

Indexes for Jumbled Pattern Matching in Strings, Trees and Graphs

Ferdinando Cicalese¹, Travis Gagie^{2,3},
Emanuele Giaquinta^{2,*}, Eduardo Sany Laber⁴,
Zsuzsanna Lipták⁵, Romeo Rizzi⁵,
and Alexandru I. Tomescu^{2,3,**}

¹ Department of Computer Science, University of Salerno, Italy

² Department of Computer Science, University of Helsinki, Finland

³ Helsinki Institute for Information Technology, Finland

⁴ Department of Computer Science, PUC Rio de Janeiro, Brazil

⁵ Department of Computer Science, University of Verona, Italy

Abstract. We consider how to index strings, trees and graphs for jumbled pattern matching when we are asked to return a match if one exists. For example, we show how, given a tree containing two colours, we can build a quadratic-space index with which we can find a match in time proportional to the size of the match. We also show how we need only linear space if we are content with approximate matches.

1 Introduction

Suppose we are given a connected graph G on n coloured nodes and a multiset M of colours and asked to find a connected subgraph of G whose nodes' colours are exactly those in M , if such a subgraph exists. This problem is commonly referred to as *jumbled pattern matching*, and has recently aroused much interest in the case of strings [7,5,8,6,17,13,14,2]: There we are looking for substrings of the text which have the same multiplicity of each character as the query, also referred to as its *Parikh vector*. (The boolean version is the Parikh fingerprint or character set [1,10].) Parikh vectors appear frequently in applications in computational biology [3,5,11,4], as do jumbled patterns in graphs [16].

Even when G is a tree, there can be exponentially many such matching subgraphs. When G is a path, however, there are $\mathcal{O}(n)$ matches and we can find them all in $\mathcal{O}(n)$ time [7]. When G is a path containing a constant number of colours—i.e. the nodes are coloured with only two colours—then in $\mathcal{O}(n^2)$ time we can build a $o(n^2)$ -space index with which we can determine in $o(n)$ time whether there is a match [15]. When G is a path containing only two colours, in $\mathcal{O}(n^2/\log^2 n)$ time we can build an $\mathcal{O}(n)$ -bit index with which we can determine in $\mathcal{O}(1)$ time whether there is a match [8,6,17,13]. Moreover, in $\mathcal{O}(n^2/\log^2 n)$ time we can build an index of size $\mathcal{O}(n \log n)$ -bits with which we can find all the

The original version of this chapter was revised: The copyright line was incorrect. This has been corrected. The Erratum to this chapter is available at DOI: [10.1007/978-3-319-02432-5_33](https://doi.org/10.1007/978-3-319-02432-5_33)

* Supported by Academy of Finland grant 118653 (ALGODAN).

** Supported by Academy of Finland grant 250345 (CoECCR).

matches using $\mathcal{O}(|M|)$ worst-case time per match [13]. We can build an approximation of this index in $\mathcal{O}(n^{1+\epsilon})$ time with the quality of the approximation depending on ϵ [9]. Throughout this paper our model is the word-RAM with $\Omega(\log n)$ -bit words and we measure space in words unless stated otherwise.

Determining whether there is a match is NP-complete even when G is a tree [16]. It is also NP-complete when G contains only two colours, but takes polynomial time when G both has bounded treewidth and contains only a constant number of colours [12]. When G contains only two colours there exists an $\mathcal{O}(n)$ -bit index with which we can determine in $\mathcal{O}(1)$ time whether there is a match [13]. Building this index is NP-hard in general but, since finding a match is self-reducible, takes polynomial time when G has bounded treewidth and $\mathcal{O}(n^2/\log^2 n)$ time when G is a tree. At the cost of increasing the space to $\mathcal{O}(n)$ words, this index can be generalized to return a subset of the nodes in the matches that is also a hitting set for all the matches, using $\mathcal{O}(\log n)$ time worst-case time per match. In the worst case, however, this subset of nodes is of little use in finding even a single complete match.

We start by presenting some basic tradeoffs in Section 2, to establish what can be done naïvely on graphs. In Sections 3 to 6 we assume G contains only two colours. In Section 3 we consider the case when G is a path — i.e., a binary string — and describe an $\mathcal{O}(n)$ -space index with which we can find a match in $\mathcal{O}(\log n)$ time. In Section 4 we consider the case when G is a tree and, based on our index for binary strings, describe an $\mathcal{O}(n^2)$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time. If we are concerned only with multisets of size at most $n^{1/2}$, then we can reduce the space bound to $\mathcal{O}(n)$. In Section 5 we show that we can achieve the same space bound if we are content with approximate matches. In Section 6 we partially extend our results from trees to graphs.

2 Basic Tradeoffs

Suppose G is a graph containing a constant number c of colours and we will be given M as the vector of length c whose components are the frequencies of the colours, the Parikh vector for M ; note that the sum of its entries equals $|M|$. Since for any $1 \leq m \leq n$, there are $\binom{m+c-1}{c-1} = \mathcal{O}(m^{c-1})$ possible multisets of size m and it takes $\mathcal{O}(m)$ space to store pointers to a match for such a multiset, there exists an $\mathcal{O}(n^{c+1})$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time. When G has bounded treewidth we can build this index in polynomial time, and we can reduce the space bound to $\mathcal{O}(n)$ at the cost of increasing the query time to $|M|^{\mathcal{O}(1)}$. To do the latter, we store G itself and pre-compute and store pointers to matches only for multisets of size at most $n^{1/(c+1)}$. Given a multiset M with $|M| > n^{1/(c+1)}$, we search G in $n^{\mathcal{O}(1)} = |M|^{\mathcal{O}(1)}$ time.

For any positive constant ϵ , we can build an $\mathcal{O}(n \log^c n)$ -space approximate index with which, if M has an exact match, then in $\mathcal{O}(1)$ time we can find a substring whose Parikh vector differs from M 's by at most a factor of $1 + \epsilon$ in each component. (This index does not tell us whether M has an exact match, however, since we may find such a substring even when it does not.) Without

loss of generality, assume we are concerned only with multisets in which each character appears at least once; we can reduce the general case to $2^c = \mathcal{O}(1)$ instances of this one. We store a c -dimensional grid with each side having length $\lceil \log_{1+\epsilon} n \rceil + 1$. For each point (x_0, \dots, x_{c-1}) in this grid, we store pointers to the nodes in a connected subgraph whose Parikh vector is component-wise between $((1 + \epsilon)^{x_0}, \dots, (1 + \epsilon)^{x_{c-1}})$ and $((1 + \epsilon)^{x_0+1}, \dots, (1 + \epsilon)^{x_{c-1}+1})$. This takes a total of $\mathcal{O}(n \log^c n)$ space. Given the Parikh vector (v_0, \dots, v_{c-1}) of M , we return the subgraph stored for the point $(\lceil \log_{1+\epsilon} v_0 \rceil, \dots, \lceil \log_{1+\epsilon} v_{c-1} \rceil)$ in the grid, if that subgraph exists. We summarize these basic tradeoffs in the following proposition:

Proposition 1. *When G is a graph containing a constant number c of colours there exists an $\mathcal{O}(n^{c+1})$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time. For any positive constant ϵ there exists an $\mathcal{O}(n \log^c n)$ -space index with which in $\mathcal{O}(|M|)$ time we can find an approximate match in which each colour's frequency is within a factor of $1 + \epsilon$ of its frequency in M . When G has bounded treewidth we can build these indexes in polynomial time and, moreover, we can reduce the space of the exact index to $\mathcal{O}(n)$ at the cost of increasing the query time to $|M|^{\mathcal{O}(1)}$.*

When G is a path — which we can think of as a string over an alphabet of c characters — we can improve these bounds. Since G contains $\mathcal{O}(n^2)$ substrings and we can specify any substring by its two endpoints, we can build an $\mathcal{O}(n^2)$ -space index with which we can find a match in $\mathcal{O}(1)$ time. Calculation shows we can reduce the space bound to $\mathcal{O}(n)$ at the cost of increasing the query time to $\mathcal{O}(|M|^c)$, and we can store an approximate index in $\mathcal{O}(\log^c n)$ space.

Suppose G is a string over a constant-size alphabet and $0 < \epsilon \leq 1$. Then in $\mathcal{O}(n^{1+\epsilon})$ expected time we can build an index with which, given a multiset M of characters, we can find all *occ* matches of M in $\mathcal{O}(|M|^{1/\epsilon} + occ)$ worst-case time. To do this, we store G itself and, for $1 \leq m \leq n^\epsilon$, we make a pass over G and store, for each multiset of size m that has a match in G , a list of all the locations of that multiset's matches. Notice the lists for multisets of size m are disjoint and have total length $n - m + 1$; therefore, with dynamic perfect hashing we use a total of $\mathcal{O}(n^{1+\epsilon})$ expected time and $\mathcal{O}(n^{1+\epsilon})$ space. Given a multiset M with $|M| \leq n^\epsilon$, we return our pre-computed list of the locations of M matches in $\mathcal{O}(|M| + occ)$ time, or $\mathcal{O}(occ)$ time if we are given M as a Parikh vector. Given a multiset M with $|M| > n^\epsilon$, we search G in $\mathcal{O}(n) = \mathcal{O}(|M|^{1/\epsilon})$ time.

As an aside, we note that we can extend our approximate indexes to support approximate scaled-then-permuted pattern matching (see [7]). To do this, for each point (x_0, \dots, x_{c-1}) in the grid for which there is no subgraph whose Parikh vector is between $((1 + \epsilon)^{x_0}, \dots, (1 + \epsilon)^{x_{c-1}})$ and $((1 + \epsilon)^{x_0+1}, \dots, (1 + \epsilon)^{x_{c-1}+1})$, we store pointers to the nodes in a connected subgraph (if there is one) whose Parikh vector is a multiple of a one between $((1 + \epsilon)^{x_0}, \dots, (1 + \epsilon)^{x_{c-1}})$ and $((1 + \epsilon)^{x_0+1}, \dots, (1 + \epsilon)^{x_{c-1}+1})$. The query time is still proportional to the size of the match returned but that may now be larger than $|M|$.

3 An Index for Binary Strings

Suppose G is a binary string, i.e., $G[1..n] \in \{0, 1\}^*$. It holds for any m : If there are p copies of 1 in $G[i..i+m-1]$ and r copies of 1 in $G[k..k+m-1]$, then for every value q between p and r there is a position j between i and k such that $G[j..j+m-1]$ contains q copies of 1. This observation was the basis for the index in [8] and is the basis for ours as well.

We store an $\mathcal{O}(1)$ -time rank data structure for G and, for $1 \leq m \leq n$, we store the endpoints of two substrings of length m in G with the most and with the fewest copies of 1. This takes a total of $\mathcal{O}(n)$ space. Given a Parikh vector (v_0, v_1) , we look up the left endpoints i and j of the substrings of length $v_0 + v_1$ in G with the most and with the fewest copies of 1. We set i and j as the initial endpoints for a binary search: at each step, we use two rank queries to find the number q of 1s in $G[\lfloor \frac{i+j}{2} \rfloor .. \lfloor \frac{i+j}{2} \rfloor + v_0 + v_1 - 1]$; if $q = v_1$ then we stop and report this substring by its endpoints; if $q < v_1$ then we set $i = \lfloor (i+j)/2 \rfloor$ and continue; if $q > v_1$ then we set $j = \lfloor (i+j)/2 \rfloor$ and continue. This search takes a total of $\mathcal{O}(\log n)$ time.

Theorem 1. *When G is a path containing only two colours, we can build an $\mathcal{O}(n)$ -space index with which we can find a match in $\mathcal{O}(\log n)$ time.*

4 Exact Indexes for Trees with Two Colours

Suppose G is a tree containing only two colours, black and white. Gagie, Hermelin, Landau and Weimann [13] noted that the observation in Section 3 can be extended to connected graphs: For any m , if there are connected subgraphs H_p and H_r in G with m nodes each and p and r white nodes, respectively, then for every value q between p and r , there is a connected subgraph H_q with m nodes and q white nodes.

To see why, notice that we can construct a sequence of connected subgraphs with m nodes such that the sequence starts with H_p and ends with H_r and any consecutive pair of subgraphs in the sequence differ on two nodes. To build this sequence, we find a path between H_p and H_r . We root H_p and H_r , which are trees themselves, at the first and last nodes in the path (or at a shared node, if they are not disjoint). One by one, we remove nodes bottom-up in H_p and add nodes along the path; remove nodes nearest to H_p in the path and add nodes further along the path; then remove nodes from the path and add nodes top-down in H_r .

Suppose p and r are the minimum and maximum numbers of white nodes in any connected subgraphs of size m , and we store a path consisting of the nodes in H_p in bottom-up order, followed by the nodes in the path, followed by the nodes in H_r in top-down order. If we apply Theorem 1 to this path, then we obtain an $\mathcal{O}(n)$ -space index with which, given the Parikh vector for a multiset M with $|M| = m$, we can find a match in the graph G in $\mathcal{O}(\log n + |M|)$ time. Notice that, if $|M| < \log n$, then we can simply store an $\mathcal{O}(\log^2 n)$ -space lookup

table with which we can find a match in $\mathcal{O}(|M|)$ time. Therefore, applying this construction for $1 \leq m \leq n$, we obtain the following theorem:

Theorem 2. *When G is a tree containing only two colours, we can build an $\mathcal{O}(n^2)$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time.*

When $m \approx n$, we need $\mathcal{O}(n)$ space to store subgraphs with the minimum and maximum numbers of white nodes and the path between them. When $m \ll n$, however, those subgraphs are small and most of the space is taken up by the path. We now show how we can store G such that we can support fast rank queries on paths.

Lemma 1. *We can store G in $\mathcal{O}(n)$ space such that q rank queries on the path between any two nodes take a total of $\mathcal{O}(\log n + q)$ time.*

Proof. We compute the heavy-path decomposition [18] of G and store $\mathcal{O}(1)$ -time rank data structures for each of the heavy paths, which takes $\mathcal{O}(n)$ space. The path between any two nodes u and v is a sequence of $\mathcal{O}(\log n)$ intervals of heavy paths. Given u and v , for each of these intervals we compute the number of white nodes in that interval and to either side of it in the heavy path; this takes a total of $\mathcal{O}(\log n)$ time and rank queries on heavy paths. With this information we can perform any rank query on the path from u to v using a single rank query on a heavy path. \square

If we store G with Lemma 1 and store subgraphs with the minimum and maximum numbers of white nodes only for $1 \leq m \leq n^{1/2}$, then our index takes only $\mathcal{O}(n)$ space but supports queries only for $|M| \leq n^{1/2}$. When $|M| > n^{1/2}$ we can use an algorithm by Gagie et al. to find a match in $\mathcal{O}(|M|n) = \mathcal{O}(|M|^3)$ time.

Corollary 1. *When G is a tree containing only two colours, we can build an $\mathcal{O}(n)$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time when $|M| \leq n^{1/2}$ and in $\mathcal{O}(|M|^3)$ time otherwise.*

5 An Approximate Index for Trees with Two Colours

In this section we present our most technical result, which is how to store in $\mathcal{O}(n)$ space an approximate index for a tree containing only two colours. Again, an approximate match is one whose Parikh vector differs from M 's by a factor of at most $1 + \epsilon$ in each component. (In contrast, with Proposition 1 we would use $\mathcal{O}(n \log^2 n)$ space.) Without loss of generality, assume we are only concerned with multisets in which there are at least as many black nodes as white nodes; we can build a symmetric index for the other case. Notice that in this case, if we can find a connected subgraph H with the same size as the given multiset M and in which the number of white nodes is within a factor of $1 + \epsilon$ of the number in M , then the number of black nodes in H is also within a factor of $1 + \epsilon$ of the number in M .

Our main idea is to store an $\mathcal{O}(n)$ -space data structure with which, given a size m , we can find two connected subgraphs with size m that have approximately the minimum and maximum numbers of white nodes. Suppose we store a subgraph with the minimum number of white nodes for each size that is a power of two and for each size such that the minimum number of white nodes is a factor of $1 + \epsilon$ greater than the number in the preceding stored subgraph. That is, we store a sequence of $\lg n$ subgraphs with total size $\mathcal{O}(n)$ and a sequence of $\log_{1+\epsilon} n$ subgraphs with total size $\mathcal{O}(n \log n)$. The latter sequence of subgraphs has total size $\mathcal{O}(n \log n)$ in the worst case because the minimum number of white nodes may stay low until we reach size nearly n and then increase rapidly, causing us to store about $\log_{1+\epsilon} n$ subgraphs each of size nearly n . However, we can store this sequence of subgraphs in a total of $\mathcal{O}(n)$ space using the following lemma. Similarly, we also store a subgraph with the maximum number of white nodes for each size that is a power of two and for each size such that the maximum number of white nodes is a factor of $1 + \epsilon$ greater than the number in the preceding stored subgraph; this also takes $\mathcal{O}(n)$ total space if we store the subgraphs with the following lemma.

Lemma 2. *We can store G in $\mathcal{O}(n)$ space such that, if G contains a connected subgraph of size m with w white nodes, then we can represent some such subgraph in $\mathcal{O}(w)$ space such that recovering this subgraph takes $\mathcal{O}(m)$ time.*

Proof. We store the adjacency lists for G 's nodes, with each list ordered such that black neighbours precede white neighbours. With this representation, we can expand a subgraph by adding only black nodes as long as this is possible, using $\mathcal{O}(1)$ time per added node.

Let H be a connected subgraph of size m with w white nodes. We store pointers to the white nodes in H , which takes $\mathcal{O}(w)$ space. Since G is a tree, we can find the unique paths between these nodes in a total of m time; notice these paths are contained in H and consist of black nodes. If the subgraph consisting of the white nodes and these paths has fewer than m nodes, then we add black nodes until it has m nodes, which takes a total of $\mathcal{O}(m)$ time. It is possible to add enough black nodes without adding any white nodes because, e.g., we could add the remaining black nodes in H . \square

If we are given a multiset M such that we have subgraphs of size $|M|$ sampled, then we can proceed as in the proof of Theorem 2 and find an exact match if there is one. If we do not have subgraphs of size $|M|$ sampled, then we use our sampled subgraphs to build subgraphs H_{\min} and H_{\max} of size $|M|$ with approximately minimum and maximum numbers of white nodes, then proceed almost as in the proof of Theorem 2: if the number of white nodes H_{\min} is larger but within a factor of $1 + \epsilon$ of the number in M , then we return H_{\min} ; if the number in H_{\min} is more than a factor of $1 + \epsilon$ larger than the number in M , then there is no exact match and we return nothing; if the number of white nodes H_{\max} is smaller but within a factor of $1 + \epsilon$ of the number in M , then we return H_{\max} ; if the number in H_{\max} is less than a factor of $1 + \epsilon$ smaller than the number in M , then there

is no exact match and we return nothing; in all other cases, we proceed as in Theorem 2.

To build H_{\min} we take the next larger subgraph with a minimum number of white nodes and discard nodes until it has size $|M|$ while leaving it connected. This next larger subgraph has size less than $2|M|$, because we sampled for every size that is a power of two; has at most $1 + \epsilon$ times more white nodes than the subgraph of size $|M|$ with the minimum number of white nodes, because we sampled whenever the minimum number of white nodes increased by a factor of $1 + \epsilon$; and is a tree, because it is a connected subgraph of a tree. It follows that discarding nodes takes $\mathcal{O}(|M|)$ time and, since discarding nodes cannot increase the number of white nodes, H_{\min} contains at most $1 + \epsilon$ times the minimum number of white nodes. To build H_{\max} we take the next smaller subgraph with a maximum number of white nodes and add nodes until it has size $|M|$. By symmetric arguments, this takes $\mathcal{O}(|M|)$ time and, since adding nodes cannot decrease the number of white nodes, the maximum number of white nodes in a subgraph of size $|M|$ is at most $1 + \epsilon$ times the number in H_{\max} . Finding the path from H_{\min} to H_{\max} takes $\mathcal{O}(|M|)$ time using the representation from Lemma 1.

Theorem 3. *When G is a tree containing only two colours, for any positive constant ϵ we can build an $\mathcal{O}(n)$ -space index with which in $\mathcal{O}(|M|)$ time we can find an approximate match in which each colour's frequency is within a factor of $1 + \epsilon$ of its frequency in M .*

6 Indexes for Graphs with Two Colours

Suppose G is a graph containing only two colours, black and white. Theorem 2 applies in this case as well, if we consider spanning trees of H_p and H_r instead of the connected subgraphs themselves, but we can build the index in polynomial time only in special cases, such as when G has bounded treewidth.

Theorem 4. *When G is a graph containing only two colours, there exists an $\mathcal{O}(n^2)$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time.*

When G has bounded treewidth we can find a match in $n^{\mathcal{O}(1)}$ time, so we can prove a weaker version of Corollary 1 for graphs. To do this, we build a spanning tree for G and apply Lemma 1 to that spanning tree. We can build the resulting index in polynomial time.

Corollary 2. *When G is a graph with bounded treewidth containing only two colours, we can build an $\mathcal{O}(n)$ -space index with which we can find a match in $\mathcal{O}(|M|)$ time when $|M| \leq n^{1/2}$ and in $|M|^{\mathcal{O}(1)}$ time otherwise.*

We cannot quite extend Theorem 3 to graphs because Lemma 2 does not apply (since we may not be able to recover the correct paths between the white nodes in the subgraph H). However, if we store connected subgraphs explicitly instead of with Lemma 2, then calculation shows the index takes $\mathcal{O}(n \log n)$ space. Again, this index can be built in polynomial time when G has bounded treewidth.

Theorem 5. *When G is a graph containing only two colours, for any positive constant ϵ we can build an $\mathcal{O}(n \log n)$ -space index with which in $\mathcal{O}(|M|)$ time we can find an approximate match in which each colour's frequency is within a factor of $1 + \epsilon$ of its frequency in M .*

References

1. Amir, A., Apostolico, A., Landau, G.M., Satta, G.: Efficient text fingerprinting via Parikh mapping. *J. Discrete Algorithms* 1(5-6), 409–421 (2003)
2. 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)
3. Benson, G.: Composition alignment. In: Benson, G., Page, R.D.M. (eds.) *WABI 2003. LNCS (LNBI)*, vol. 2812, pp. 447–461. Springer, Heidelberg (2003)
4. Böcker, S.: Sequencing from compomers: Using mass spectrometry for DNA de novo sequencing of 200+ nt. *J. of Computational Biology* 11(6), 1110–1134 (2004)
5. Böcker, S.: Simulating multiplexed SNP discovery rates using base-specific cleavage and mass spectrometry. *Bioinformatics* 23(2), 5–12 (2007)
6. Burcsi, P., Cicalese, F., Fici, G., Lipták, Z.: Algorithms for jumbled pattern matching in strings. *Int. J. Found. Comput. Sci.* 23(2), 357–374 (2012)
7. Butman, A., Eres, R., Landau, G.M.: Scaled and permuted string matching. *Inf. Process. Lett.* 92(6), 293–297 (2004)
8. Cicalese, F., Fici, G., Lipták, Z.: Searching for jumbled patterns in strings. In: *Proc. Prague Stringology Conference (PSC 2009)*, pp. 105–117 (2009)
9. 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)
10. Didier, G., Schmidt, T., Stoye, J., Tsur, D.: Character sets of strings. *J. Discrete Algorithms* 5(2), 330–340 (2007)
11. Eres, R., Landau, G.M., Parida, L.: Permutation pattern discovery in biosequences. *Journal of Computational Biology* 11(6), 1050–1060 (2004)
12. 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)
13. Gagie, T., Hermelin, D., Landau, G.M., Weimann, O.: Binary jumbled pattern matching on trees and tree-like structures. In: Bodlaender, H.L., Italiano, G.F. (eds.) *ESA 2013. LNCS*, vol. 8125, pp. 517–528. Springer, Heidelberg (2013)
14. Giaquinta, E., Grabowski, S.: New algorithms for binary jumbled pattern matching. *Inf. Process. Lett.* 113(14-16), 538–542 (2013)
15. Kociumaka, T., Radoszewski, J., Rytter, W.: Efficient indexes for jumbled pattern matching with constant-sized alphabet. In: Bodlaender, H.L., Italiano, G.F. (eds.) *ESA 2013. LNCS*, vol. 8125, pp. 625–636. Springer, Heidelberg (2013)
16. 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)
17. Moosa, T.M., Rahman, M.S.: Sub-quadratic time and linear space data structures for permutation matching in binary strings. *J. Discr. Alg.* 10, 5–9 (2012)
18. Sleator, D.D., Tarjan, R.E.: A data structure for dynamic trees. *J. Comput. Syst. Sci.* 26(3), 362–391 (1983)