

Binary Jumbled Pattern Matching on Trees and Tree-Like Structures

Travis Gagie¹, Danny Hermelin², Gad M. Landau^{3,4,*}, and Oren Weimann³

¹ University of Helsinki

travis.gagie@cs.helsinki.fi

² Ben-Gurion University

hermelin@bgu.ac.il

³ University of Haifa

{landau,oren}@cs.haifa.ac.il

⁴ NYU poly

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 grammar of size g . This solution improves the known bounds when the string is compressible under many popular compression schemes. Finally, we prove that the problem is fixed-parameter tractable with respect to the treewidth w of the graph, even for a constant number of different vertex-labels, thus improving the previous best $n^{O(w)}$ algorithm.

1 Introduction

Jumbled pattern matching is an important variant of classical pattern matching with several applications in computational biology, ranging from alignment [4] and SNP discovery [6], to the interpretation of mass spectrometry data [9] and metabolic network analysis [21]. 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

* Partially supported by the National Science Foundation Award 0904246, Israel Science Foundation grant 347/09, and Grant No. 2008217 from the United States-Israel Binational Science Foundation (BSF).

is of length i and has j 1s). Given a binary string of length n , Cicalese, Fici and Lipták [13] 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.* [9] (see also [10,11]) and Moosa and Rahman [22] independently improved the construction time to $O(n^2/\log n)$, then Moosa and Rahman [23] further improved it to $O(n^2/\log^2 n)$ in the RAM model. Currently, faster algorithms than $O(n^2/\log^2 n)$ exist only when the string compresses well using run-length encoding [3,19] or when we are willing to settle for approximate indexes [14].

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 [21], even if our query contains at most one occurrence of each letter [17]. It is not even fixed-parameter tractable when parameterized by the alphabet size [17]. The fixed-parameter tractability of the problem was further studied when extending the problem from trees to graphs [2,5,15,16]. In particular, the problem (also known as the *graph motif problem*) was recently shown by Fellows *et al.* [17] to be polynomial-time solvable when the number of letters in the alphabet as well as the treewidth of the graph are both fixed.

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. The problem on such trees turns out to be more challenging than on strings and requires substantially different ideas and techniques.

- **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 [23] 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 grammar of size g , 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 [12,24].

- **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.* [17] 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 sections 2, 3 and 4 respectively. Proofs of all lemmas are given in the full version of this paper.

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 [22,23]. 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 on 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 .*

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 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 children of v_1 , and make $v_{\ell-1}$ and $u_{\ell+1}$ be the children of v_ℓ for each $\ell = 2, \dots, k - 1$. We call the nodes v_2, \dots, v_k *dummy nodes*. This procedure at most doubles the size of T . To avoid cumbersome notation, we henceforth use T and n to denote the resulting binary rooted tree and its number of nodes. 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 $|T_v| + 1$. The entry $A_v[i]$ will store the maximum number of black nodes that appear in a connected subgraph of size i that includes v and another $i - 1$ nodes in T_v . 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 in 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 . 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 ℓ index starts with 1 to indicate that the non-dummy copy of v is already included in the pattern.

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

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 $A[i]$ 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 bits, where $B_v[0] = 0$, and $B_v[i] = A_v[i] - A_v[i - 1]$ for all $i = 1, \dots, n - 1$. 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 [20].

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 is 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) .*

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.

From trees to strings. Recall that we can represent every A_v by a binary string B_v where $B_v[i] = A_v[i] - A_v[i - 1]$. 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 $S_v = X_v \circ col(v) \circ Y_v$ (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 $col(v)$ the *split position* of S_v . Recall that $A_v[i] = col(v) + \max_{0 \leq \ell \leq i-1} \{A_u[\ell] + A_w[i - 1 - \ell]\}$. This is equal to the maximum number of 1s in a window of S 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 [23] for the standard string problem we can also solve our problem and compute A_v in $O(|S|^2/\log^2 n)$ 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 $O((|X_v| + |Y_v|)^2/\log^2 n)$.

First Speedup. The $O(\log^2 n)$ -factor speedup for jumbled pattern matching on strings [23] is achieved by a clever combination of lookup tables. One log factor is achieved by computing the maximum 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 windows of length 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 [23] 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.

Second Speedup. In strings, an additional logarithmic improvement shown in [23] 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 $\leq s$ (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.

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 $parent(C)$ (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 .

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. The array A_C stores the maximum 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} . The array A_b stores

¹ The root of the macro tree is unique as it might have a top boundary node connected to two 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 black nodes correspond to boundary nodes and the white nodes to non-boundary nodes.

the maximum 1s in every pattern that includes the boundary node $b(C)$ and other nodes of C , $T_{\ell(C)}$, and $T_{r(C)}$. The array A_t stores the maximum 1s in every pattern that includes the boundary node $t(C)$ and other nodes of C and $T_{child(C)}$. Finally, the array A_{tb} stores the maximum 1s in every pattern that includes *both* boundary nodes $t(C)$ and $b(C)$ and other nodes of C , $T_{\ell(C)}$, 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 1s in every pattern that includes $t(C)$ and nodes of C . This can be done in $O(|C|^2)$ time for each C by rooting C at $t(C)$ and running the algorithm from the previous subsection. Similarly, we initialize the array A_b to hold the maximum 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 1s in all patterns that are completely inside a micro tree. 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 $\ell(C)$, $r(C)$ and $child(C)$. We first compute the maximum 1s in all patterns that include $b(C)$ and vertices of $T_{\ell(C)}$ and $T_{r(C)}$. This can be done using the aforementioned string speedups in $O(|T_{\ell(C)}| \cdot |T_{r(C)}| / \log^2 n)$ time when both $|T_{\ell(C)}| > \log n$ 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 in time $O((|T_{\ell(C)}| + |T_{r(C)}|)/\log n) = O(n/\log n)$. 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 $O((|T_{\ell(C)}| + |T_{r(C)}| + |C|) \cdot |T_{child(C)}|/\log^2 n)$ time if both $|T_{\ell(C)}| + |T_{r(C)}| + |C| > \log n$ 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 $b(C)$.

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* 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, the time to construct all the arrays is bounded by $T(n) = 2T(n/2) + O(n^2/\log^2 n) = O(n^2/\log^2 n)$.

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 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 . Finally, we note that the above can be extended so that for every occurrences of (i, j) one node that is part of this occurrence is reported.

3 Jumbled Pattern Matching on Grammars

In grammar 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 $n = |S|$, and is always at most $O(n/\log n)$. 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. The proof is given in the appendix.

Theorem 2. *Given a binary string S of length n compressed by a grammar $G(S)$ of size g , 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.*

We also note that similarly to the case of trees (subsection 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 and white, and a query (i, j) , determine whether G has a connected subgraph G' with i white vertices and j black vertices². This problem is also known as the (binary) graph motif problem in the literature. Fellows *et al.* [17] 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 .*

² The difference between the meaning of the query here and elsewhere in the paper is for ease of the presentation.

The function $f(w)$ in the theorem above can be replaced with $w^{O(w)}$ in case a tree decomposition of width w (see below) is provided with the input graph, and otherwise it can be replaced by $2^{O(w^3)}$. Also, 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.

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 a 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 v , the set of all bags including v 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 [7] 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 next describe the information we store for each bag in the tree decomposition of G . Let X be an arbitrary bag. A partition $\Pi_X = \{X_0, X_1, \dots, X_x\}$ 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 and white vertices in $G' = G_1 \cup \dots \cup G_x$ is i and j respectively, and (2) $V(G') \cap X_0 = \emptyset$ and $V(G_\ell) \cap X = X_\ell$ for each $\ell = 1, \dots, x$. Here we slightly abuse our terminology and allow X_0 to be the empty set. The information we compute for each bag X is an array A_X which has an entry for each possible query (i, j) , where the entry $A_X[i, j]$ contains the set of all positive partitions of X for (i, j) in G_X . Note that the 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) . Since (i, j) appears in G iff (i, j) appears in G_X for some bag $X \in \mathcal{X}$, computing the arrays A_X for all bags allows us to determine whether (i, j) appears in G . Our algorithm computes all arrays A_X in a bottom-top fashion from the leaves to the root of \mathcal{T} . It is easy to verify that the size of each array A_X is bounded by $w^{O(w)}n^2$. 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 with respect to the child array sizes.

We will work with a specific kind of tree decompositions, namely *nice tree decompositions* [8]. A nice tree decomposition is a binary rooted tree decomposition \mathcal{T} with four types of bags: *Leaf*, *forget*, *introduce*, and *join*. Leaf bags are the leaves of \mathcal{T} and include a single vertex of G , and so computing A_X for leaf bags is trivial. A forget bag X has a single child Y with $X = Y \setminus \{v\}$ for some vertex v of G . Computing 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 v from the class it belongs to in Π_Y . An introduce bag X also has a single child Y , but this time we have $X = Y \cup \{v\}$ for some vertex $v \notin Y$ of G . By the properties of a tree decomposition, we know that v is only adjacent to vertices of Y in G_X . Computing A_X from A_Y in this case requires the consideration of all partitions of X which are formed from positive partitions Π_Y of Y by adding

v to a class in Π_Y with one of its neighbors (or adding $\{v\}$ as a new singleton class). We leave the precise details to the full version of this paper, but it should be easy to see that computing A_X in this case, as well as in all cases above, can be done in $w^{O(w)}n^{O(1)}$.

The more challenging case is when X is a join bag. A join bag X has two children Y and Z in \mathcal{T} , with $X = Y = Z$. Consider two partitions $\Pi_Y = \{Y_0, \dots, Y_y\}$ and $\Pi_Z = \{Z_0, \dots, Z_z\}$ for Y and Z . We define the partition $\Pi_Y \oplus \Pi_Z$ as follows: First we set X_0 to be $Y_0 \cap Z_0$. The remaining classes are constructed such that any pair of vertices in X belong to the same class in $\Pi_X \setminus \{X_0\}$ iff they belong to the same class in $\Pi_Y \setminus \{Y_0\}$ or to the same class in $\Pi_Z \setminus \{Z_0\}$.

Let i_0 and j_0 respectively denote the number of white and black vertices in X . 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 $G_1^Y \cup \dots \cup G_y^Y \cup G_1^Z \cup \dots \cup G_z^Z$, 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 $i_1 + i_2 - i_0$ and $j_1 + j_2 - j_0$, 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 .

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 $(i, j) = (i_1 + i_2 - i_0, j_1 + j_2 - j_0)$ 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 examining all such pairs (i_1, j_1) and (i_2, j_2) , and computing the partition $\Pi_Y \oplus \Pi_Z$ for each pair of positive partitions $\Pi_Y \in A_Y[i_1, j_1]$ and $\Pi_Z \in A_Z$. This requires $w^{O(w)}n^{O(1)}$ time.

To summarize we compute each array A_X in $w^{O(w)}n^{O(1)}$ time. As the total number of bags is $O(wn)$, 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} . We note that the running time of our algorithm can be improved slightly by using an extension of Lemma 1 to graphs. Also, our result 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.

References

1. Alstrup, S., Secher, J., Sporkn, M.: Optimal on-line decremental connectivity in trees. *Information Processing Letters* 64(4), 161–164 (1997)
2. Ambalath, A.M., Balasundaram, R., Rao H., C., Koppula, V., Misra, N., Philip, G., Ramanujan, M.S.: On the kernelization complexity of colorful motifs. In: Raman, V., Saurabh, S. (eds.) *IPEC 2010*. LNCS, vol. 6478, pp. 14–25. Springer, Heidelberg (2010)
3. 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)
4. Benson, G.: Composition alignment. In: Benson, G., Page, R.D.M. (eds.) *WABI 2003*. LNCS (LNBI), vol. 2812, pp. 447–461. Springer, Heidelberg (2003)
5. 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. Biology Bioinform.* 8(5), 1296–1308 (2011)

6. Böcker, S.: Simulating multiplexed SNP discovery rates using base-specific cleavage and mass spectrometry. *Bioinformatics* 23(2), 5–12 (2007)
7. Bodlaender, H.L.: A linear time algorithm for finding tree-decompositions of small treewidth. *SIAM Journal on Computing* 25, 1305–1317 (1996)
8. Bodlaender, H.L.: Treewidth. Algorithmic techniques and results. In: Privara, I., Ružička, P. (eds.) *MFCS 1997*. LNCS, vol. 1295, pp. 19–36. Springer, Heidelberg (1997)
9. Burcsi, P., Cicalese, F., Fici, G., Lipták, Z.: On table arrangement, scrabble freaks, and jumbled pattern matching. In: Boldi, P. (ed.) *FUN 2010*. LNCS, vol. 6099, pp. 89–101. Springer, Heidelberg (2010)
10. Burcsi, P., Cicalese, F., Fici, G., Lipták, Z.: Algorithms for jumbled pattern matching in strings. *International Journal of Foundations of Computer Science* 23(2), 357–374 (2012)
11. Burcsi, P., Cicalese, F., Fici, G., Lipták, Z.: On approximate jumbled pattern matching in strings. *Theory of Computing Systems* 50(1), 35–51 (2012)
12. Charikar, M., Lehman, E., Liu, D., Panigrahy, R., Prabhakaran, M., Sahai, A., Shelat, A.: The smallest grammar problem. *IEEE Transactions on Information Theory* 51(7), 2554–2576 (2005)
13. Cicalese, F., Fici, G., Lipták, Z.: Searching for jumbled patterns in strings. In: *Proc. of the Prague Stringology Conference*, pp. 105–117 (2009)
14. Cicalese, F., Laber, E., Weimann, O., Yuster, R.: Near linear time construction of an approximate index for all maximum consecutive sub-sums of a sequence. In: Kärkkäinen, J., Stoye, J. (eds.) *CPM 2012*. LNCS, vol. 7354, pp. 149–158. Springer, Heidelberg (2012)
15. Dondi, R., Fertin, G., Vialette, S.: Complexity issues in vertex-colored graph pattern matching. *J. Discrete Algorithms* 9(1), 82–99 (2011)
16. Dondi, R., Fertin, G., Vialette, S.: Finding approximate and constrained motifs in graphs. In: Giancarlo, R., Manzini, G. (eds.) *CPM 2011*. LNCS, vol. 6661, pp. 388–401. Springer, Heidelberg (2011)
17. 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)
18. Gawrychowski, P.: Faster algorithm for computing the edit distance between SLP-compressed strings. In: Calderón-Benavides, L., González-Caro, C., Chávez, E., Ziviani, N. (eds.) *SPIRE 2012*. LNCS, vol. 7608, pp. 229–236. Springer, Heidelberg (2012)
19. Giaquinta, E., Grabowski, S.: New algorithms for binary jumbled pattern matching. *Inf. Process. Lett.* 113(14–16), 538–542 (2013)
20. Jacobson, G.: Space-efficient static trees and graphs. In: *Proc. of the 30th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 549–554 (1989)
21. Lacroix, V., Fernandes, C.G., Sagot, M.-F.: Motif search in graphs: Application to metabolic networks. *IEEE/ACM Trans. Comput. Biology Bioinform.* 3(4), 360–368 (2006)
22. Moosa, T.M., Rahman, M.S.: Indexing permutations for binary strings. *Information Processing Letters* 110(18–19), 795–798 (2010)
23. Moosa, T.M., Rahman, M.S.: Sub-quadratic time and linear space data structures for permutation matching in binary strings. *Journal of Discrete Algorithms* 10, 5–9 (2012)
24. Rytter, W.: Application of Lempel-Ziv factorization to the approximation of grammar-based compression. *Theoretical Computer Science* 302(1–3), 211–222 (2003)