattern matching relate to binary strings (where a query pattern (i, j) asks for a substring of T that is of length i and has j 1s). Given a binary string of length n , Cicalese, Fici and Lipták [14] showed how one can build in $O(n^2)$ time an $O(n)$ -space index that answers jumbled pattern matching queries in $O(1)$ time. Their key observation was that if one substring of length i contains fewer than j 1s, and another substring of length i contains more than j 1s, then there must be a substring of length i with exactly j 1s. Using this observation, they construct an index that stores the maximum and minimum number of 1s in any i -length substring, for each possible i . Burcsi et al. [10] (see also [11, 12]) and Moosa and Rahman [27] independently improved the construction time to $O(n^2 / \log n)$, then Moosa and Rahman [28] further improved it to $O(n^2 / \log^2 n)$ in the word RAM model. Currently, faster algorithms than $O(n^2 / \log^2 n)$ exist only when the string compresses well using run-length encoding [4, 23] or when we are willing to settle for approximate indices [16]. Regarding non-binary alphabets, the recent solution of Kociumaka et al. [25] for constant alphabets requires $o(n^2)$ space and $o(n)$ query time. For general alphabets, expected sublinear query time was achieved by Burcsi et al. [11] for large query patterns but in the worst case a query takes superlinear time. In fact, a recent result of Amir et al. [3] shows that under a 3SUM hardness assumption, jumbled indexing for alphabets of size $\omega(1)$ requires either $\Omega(n^{2-\varepsilon})$ preprocessing time or $\Omega(n^{1-\delta})$ query time for any $\varepsilon, \delta > 0$.

The natural extension of jumbled pattern matching from strings to trees is much harder. In this extension, we are asked to determine whether a vertex-labeled input tree has a connected subgraph where each label occurs the same number of times as specified by the input query. The difficulty here stems from the fact that a tree can have an exponential number of connected subgraphs as opposed to strings. Hence, a sliding

window approach becomes intractable. Indeed, the problem is NP-hard [26], even if our query contains at most one occurrence of each letter [21]. It is not even fixed-parameter tractable when parameterized by the alphabet size [21]. The fixed-parameter tractability of the problem was further studied when extending the problem from trees to graphs [2, 6, 17, 18]. In particular, the problem (also known as the *graph motif problem*) was recently shown by Fellows et al. [21] to be polynomial-time solvable when the number of letters in the alphabet as well as the treewidth of the graph are both fixed. They also gave a fixed-parameter algorithm when the size of the pattern is taken as a parameter, and showed that no such algorithm is likely to exist when the problem is parameterized by the alphabet size, even when the input graph is a tree.

Our results In this paper we extend the currently known state-of-the-art for binary jumbled pattern matching. Our results focus on trees, and tree-like structures such as grammars and bounded treewidth graphs. We use the word RAM model of computation with the standard assumption that the word-length is at least $\log n$.

- *Trees* For a tree T of size n , we present an index of size $O(n)$ bits that is constructed in $O(n^2/\log^2 n)$ time and answers binary jumbled pattern matching queries in $O(1)$ time. This matches the performance of the best known index for binary strings. In fact, our index for trees is obtained by multiple applications of an efficient algorithm for strings [28] under a more careful analysis. This is combined with both a micro–macro [1] and centroid decomposition of the input tree. Our index can also be used as an $O(ni/\log^2 n)$ -time algorithm for the pattern matching (as opposed to the indexing) problem, where i denotes the size of the pattern. Finally, by increasing the space of our index to $O(n \log n)$ bits, we can output in $O(\log n)$ time a node of T that is part of the pattern occurrence.
- *Grammars* For a binary string S of length n derived by a context free grammar (CFG) of size g in Chomsky normal form, we show how to construct in $O(g^{2/3}n^{4/3}/\log^{4/3} n)$ time an index of size $O(n)$ bits that answers jumbled pattern matching queries on S in $O(1)$ time. The size of the grammar g can be exponentially smaller than n and is always at most $O(n/\log n)$. This means that our time bound is $O(n^2/\log^2 n)$ even when S is not compressible. If S is compressible but with other compression schemes such as the LZ-family, then we can transform it into a grammar-based compression with little or no expansion [13, 30].
- *Bounded Treewidth Graphs:* For a graph G with treewidth bounded by w , we show how to improve on the $O(n^{O(w)})$ time algorithm of Fellows et al. [21] to an algorithm which runs in $2^{O(w^3)}n + w^{O(w)}n^{O(1)}$ time. Thus, we show that for a binary alphabet, jumbled pattern matching is fixed-parameter tractable when parameterized only by the treewidth. This result extends easily to alphabets of constant sizes.

We present our results for trees, grammars, and bounded treewidth graphs in Sects. 2, 3 and 4 respectively.

2 Jumbled Pattern Matching on Trees

In this section we consider the natural extension of binary jumbled pattern matching to trees. Recall that in this extension we are given a tree T with n nodes, where each

node is labeled by either 1 or 0. We will refer to the nodes labeled 1 as *black* nodes, and the nodes labeled 0 as *white* nodes. Our goal is to construct a data structure that on query (i, j) determines whether T contains a connected subgraph with exactly i nodes, j of which are black. Such a subgraph of T is referred to as a *pattern* and (i, j) is said to *appear* in T . The main result of this section is stated below.

Theorem 1 *Given a tree T with n nodes that are colored black or white, we can construct in $O(n^2/\log^2 n)$ time a data structure of size $O(n)$ bits that given a query (i, j) determines in $O(1)$ time if (i, j) appears in T .*

Notice that the bounds of Theorem 1 match the currently best bounds for the case where T is a string [27,28]. This is despite the fact that a string has only $O(n^2)$ substrings while a tree can have $\Omega(2^n)$ connected subgraphs. The following lemma indicates an important property of string jumbled pattern matching that carries over to trees. It gives rise to a simple index described below.

Lemma 1 *If (i, j_1) and (i, j_2) both appear in T , then for every $j_1 \leq j \leq j_2$, (i, j) appears in T .*

Proof Let j be an arbitrary integer with $j_1 \leq j \leq j_2$, and let T_1 and T_2 be two patterns in T corresponding to (i, j_1) and (i, j_2) respectively. The lemma follows from the fact that there exists a sequence of patterns starting with T_1 and ending with T_2 such that every pattern has exactly i nodes and two consecutive patterns differ by removing a leaf from the first pattern and adding a different node instead. This means that the numbers of black nodes in two consecutive patterns differ by at most 1. \square

2.1 A Simple Index

As in the case of strings, the above lemma suggests an $O(n)$ -size data structure: For every $i = 1, \dots, n$, store the minimum and maximum values i_{min} and i_{max} such that (i, i_{min}) and (i, i_{max}) appear in T . This way, upon query (i, j) , we can report in constant time whether (i, j) appears in T by checking if $i_{min} \leq j \leq i_{max}$. However, while $O(n^2)$ construction-time is trivial for strings (for every $i = 0, \dots, n$, slide a window of length i through the text in $O(n)$ time) it is harder on trees.

To obtain $O(n^2)$ construction time, we begin by converting our tree into a rooted ordered binary tree. We arbitrarily root the tree T . To convert it to a binary tree, we duplicate each node with more than two children as follows: Let v be a node with children $u_1, \dots, u_k, k \geq 3$. We replace v with $k - 1$ new nodes v_1, \dots, v_{k-1} , make u_1 and u_2 be the left and right children of v_1 respectively, and for each $\ell = 2, \dots, k - 1$ we make $v_{\ell-1}$ and $u_{\ell+1}$ be the left and right children of v_ℓ respectively. If v is not the root then we set the parent of v_{k-1} to be the parent of v (otherwise, v_{k-1} is the root). The node v_1 gets the same color as the the node v . The other nodes v_2, \dots, v_k are called *dummy nodes* and have no color. This procedure at most doubles the size of T . To avoid cumbersome notation, we henceforth use T and n to denote the resulting rooted binary tree and its number of nodes respectively. For a node v , we let T_v denote the subtree of T rooted at v (*i.e.* the connected subgraph induced by v and all its descendants).

Next, in a bottom-up fashion, we compute for each node v of T an array A_v of size $|T_v| + 1$.

For $i > 0$, $A_v[i]$ will store the maximum number of black nodes that appear in a connected subgraph of size i that (1) includes i non-dummy nodes in T_v , (2) includes v , and (3) includes dummy nodes such that for every included dummy node its corresponding non-dummy node must also be included. The entry $A_v[0]$ is set to zero. Computing the minimum (rather than maximum) number of black nodes is done similarly.

Throughout the execution, we also maintain a global array A such that $A[i]$ stores the maximum $A_v[i]$ over all nodes v considered so far. Notice that at the end of the execution, $A[i]$ holds the desired value i_{max} since every connected subgraph of T of size i includes some node v and $i - 1$ nodes in T_v .

We now show how to compute $A_v[i]$ for a node v and a specific value $i \in \{1, \dots, |T_v|\}$. If v has a single child u , then v is necessarily not a dummy node and we set $A_v[i] = col(v) + A_u[i - 1]$, where $col(v) = 1$ if v is black and $col(v) = 0$ otherwise. If v has two children u and w , then any pattern of size i that appears in T_v and includes v is composed of v , a pattern of size ℓ in T_u that includes u , and a pattern of size $i - 1 - \ell$ in T_w that includes w . Notice that we allow $\ell = 0$ which corresponds to a pattern that does not include u , and we allow $i - 1 - \ell = 0$ which corresponds to a pattern that does not include w . We therefore set $A_v[i] = col(v) + \max_{0 \leq \ell \leq i-1} \{A_u[\ell] + A_w[i - 1 - \ell]\}$ and $A_v[i] = \max_{1 \leq \ell \leq i-1} \{A_u[\ell] + A_w[i - 1 - \ell]\}$ when v is a dummy node. Observe that in the latter the index ℓ starts with 1 to indicate that the non-dummy copy of v (i.e., v_1) must be included in the pattern.

We next analyze and then improve the above algorithm (first by one log factor and then by another log factor). In the rest of this section, like the above algorithm, all of our algorithms will compute the A_v arrays for each v in T in a bottom-up fashion. In all these algorithms, just like in the above algorithm, special attention has to be given to the case where v is a dummy node. Handling dummy nodes is done similarly to the above. To make the presentation simpler we will assume that there are no dummy nodes at all.

Lemma 2 *The above algorithm runs in $O(n^2)$ time.*

Proof The computation done on nodes with one child requires $O(n)$ time, hence the total time required to compute all arrays A_v for such nodes is $O(n^2)$. The time required to compute all arrays for nodes with two children is asymptotically bounded by the sum $\sum_v \alpha(v)\beta(v)$, where $\alpha(v)$ and $\beta(v)$ denote the sizes of the two subtrees rooted at each of the children of v , and the sum is taken over all nodes v with two children. For a tree rooted at r , we let $cost(r)$ denote this sum over all nodes in T_r and argue by induction that $cost(r)$ is bounded by $|T_r|^2 = O(n^2)$.

Let r be the root of a tree with n nodes, and let u and v denote the two children of r . Let x denote the size of the subtree rooted at u . Then $x < n$, and the size of the subtree rooted at v is $n - 1 - x$. By induction, we have $cost(u) \leq x^2$ and $cost(v) \leq (n - 1 - x)^2$. Thus, $cost(r) = x(n - 1 - x) + cost(u) + cost(v) < n^2 - x(n - x) \leq n^2$. \square

Note that if at any time the algorithm only stores arrays A_v which are necessary for future computations, then the total space used by the algorithm is $O(n)$. The space can be made $O(n)$ bits by storing the A_v arrays in a succinct fashion (this will also

prove useful later for improving the running time): Observe that $A_v[i + 1]$ is either equal to $A_v[i]$ or to $A_v[i] + 1$. This is because any pattern of size i with b black nodes can be turned into a pattern of size $i - 1$ with at least $b - 1$ black nodes by removing a leaf. We can therefore represent A_v as a binary string B_v of $n + 1$ bits, where $B_v[0] = 0$, and $B_v[i] = A_v[i] - A_v[i - 1]$ for all $i = 1, \dots, n$. Notice that since $A_v[i] = \sum_{\ell=0}^i B_v[\ell]$, each entry of A_v can be retrieved from B_v in $O(1)$ time using *rank* queries [24,29].

2.2 Pattern Matching

Before improving the above algorithm, we show that it can already be analyzed more carefully to get a bound of $O(n \cdot i)$ when the pattern size is known to be at most i . This means that in $O(n)$ space and $O(n \cdot i)$ construction time we can build an index that answer queries in $O(1)$ time provided the pattern size is bounded by i . It is also useful for the *pattern matching* problem: without preprocessing, decide whether a given pattern (i, j) appears in T .

In the case of strings, this problem can trivially be solved in $O(n)$ time by sliding a window of length i through the string thus effectively considering every substring of length i . This sliding-window approach however does not extend to trees since we cannot afford to examine all connected subgraphs of T . We next show that, in trees, searching for a pattern of size i can be done in $O(n \cdot i)$ time by using our above indexing algorithm. This is useful when the pattern is small (i.e., when $i = o(n)$). Obtaining $O(n)$ time remains our main open problem.

Lemma 3 *Given a tree T with n nodes that are colored black or white and a query pattern (i, j) , we can check in $O(n \cdot i)$ time and $O(n)$ space if T contains the pattern (i, j) .*

Proof In our indexing algorithm, every node v computes an array A_v of size $|T_v|$. When the pattern size is known to be i we can settle for an array A_v of size $\min\{|T_v|, i\}$. Recall from the above discussion that we can assume T is a binary tree. Consider some node v that has only one child u . We can compute A_v from A_u in time $O(\min\{|T_v|, i\}) = O(i)$. Summing over all such nodes v gives at most $O(n \cdot i)$. If on the other hand, node v has two children u and w then A_v is computed from A_u and A_w in $O(\min\{|T_u|, i\} \cdot \min\{|T_w|, i\})$ time. We claim that summing this term over all nodes in T that have two children gives $O(n \cdot i)$.

To see this, first consider the subset of nodes $V = \{v \in T : |T_v| < i \text{ and } |T_{\text{parent}(v)}| \geq i\}$, where $\text{parent}(v)$ denotes the parent of v in T . Notice that each subtree $T_v \in \{T_v : v \in V\}$ is of size less than i and that these subtrees are disjoint. By the proof of Lemma 2 we know that computing A_v (along with every A_u for vertices $u \in T_v$) is done in $O(|T_v|^2)$ time. The total time to compute A_v for all nodes $v \in V$ and their descendants is therefore $\text{cost}(v) = \sum_{v \in V} |T_v|^2$. Since every $|T_v| < i$ and $\sum_{v \in V} |T_v| \leq n$, we have that $\text{cost}(v)$ is upper bounded by $O(n \cdot i)$ that is achieved when all $|T_v|$ s are equal to i and $|V| = n/i$.

The remaining set of nodes S consists of all nodes v such that v has two children u, w and $|T_v| \geq i$. We partition these nodes into $S_1 = \{v \in S : |T_u| \geq i \text{ and } |T_w| \geq i\}$

and $S_2 = S \setminus S_1$. Notice that $|S_1| = O(n/i)$. Therefore, computing A_v for all nodes $v \in S_1$ can be done in $O(|S_1| \cdot i^2) = O(n \cdot i)$ time. We are left only with the vertices of S_2 . These are all vertices v such that at least one of their children is in V . Denote this child as $d(v)$. Computing A_v for all nodes in S_2 can therefore be done in time

$$\sum_{v \in S_2} O(|T_{d(v)}| \cdot i) = i \cdot \sum_{v \in S_2} O(|T_{d(v)}|) = i \cdot \sum_{u \in V} O(|T_u|) = O(i \cdot n).$$

□

2.3 An Improved Index

In this subsection, we will gradually improve the construction time from $O(n^2)$ to $O(n^2/\log^2 n)$. For simplicity of the presentation, we will assume the input tree T is a rooted binary tree. This extends to arbitrary trees using a similar dummy-nodes trick as above.

2.3.1 From Trees to Strings

Recall that we can represent every A_v by a binary string B_v of $n + 1$ bits where $B_v[0]$ is always zero and for $i = 1, \dots, n$, $B_v[i] = A_v[i] - A_v[i - 1]$. We begin by showing that if v has two children u, w then the computation of B_v can be done by solving a variant of jumbled pattern matching on the string $S_v = X_v \circ \text{col}(v) \circ Y_v$ (here \circ denotes concatenation) of length $|S_v| = |T_u| + |T_w| + 1$, where X_v is obtained from B_u by reversing it and removing its last bit, and Y_v is obtained from B_w by removing its first bit. We call the position in S_v with $\text{col}(v)$ the *split position* of S_v . Recall that $A_v[i] = \text{col}(v) + \max_{0 \leq \ell \leq i-1} \{A_u[\ell] + A_w[i - 1 - \ell]\}$. This is equal to the maximum number of 1s in a window of S_v that is of length i and includes the split position of S_v .

We are therefore interested only in windows including the split position, and this is the important distinction from the standard jumbled pattern matching problem on strings. Clearly, using the fastest $O(n^2/\log^2 n)$ -time algorithm [28] for the standard string problem we can also solve our problem and compute A_v in $O(|S_v|^2/\log^2 n)$ time. However, recall that for our total analysis (over all nodes v) to give $O(n^2/\log^2 n)$ we need the time to be $O(|X_v| \cdot |Y_v|/\log^2 n)$ and not $O((|X_v| + |Y_v|)^2/\log^2 n)$.

2.3.2 First Speedup

The $O(\log^2 n)$ -factor speedup for jumbled pattern matching on strings [28] is achieved by a clever combination of lookup tables. One log factor is achieved by computing the maximum number of 1s in a window of length i only when i is a multiple of $s = (\log n)/6$. Using a lookup table over all possible pairs of length- s windows and all values between 1 and s , a sliding window of size i can be extended in $O(1)$ time to all windows of sizes $i + 1, \dots, i + s - 1$ that start at the same location (see [28] for details). Their algorithm can output in $O(n^2/\log n)$ time an array of $O(n/\log n)$ words. For each i that is a multiple of s , the array keeps one word storing the maximum number of

1s over all windows of length i and another word storing the binary increment vector for the maximum number of 1s in all windows of length $i + 1, \dots, i + s - 1$.

By only considering windows that include the split position of S_v , this idea easily translates to an $O(|X_v| \cdot |Y_v| / \log n)$ -time algorithm to compute A_v and implicitly store it in $O((|X_v| + |Y_v|) / \log n)$ words. From this it is also easy to obtain an $O((|X_v| + |Y_v|) / \log n)$ -words representation of B_v . Notice that if v has a single child then the same procedure works with $|X_v| = 0$ in time $O(|Y_v| / \log n) = O(n / \log n)$. Summing over all nodes v , we get an $O(n^2 / \log n)$ -time solution for binary jumbled indexing on trees.

2.3.3 Second Speedup

In strings, an additional logarithmic improvement shown in [28] can be obtained as follows: When sliding a window of length i (i is a multiple of s) the window is shifted s locations in $O(1)$ time using a lookup table over all pairs of binary substrings of length s (representing the leftmost and rightmost bits in all these s shifts). This further improvement yields an $O(n^2 / \log^2 n)$ -time algorithm for strings. In trees however this is not the case. While we can compute A_v in $O((|X_v| + |Y_v|)^2 / \log^2 n)$ time, we can guarantee $O(|X_v| \cdot |Y_v| / \log^2 n)$ time only if both $|X_v|$ and $|Y_v|$ are greater than s . Otherwise, say $|X_v| < s$ and $|Y_v| \geq s$, we will get $O(|X_v| \cdot |Y_v| / |X_v| \log n) = O(|Y_v| / \log n)$ time. This is because our windows must include the $col(v)$ index and so we never shift a window by more than $|X_v|$ locations. Overcoming this obstacle is the main challenge of this subsection. It is achieved by carefully ensuring that the $O(|Y_v| / \log n) = O(n / \log n)$ costly constructions will be done only $O(n / \log n)$ times.

2.3.4 A Micro–Macro Decomposition


A *micro–macro decomposition* [1] is a partition of T into $O(n / \log n)$ disjoint connected subgraphs called *micro trees*. Each micro tree is of size at most $\log n$, and at most two nodes in a micro tree are adjacent to nodes in other micro trees. These nodes are referred to as *top* and *bottom boundary* nodes. The top boundary node is chosen as the root of the micro tree. The *macro tree* is a rooted tree of size $O(n / \log n)$ whose nodes correspond to micro trees as follows (See Fig. 1): The top boundary node $t(C)$ of a micro tree C is connected to a boundary node in the parent micro tree (apart from the root). The boundary node $t(C)$ might also be connected to a top boundary node of a child micro tree $child(C)$.¹ The bottom boundary node $b(C)$ of C is connected to top boundary nodes of at most two child micro trees $\ell(C)$ and $r(C)$ of C . Such a micro–macro decomposition can be found in linear time [1].

2.3.5 A Bottom up Traversal of the Macro Tree

With each micro tree C we associate an array A_C . Let T_C denote the union of micro tree C and all its descendant micro trees (including the edges between them). The array

¹ The root of the macro tree is an exception as it might have a top boundary node connected to two (rather than one) child micro trees. We focus on the other nodes. Handling the root is done in a very similar way.

Fig. 1 A micro tree C and its neighboring micro trees in the macro tree. Inside each micro tree, the *solid nodes* correspond to boundary nodes and the *hollow nodes* to non-boundary nodes



A_C stores the maximum number of 1s (black nodes) in every pattern that includes the boundary node $t(C)$ and other nodes of T_C . We also associate three auxiliary arrays: A_b , A_t and A_{tb} and set $A_b[0] = A_t[0] = A_{tb}[0] = A_{tb}[1] = 0$. The array A_b stores the maximum number of 1s in every pattern that includes the boundary node $b(C)$ and possibly other nodes of C , $T_{l(C)}$, and $T_{r(C)}$. The array A_t stores the maximum number of 1s in every pattern that includes the boundary node $t(C)$ and possibly other nodes of C and $T_{child(C)}$. Finally, the array A_{tb} stores the maximum number of 1s in every pattern that includes *both* boundary nodes $t(C)$ and $b(C)$ and possibly other nodes of C , $T_{l(C)}$, and $T_{r(C)}$.

We initialize for every micro tree C its $O(|C|) = O(\log n)$ sized arrays. Arrays A_C and A_t are initialized to hold the maximum number of 1s in every pattern that includes $t(C)$ and nodes of C . For each C , the required values are calculated in $O(|C|^2)$ time as a by-product of running the algorithm from the previous subsection after rooting C at $t(C)$. Similarly, we initialize the array A_b to hold the maximum number of 1s in every pattern that includes $b(C)$ and nodes of C . The array A_{tb} is initialized as follows: First we check how many nodes are 1s and how many are 0s on the unique path between $t(C)$ and $b(C)$. If there are i 1s and j 0s we set $A_{tb}[k] = 0$ for every $k < i + j$ and we set $A_{tb}[i + j] = i$. We compute $A_{tb}[k]$ for all $k > i + j$ in total $O(|C|^2)$ time by contracting the $b(C)$ -to- $t(C)$ path into a single node and running the previous algorithm rooting C in this contracted node. The total running time of the initialization step is therefore $O(n \cdot |C|^2 / \log n) = O(n \log n)$ which is negligible. Notice that during this computation we have computed the maximum number of 1s in all patterns that are completely inside a micro tree. We initialize the array A (that is, only the first $O(\log n)$ entries of A) with these values. In particular, this takes care of all patterns that do not contain any boundary node. We are now done with the leaf nodes of the macro tree.

We next describe how to compute the arrays of an internal node C of the macro tree given the arrays of $l(C)$, $r(C)$ and $child(C)$. We first compute the maximum number of 1s in all patterns that include $b(C)$ and possibly other vertices of $T_{l(C)}$ and $T_{r(C)}$. This can be done using the aforementioned string speedups in $O(|T_{l(C)}| \cdot |T_{r(C)}| / \log^2 n)$ time when both $|T_{l(C)}| > \log n$ and $|T_{r(C)}| > \log n$ and in $O(n / \log n)$ time otherwise. Using this and the initialized array A_b of C (that is of size $|C| \leq \log n$) we can compute the final array A_b of C . This is done by using the aforementioned string algorithm (on a string S of length $|T_{l(C)}| + |T_{r(C)}| + 1 + |C|$) restricted to the case where windows must include the split position (the split position separates S to a substring of length

$|T_{\ell(C)}| + |T_{r(C)}| + 1$ and a substring of length $|C| \leq \log n$). Using only the first speedup, this takes time $O(|S|/\log n) = O(n/\log n)$. Similarly, using the initialized A_{tb} of C , we can compute the final array A_{tb} of C in $O(n/\log n)$ time.

Next, we compute the array A_t using the initialized array A_t of C and the array A_t of $child(C)$ in time $O(n/\log n)$. Finally, we compute A_C of C using A_{tb} of C and A_t of $child(C)$ in $O((|T_{\ell(C)}| + |T_{r(C)}| + 1 + |C|) \cdot |T_{child(C)}|/\log^2 n)$ time if both $|T_{\ell(C)}| + |T_{r(C)}| + 1 + |C| > \log n$ and $|T_{child(C)}| > \log n$ and in $O(n/\log n)$ otherwise. To finalize A_C we must then take the entry-wise maximum between the computed A_C and A_t . This is because a pattern in T_C may or may not include $b(C)$. Finally, once A_C is computed, we update the global array A accordingly (by taking the entry-wise maximum between A and A_C).

To bound the total time complexity over all clusters C , notice that some computations required $O(\alpha(v) \cdot \beta(v)/\log^2 n)$ when $\alpha(v) > \log n$ and $\beta(v) > \log n$ are the subtree sizes of two children of some node $v \in T$. We have already seen that the sum of all these terms over all nodes of T is $O(n^2/\log^2 n)$. The other type of computations each require $O(n/\log n)$ time but there are at most $O(n/\log n)$ such computations ($O(1)$ for each micro tree) for a total of $O(n^2/\log^2 n)$. This completes the proof of Theorem 1.

2.4 Finding the Query Pattern

In this subsection we extend the index so that on top of identifying in $O(1)$ time if a pattern (i, j) appears in T , it can also locate in $O(\log n)$ time a node $v \in T$ that is part of such a pattern appearance. We call this node an *anchor* of the appearance. This extension increases the space of the index from $O(n)$ bits to $O(n \log n)$ bits (i.e., $O(n)$ words).

Recall that given a tree T we build in $O(n^2/\log^2 n)$ time an array A of size $n = |T|$ where $A[i]$ stores the minimum and maximum values i_{min} and i_{max} such that (i, i_{min}) and (i, i_{max}) appear in T . Now consider a *centroid decomposition*² of T : A centroid node c in T is a node whose removal leaves no connected component with more than $n/2$ nodes. We first construct the array A of T in $O(n^2/\log^2 n)$ time and store it in node c . We then recurse on each remaining connected component. This way, every node $v \in T$ will compute the array corresponding to the connected component whose centroid was v . Notice that this array is not the array A_v since we do not insist the pattern uses v . Observe that since each array A is implicitly stored in an n -sized bit array B , and since the recursion tree is balanced the total space complexity is $O(n \log n)$ bits. Furthermore, since every node in T has degree at most three, removing the centroid leaves at most three connected components and so the time to construct all the arrays is bounded by $T(n) = T(n_1) + T(n_2) + T(n_3) + O(n^2/\log^2 n)$ where $n_1 + n_2 + n_3 = n$ and every $n_i \leq n/2$. This yields the time complexity $T(n) = O(n^2/\log^2 n)$.

Let c denote the centroid of T whose removal leaves at most three connected components T_1 , T_2 , and T_3 (recall we assume degree at most 3). Upon query (i, j) we first check the array of c if pattern (i, j) appears in T (i.e., if $i_{min} \leq j \leq i_{max}$). If it does then we check the centroids of T_1 , T_2 and T_3 . If (i, j) appears in any of them then

² A centroid decomposition can be found in linear time.

we continue the search there. This way, after at most $O(\log n)$ steps we reach the first node v whose connected component includes (i, j) but none of its child components do. We return v as the anchor node since such a pattern must include v . We note that the above can be extended so that for every occurrence of (i, j) one node that is part of this occurrence is reported. Finally, we note that it was recently observed in [15] that if we are willing to settle for an index of size $O(n^2)$ then we can locate the entire match (not just an anchor) in time proportional to the size of the match.

3 Jumbled Pattern Matching on Grammars

In grammar-based compression, a binary string S of length n is compressed using a context-free grammar $G(S)$ in Chomsky normal form that generates S and only S . Such a grammar has a unique parse tree that generates S . Identical subtrees of this parse tree indicate substring repeats in S . The size of the grammar $g = |G(S)|$ is defined as the total number of variables and production rules in the grammar. Note that g can be exponentially smaller than $n = |S|$. We show how to solve the jumbled pattern matching problem on S by solving it on the parse tree of $G(S)$, taking advantage of subtree repeats. We obtain the following bounds:

Theorem 2 *Given a binary string S of length n compressed by a context free grammar $G(S)$ of size g in Chomsky normal form, we can construct in $O(g^{2/3}n^{4/3}/(\log n)^{4/3})$ time a data structure of size $O(n)$ bits that on query (i, j) determines in $O(1)$ time if S has a substring of length i with exactly j 1s.*

Proof We will show how to compute the array A such that $A[i]$ holds the maximum number of 1s in a substring of S of size i . The minimum is found similarly. We use a recent result of Gawrychowski [22] who showed how for any ℓ , we can modify $G(S)$ in $O(n)$ time by adding $O(g)$ new variables such that every new variable generates a string of length at most ℓ , and S can be written as the concatenation of substrings generated by these $O(g)$ new variables. Thus, we can write S as the concatenation of blocks $S = B_1 \circ \dots \circ B_b$ with $b = O(n/\ell)$ and $|B_j| \leq \ell$, such that amongst these blocks there are only $d = O(g)$ distinct blocks B_1^*, \dots, B_d^* . We refer to these d blocks as *basic blocks*. For each basic block B_k^* , $1 \leq k \leq d$, we first build an array A_k^* where $A_k^*[i]$ stores the maximum number of 1s over all substrings of B_k^* of length i . This is done in $O(\ell^2/\log^2 n)$ time per block (by using the algorithm of [28] for strings) for a total of $O(g \cdot \ell^2/\log^2 n)$.

We next handle substrings that span over two adjacent blocks. Namely, for each possible pair of basic blocks B_k^* and B_m^* , $1 \leq k \leq m \leq d$, we build an array $A_{k,m}^*$ where $A_{k,m}^*[i]$ stores the maximum number of 1s over all substrings of $B_k^* \circ B_m^*$ of length i that start in B_k^* and end in B_m^* . This is done in $O(\ell^2/\log^2 n)$ time for each pair for a total of $O(g^2\ell^2/\log^2 n)$. Recall that, since we use the algorithm of [28], the array $A_{k,m}^*$ is implicitly represented by an array of $O(\ell/\log n)$ words: For each i that is a multiple of $\log n$, the array keeps one word storing the maximum number of 1s over all substrings of length i , and another word storing the binary increment vector for substrings of length $i + 1, \dots, i + \log n - 1$.




Fig. 2 A string S partitioned into blocks $B_1 \circ \dots \circ B_b$, each of length at most ℓ . The shaded substring is one of the substrings considered for $A_{k,m}[i]$ as it includes the substring $S_{k,m}$ of length $i_{k,m}$ and a prefix and suffix of total length i

Finally, we consider substrings that span over more than two blocks. For each pair of (non-basic) blocks B_k and B_m , $1 \leq k < m \leq b$, let $S_{k,m} = B_{k+1} \circ \dots \circ B_{m-1}$ be a substring of S that is of length $i_{k,m} = |S_{k,m}|$ and has $j_{k,m}$ 1’s. Note that we can easily compute $i_{k,m}$ and $j_{k,m}$ of all $1 \leq k < m \leq b$ in total time $O(n^2/\ell^2)$. For every $1 \leq k < m \leq b$, we build an array $A_{k,m}$ of size $O(\ell)$ where $A_{k,m}[i]$ stores the maximum number of 1s over all substrings of $B_k \circ \dots \circ B_m$ of length $i + i_{k,m}$ that start in B_k and end in B_m . Notice that all such substring include $S_{k,m}$ as well as a suffix of B_k and a prefix of B_m whose total length is i (see Fig. 2). Therefore, for each $A_{k,m}$ we set $A_{k,m}[i]$ to be $j_{k,m}$ plus the maximal number of 1’s in a suffix of B_k and a prefix of B_m whose total length is i . In other words, we set $A_{k,m}[i] = j_{k,m} + A_{k',m'}^*[i]$ where k' (resp. m') is such that the block B_k (resp. B_m) corresponds to the basic block $B_{k'}^*$ (resp. $B_{m'}^*$). The computation of (an implicit representation of) each $A_{k,m}$ can be done in $O(\ell/\log n)$ time by only setting $A_{k,m}[i]$ for i ’s that are multiples of $\log n$ (the binary increment vectors of $A_{k,m}$ remain as in $A_{k',m'}^*$). Since there are $O((n/\ell)^2)$ pairs of blocks and each pair requires $O(\ell/\log n)$ time, we get a total of $O(n^2/(\ell \log n))$ time.

Finally, once we have the implicit representation of all $A_{k,m}$ ’s we can compute the desired array A from them in $O(n^2/(\ell \log n))$ time: For each i that is a multiple of $\log n$ and each $A_{k,m}$ we set $A[i]$ to be the maximum of $A[i]$ and $A_{k,m}[i - i_{k,m}]$ in $O(1)$ time. The next $\log n$ entries of A are computed in $O(1)$ time (as done in [28], Section 3.3) from the increment vectors of $A[i]$ and $A_{k,m}[i - i_{k,m}]$. To conclude, we get a total running time of $O(g^2\ell^2/\log^2 n + n^2/(\ell \log n)) = O(g^{2/3}n^{4/3}/(\log n)^{4/3})$ when ℓ is chosen to be $(n/g)^{2/3}(\log n)^{1/3}$. □

We also note that similarly to the case of trees (Subsect. 2.4), if we are willing to increase our index space to $O(n \log n)$ bits, then it is not difficult to turn indexes for detecting jumbled pattern matches in grammars into indexes for locating them. To obtain this, we build an index for S and recurse (build indexes) on $S_1 = B_1 \circ \dots \circ B_k$ and $S_2 = B_{k+1} \circ \dots \circ B_d$ where $|S_1|$ and $|S_2|$ are roughly $n/2$. This way, like in the centroid decomposition for trees, we can get in $O(\log n)$ time an anchor index of S . That is, an index of S that is part of a pattern appearance. Furthermore, as opposed to trees, we can then find the actual appearance (not just the anchor) in additional $O(i)$ time by sliding a window of size i that includes the anchor.

4 Jumbled Pattern Matching on Bounded Treewidth Graphs

In this section we consider the extension of binary jumbled pattern matching to the domain of graphs: Given a graph G whose vertices are colored either black

or white, and a query (i, j) , determine whether G has a connected subgraph G' with i white vertices and j black vertices³. This problem is also known as the (binary) graph motif problem in the literature. Fellows *et al.* [21] provided an $n^{O(w)}$ algorithm for this problem, where w is the treewidth of the input graph. Here we will substantially improve on this result by proving the following theorem, asserting that the problem is fixed-parameter tractable in the treewidth of the graph.

Theorem 3 *Binary jumbled pattern matching can be solved in $f(w) \cdot n^{O(1)}$ time on graphs of treewidth w . The function $f(w)$ can be bounded by $w^{O(w)}$ in case a tree decomposition of width w (see below) is provided with the input graph, and otherwise $f(w) = 2^{O(w^3)}$.*

Note that the algorithm in the theorem actually computes all queries (i, j) that appear in G , and can thus be easily converted to an index for the input graph.

4.1 Tree Decompositions

We begin by first introducing some necessary notation and terminology. Let $G = (V(G), E(G))$ be a graph. A *tree decomposition* of G is defined by a rooted tree \mathcal{T} whose nodes are subsets of $V(G)$, called *bags*, with the following two properties: (i) the union of all subgraphs induced by the bags of \mathcal{T} is G , and (ii) for any vertex $x \in V(G)$, the set of all bags including x induces a connected subgraph in \mathcal{T} . We use \mathcal{X} to denote the set of bags in a given tree decomposition. The *width* of the decomposition is defined as $\max_{X \in \mathcal{X}} |X| - 1$. The *treewidth* of G is the smallest possible width of any tree decomposition of G . Given a bag X of a given tree decomposition \mathcal{T} , we let G_X denote the subgraph induced by the union of all bags in \mathcal{T}_X . Bodlaender [8] gave an algorithm for computing a width- w tree decomposition of a given graph with treewidth w in $2^{O(w^3)}n$ time. We refer readers interested in further details to [20].

We will work with a specific kind of tree decompositions, namely *nice tree decompositions* [9]. A nice tree decomposition is a binary tree decomposition \mathcal{T} with four types of bags: *Leaf*, *forget*, *introduce*, and *join*. Leaf bags are the leaves of \mathcal{T} and are singleton sets which include a single vertex of G . A forget bag X has one child Y such that $X = Y \setminus \{x\}$ for some vertex x of G . Thus, X *forgets* the vertex x . Similarly, an introduce bag X has one child Y such that $X = Y \cup \{x\}$ for some vertex $x \notin Y$ of G . In this case, we say X *introduces* the vertex x . Finally a join bag X has two children Y and Z in \mathcal{T} with $X = Y = Z$. It is well known that given a tree decomposition of any graph, one can compute in polynomial-time a nice tree decomposition of the same graph with equal width and with an at most linear increase in its number of nodes [9]. Thus, from this point onwards we may assume that we are given a nice tree decomposition \mathcal{T} of G with width w .

³ The difference between the meaning of the query here and elsewhere in the paper is for ease of the presentation.




Fig. 3 A positive partition for the query (7, 6) in G_X , where $X = \{1, 2, 3, 4, 5, 6\}$. The partition is defined by $\{\{2\}, \{1, 4\}, \{5\}, \{3, 6\}\}$. Notice that vertex 2 is not in any of the connected graphs witnessing the partial occurrence. Also note that these graphs may or may not have edges between them

4.2 Positive Partitions

We next describe the main data structure that we compute in our algorithm. Let X be an arbitrary bag. A partition⁴ $\Pi_X = \{X_0, X_1, \dots, X_x\}$ of X is *positive* for a given query (i, j) in G_X if there are x disjoint connected subgraphs G_1, \dots, G_x of G_X such that (1) the total number of black (resp. white) vertices in $G' = G_1 \cup \dots \cup G_x$ is i (resp. j), and (2) $V(G') \cap X_0 = \emptyset$ and $V(G_\ell) \cap X = X_\ell$ for each $\ell = 1, \dots, x$ (see Fig. 3). Thus, positive partitions capture *partial occurrences* that intersect X at exactly $X \setminus X_0$. These may not be actual occurrences as we do not require any edges between the different G_i 's, and so G' itself may not be connected. We let $A_X[i, j]$ denote the set of all positive partitions for a query (i, j) , and let A_X denote the array with an entry for each possible query (i, j) . We will require that the trivial partition where $X_0 = X$ is only positive for the query $(0, 0)$.

Note that by definition, a query (i, j) appears in G_X iff there exists some partition into two sets $\{X_0, X_1\}$ that is positive for (i, j) in G_X . Since (i, j) appears in G iff (i, j) appears in G_X for some bag $X \in \mathcal{X}$, this means that it is also positive for (i, j) in some G_X . Thus, computing the arrays A_X for all bags $X \in \mathcal{X}$ suffices for solving our problem. We do this by computing all arrays A_X in a bottom-up fashion from the leaves to the root of \mathcal{T} . Note that the size of each array A_X can easily be bounded by $w^{O(w)}n^2$, considering that the w 'th Bell number is bounded by $w^{O(w)}$. Thus, to get a similar term in our running time, we will show that computing the array A_X from the arrays of the children of X can be done in polynomial-time. The computation on leaf bags is trivial (as they are singletons), and the computation on forget nodes is almost equally easy: If X is a forget bag with child Y , then computing

⁴ Here we slightly abuse our terminology and allow X_0 to be the empty set.

A_X from A_Y in this case amounts to converting each positive partition Π_Y of Y to a corresponding positive partition Π_X of X by removing x , the vertex forgotten by X , from the class it belongs to in Π_Y . We thus focus below on introduce nodes and join nodes.

4.3 Introduce Nodes

Let X be an introduce bag with child Y in \mathcal{T} , and let x be the vertex introduced by X . Let us assume for ease of presentation that x is colored white (the case where it is colored black is symmetric). By the properties of a tree decomposition, we know that x is only adjacent to vertices $y \in Y$ in G_X [20]. Let y_1, \dots, y_ℓ denote these neighbors of x , and let G_X^k denote the graph obtained by deleting the edges $\{x, y_{k+1}\}, \dots, \{x, y_\ell\}$ from G_X for each $k = 0, \dots, \ell$ ($G_X^\ell = G_X$). Similarly, let $A_X^k[i, j]$ denote the set of all positive partitions of (i, j) in G_X^k . We will compute $A_X^0[i, j]$ from A_Y , and $A_X^k[i, j]$ from $A_X^{k-1}[i, j]$ for each $k > 0$. Finally, we will set $A_X[i, j] = A_X^\ell[i, j]$.

We begin with $k = 0$. In this case, x is an isolated vertex in G_X^0 . Hence, there are only two types of positive partitions of (i, j) in G_X :

- A partition Π_X obtained by taking $\Pi_X = \Pi_Y \cup \{\{x\}\}$ for some $\Pi_Y \in A_Y[i - 1, j]$ (thus, $\{x\}$ is a singleton set in Π_X).
- A partition Π_X obtained by taking a partition $\Pi_Y \in A_Y[i, j]$ and adding x to $Y_0 \in \Pi_Y$ (thus, x is in the set of vertices not included in the partial occurrence captured by Π_X).

It is easy to see that since x is an isolated vertex the above description indeed captures all types of positive partitions for G_X^0 , and so we can compute $A_X^0[i, j]$ from A_Y in polynomial time.

Assume that $k > 0$. Then any positive partition for (i, j) in G_X^{k-1} is also positive in G_X^k . Moreover, the only new positive partitions for (i, j) in G_X^k that were not positive in G_X^{k-1} are partitions where x and y_k belong to the same class (although, there might be partitions of this type which were positive in G_X^{k-1}). Thus, we compute $A_X^k[i, j]$ by first setting $A_X^k[i, j] = A_X^{k-1}[i, j]$. Then for each $\Pi \in A_X^{k-1}[i, j]$ with $x \in X_i \in \Pi$ and $y_k \in X_j \in \Pi, i \neq j$, we add the partition $(\Pi \setminus \{X_i, X_j\}) \cup \{X_i \cup X_j\}$ to $A_X^k[i, j]$ (assuming it is not already there). The total amount of computation time required here is obviously polynomial in the size of $A_X^{k-1}[i, j]$.

4.4 Join Nodes

Consider a join bag X with two children Y and Z in \mathcal{T} , and recall that $X = Y = Z$. For a pair of partitions $\Pi_Y = \{Y_0, \dots, Y_y\}$ and $\Pi_Z = \{Z_0, \dots, Z_z\}$ of Y and Z , we define the partition $\Pi_Y \oplus \Pi_Z$ (the *join* of Π_Y and Π_Z) as follows: Consider the two equivalence relations induced by $\Pi_X \setminus X_0$ and $\Pi_Y \setminus Y_0$ on the set $\Pi_X \setminus X_0$. Then the equivalence relation defined by $\Pi_X \setminus X_0$ is taken as the transitive closure of the union of these two equivalence relations. As an example, suppose that $X = Y = Z =$

$\{1, \dots, 6\}$, $\Pi_Y = \{\{1\}, \{2\}, \{3, 4\}, \{5, 6\}\}$, and $\Pi_Z = \{\{2\}, \{1\}, \{3, 5\}, \{4, 6\}\}$. Then $\Pi_X = \{X_0, X_1\}$, where $X_0 = \{1, 2\}$ and $X_1 = \{3, 4, 5, 6\}$.

Let i_0 and j_0 respectively denote the number of white and black vertices in $X \setminus X_0$. We claim that if (i_1, j_1) and (i_2, j_2) are two queries for which Π_Y and Π_Z are respectively positive in G_Y and G_Z , then $\Pi_X = \Pi_Y \oplus \Pi_Z$ is positive for $(i_1 + i_2 - i_0, j_1 + j_2 - j_0)$. This can be verified by considering the connected components in the graph $G'_X = G_1^Y \cup \dots \cup G_y^Y \cup G_1^Z \cup \dots \cup G_z^Z$, where G_1^Y, \dots, G_y^Y and G_1^Z, \dots, G_z^Z are sets of graphs witnessing that Π_Y and Π_Z are positive for (i_1, j_1) in G_Y and (i_2, j_2) in G_Z . It is easy to see that the total number of white and black vertices in these components is $i = i_1 + i_2 - i_0$ and $j = j_1 + j_2 - j_0$, where i_0 white vertices and j_0 black vertices are subtracted due to double counting the vertex colors in X . Moreover, it can be verified that these components intersect X as required by Π_X , due to the fact that Π_X is the transitive closure of $\Pi_Y \cup \Pi_Z$. Thus, G'_X is a partial occurrence of (i, j) in G_X , and $\Pi_X \in A_X[i, j]$.

On the other hand, it can also be seen on the same lines that if (i, j) is a query for which Π_X is positive in G_X , then it is either in $\Pi_Y[i, j]$ or $\Pi_Z[i, j]$, or we have $(i, j) = (i_1 + i_2 - i_0, j_1 + j_2 - j_0)$ for some pair of queries (i_1, j_1) and (i_2, j_2) for which Π_Y and Π_Z are positive in G_Y and G_Z . We can therefore compute $A_X[i, j]$ by first setting $A_X[i, j]$, and then examining all pairs (i_1, j_1) and (i_2, j_2) as above. For each such pair, we compute all partitions $\Pi_Y \oplus \Pi_Z$ for $\Pi_Y \in A_Y[i_1, j_1]$ and $\Pi_Z \in A_Z[i_2, j_2]$. Note that the entire computation of A_X requires time which is polynomial in the total sizes of A_Y and A_Z .

4.5 Summary

We have shown above how to compute, for any given query (i, j) , the array A_X for each bag X of \mathcal{T} in $w^{O(w)}n^{O(1)}$ time. As the total number of bags is $O(n)$, we obtain an algorithm whose total running time is $w^{O(w)}n^{O(1)}$, excluding the time required to compute the nice tree decomposition \mathcal{T} . This completes the proof of Theorem 3. We note that our algorithm straightforwardly extends to an $w^{O(w)}n^{O(c)}$ time algorithm for the case where the vertices of G are colored with c colors.

5 Conclusions and Open Problems

In this paper we considered the binary jumbled pattern matching problem on trees, bounded treewidth graphs, and strings compressed by context-free grammars in Chomsky normal form. We gave an $\tilde{O}(g^{2/3}n^{4/3})$ -time solution for strings of length n represented by grammars of size g , an $f(w) \cdot n^{O(1)}$ -time solution for graphs with treewidth w , and an $O(n^2/\log^2 n)$ -time solution for trees. In the latter result, we showed how to determine in $O(1)$ time if a query pattern appears, and how to locate in $O(\log n)$ time a node of this appearance. With a linear-space solution, locating the entire appearance remains an open problem. Using Lemma 3, the construction time for trees can be made $O(n \cdot i/\log^2 n)$ if the query patterns are known to be of size at most i . We also note here that the construction time

can be made faster on trees that have many identical rooted subtrees⁵. This is because the bottom-up construction does not need to be applied on the same subtree twice.

Finally, the main open problems stemming from our work are: (1) To obtain a faster construction of the linear-space index for strings. Our index for trees implies that any construction speedup for strings implies a construction speedup for trees. (2) To develop an algorithm for the non-indexing variant of binary jumbled pattern matching on trees whose performance is closer to the performance of the corresponding algorithm on strings (*i.e.* the $O(n)$ sliding window algorithm).

Acknowledgments We thank the anonymous reviewers for their helpful comments.

References

1. Alstrup, S., Secher, J., Spørnk, M.: Optimal on-line decremental connectivity in trees. *Inf. Process. Lett.* **64**(4), 161–164 (1997)
2. Ambalath, A.M., Balasundaram, R., Rao, C.H., Koppula, V., Misra, N., Philip, G., Ramanujan, M.S.: On the kernelization complexity of colorful motifs. In: *Proceedings of the 5th International Symposium Parameterized and Exact Computation*, pp. 14–25, (2010)
3. Amir, A., Chan, T.M., Lewenstein, M., Lewenstein, N.: On hardness of jumbled indexing. In: *Proceedings of the 41st International Colloquium on Automata, Languages and Programming (ICALP)*, pp. 114–125, (2014)
4. Badkobeh, G., Fici, G., Kroon, S., Lipták, Z.: Binary jumbled string matching for highly run-length compressible texts. *Inf. Process. Lett.* **113**(17), 604–608 (2013)
5. Benson, G.: Composition alignment. In: *Proceedings of the 3rd International Workshop on Algorithms in Bioinformatics (WABI)*, pp. 447–461, (2003)
6. Betzler, N., van Bevern, R., Fellows, M.R., Komusiewicz, C., Niedermeier, R.: Parameterized algorithmics for finding connected motifs in biological networks. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **8**(5), 1296–1308 (2011)
7. Böcker, S.: Simulating multiplexed SNP discovery rates using base-specific cleavage and mass spectrometry. *Bioinformatics* **23**(2), 5–12 (2007)
8. Bodlaender, H.L.: A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM J. Comput.* **25**, 1305–1317 (1996)
9. Bodlaender, H.L.: Treewidth. Algorithmic techniques and results. In: *Proceedings of the 22nd International Symposium on Mathematical Foundations of Computer Science (MFCS)*, pp. 19–36, (1997)
10. Bursi, P., Cicalese, F., Fici, G., Lipták, Z.: On table arrangement, scrabble freaks, and jumbled pattern matching. In: *Proceedings of the Symposium on Fun with Algorithms*, pp. 89–101, (2010)
11. Bursi, P., Cicalese, F., Fici, G., Lipták, Z.: Algorithms for jumbled pattern matching in strings. *Int. J. Found. Comput. Sci.* **23**(2), 357–374 (2012)
12. Bursi, P., Cicalese, F., Fici, G., Lipták, Z.: On approximate jumbled pattern matching in strings. *Theory Comput. Syst.* **50**(1), 35–51 (2012)
13. Charikar, M., Lehman, E., Liu, D., Panigrahy, R., Prabhakaran, M., Sahai, A., Shelat, A.: The smallest grammar problem. *IEEE Trans. Inf. Theory* **51**(7), 2554–2576 (2005)
14. Cicalese, F., Fici, G., Lipták, Z.: Searching for jumbled patterns in strings. In: *Proceedings of the Prague Stringology Conference*, pp. 105–117, (2009)
15. Cicalese, F., Gagie, T., Giaquinta, E., Laber, E., Liptak, S., Rizzi, R., Tomescu, A.I.: Indexes for jumbled pattern matching in strings, trees and graphs. In: *Proceedings of the 20th International Symposium on String Processing and Information Retrieval (SPIRE)*, pp. 56–63, (2013)

⁵ There is a simple linear-time algorithm that given a tree finds the best way to share identical rooted subtrees [19].

16. Cicalese, F., Laber, E.S., Weimann, O., Yuster, R.: Near linear time construction of an approximate index for all maximum consecutive sub-sums of a sequence. In: Proceedings of the Symposium on Combinatorial Pattern Matching, pp. 149–158, (2012)
17. Dondi, R., Fertin, G., Vialette, S.: Complexity issues in vertex-colored graph pattern matching. *J. Discrete Algorithms* **9**(1), 82–99 (2011)
18. Dondi, R., Fertin, G., Vialette, S.: Finding approximate and constrained motifs in graphs. In: Proceedings of the 22nd Annual Symposium Combinatorial Pattern Matching, pp. 388–401, (2011)
19. Downey, P.J., Sethi, R., Tarjan, R.E.: Variations on the common subexpression problem. *J. ACM* **27**, 758–771 (1980)
20. Downey, R.G., Fellows, M.R.: Parameterized Complexity. Springer, Berlin (1999)
21. Fellows, M.R., Fertin, G., Hermelin, D., Vialette, S.: Upper and lower bounds for finding connected motifs in vertex-colored graphs. *J. Comput. Syst. Sci.* **77**(4), 799–811 (2011)
22. Gawrychowski, P.: Faster algorithm for computing the edit distance between SLP-compressed strings. In: Proceedings of the Symposium on String Processing and Information Retrieval, pp. 229–236, (2012)
23. Giaquinta, E., Grabowski, S.: New algorithms for binary jumbled pattern matching. *Inf. Process. Lett.* **113**(14–16), 538–542 (2013)
24. Jacobson, G.: Space-efficient static trees and graphs. In: Proceedings of the 30th Annual Symposium on Foundations of Computer Science (FOCS), pp. 549–554, (1989)
25. Kociumaka, T., Radoszewski, J., Rytter, W.: Efficient indexes for jumbled pattern matching with constant-sized alphabet. In: ESA, pp. 625–636, (2013)
26. Lacroix, V., Fernandes, C.G., Sagot, M.-F.: Motif search in graphs: application to metabolic networks. *IEEE/ACM Trans. Comput. Biol. Bioinform.* **3**(4), 360–368 (2006)
27. Moosa, T.M., Rahman, M.S.: Indexing permutations for binary strings. *Inf. Process. Lett.* **110**(18–19), 795–798 (2010)
28. Moosa, T.M., Rahman, M.S.: Sub-quadratic time and linear space data structures for permutation matching in binary strings. *J. Discrete Algorithms* **10**, 5–9 (2012)
29. Munro, I.: Tables. In: Proceedings of the 16th Foundations of Software Technology and Theoretical Computer Science (FSTTCS), pp. 37–42, (1996)
30. Rytter, W.: Application of Lempel–Ziv factorization to the approximation of grammar-based compression. *Theoret. Comput. Sci.* **302**(1–3), 211–222 (2003)