


A Compact DAG for Storing and Searching Maximal Common Subsequences

Alessio Conte ✉ 
Università di Pisa, Italy

Roberto Grossi ✉ 
Università di Pisa, Italy

Giulia Punzi ✉ 
National Institute of Informatics, Tokyo, Japan

Takeaki Uno ✉ 
National Institute of Informatics, Tokyo, Japan

Abstract

Maximal Common Subsequences (MCSs) between two strings X and Y are subsequences of both X and Y that are maximal under inclusion. MCSs relax and generalize the well known and widely used concept of Longest Common Subsequences (LCSs), which can be seen as MCSs of maximum length. While the number both LCSs and MCSs can be exponential in the length of the strings, LCSs have been long exploited for string and text analysis, as simple compact representations of all LCSs between two strings, built via dynamic programming or automata, have been known since the '70s. MCSs appear to have a more challenging structure: even listing them efficiently was an open problem open until recently, thus narrowing the complexity difference between the two problems, but the gap remained significant. In this paper we close the complexity gap: we show how to build DAG of polynomial size – in polynomial time – which allows for efficient operations on the set of all MCSs such as enumeration in Constant Amortized Time per solution (CAT), counting, and random access to the i -th element (i.e., rank and select operations). Other than improving known algorithmic results, this work paves the way for new sequence analysis methods based on MCSs.

2012 ACM Subject Classification Mathematics of computing → Combinatorial algorithms; Information systems → Structured text search

Keywords and phrases Maximal common subsequence, DAG, Compact data structures, Enumeration, Constant amortized time, Random access

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2023.21

Related Version *Previous Version*: <https://arxiv.org/abs/2307.13695>

Funding *Alessio Conte*: Partially supported by MUR PRIN Project n. 2022TS4Y3N – EXPAND.

Roberto Grossi: Work partially supported by MUR PRIN Project n. 2022TS4Y3N – EXPAND.

Giulia Punzi: Work partially supported by JSPS KAKENHI Grant Number JP20H05962.

Takeaki Uno: Work partially supported by JSPS KAKENHI Grant Numbers JP20H05962, JP20H00595.

Acknowledgements We thank the anonymous Referees for their comments, leading us to the current version of Theorem 13.

1 Introduction

The *Longest Common Subsequence* (LCS) [4, 12, 16, 25] have thoroughly been studied in a plethora of string comparison application domains, like spelling error correction, molecular biology, and plagiarism detection, to name a few. The LCS is a special case of *Maximal Common Subsequence* (MCS) for any two strings X, Y : it is a string S that is a subsequence of both X and Y , and is inclusion-maximal, namely, no other string S' containing S is



© Alessio Conte, Roberto Grossi, Giulia Punzi, and Takeaki Uno;
licensed under Creative Commons License CC-BY 4.0

34th International Symposium on Algorithms and Computation (ISAAC 2023).

Editors: Satoru Iwata and Naonori Kakimura; Article No. 21; pp. 21:1–21:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

also a common subsequence of X and Y . The set of all MCSs for X and Y is denoted by $MCS(X, Y)$. For example, $MCS(X, Y) = \{\text{TACA}, \mathbf{G}\}$ for $X = \text{TCACAG}$ and $Y = \text{GTACTA}$, whereas $\text{TCA} \notin MCS(X, Y)$ as it is contained in TACA , and the latter is also the only LCS. A crucial observation is the following: if S is a common subsequence, it is always contained in some MCS but this is not necessarily true for LCS (e.g. $S = \mathbf{G}$ is not contained in TACA).

In general, LCSs only provide us with information about the longest possible alignment between two strings, while MCSs offer a different range of information, possibly revealing slightly smaller but alternative alignments. For very long strings an MCS may be just slightly shorter than an LCS but provide information on parts of the strings not found by LCSs. In principle, MCS could provide helpful information in all the applications where LCS are used.

While there is a quadratic conditional lower bound for the computation of LCS, based on the Strong Exponential Time Hypothesis [1], no such bound exists for MCS: actually, an MCS between two strings of length n can be extracted in $O(n\sqrt{\log n / \log \log n})$ time [22]. Moreover it is NP-hard to compute the LCS of an arbitrary number of strings [15], whereas a recent polynomial-time algorithm for extracting an MCS from multiple strings exists [11]. It is worth noting that even though there are a few more different approaches in the literature to find LCS or common subsequences with some kind of constraints (e.g. common subsequence trees [13], common subsequence automata [7]), the above observations motivate further investigation to directly deal with MCS.

In this paper we proceed along that direction, and consider the problem of storing and searching the set $MCS(X, Y)$. The main hurdle is that $MCS(X, Y)$ could contain an exponential number of distinct strings [6]: consequently, any trie-based or immediate automaton representation would require *exponential* time and space for its construction. We improve significantly over this direction as we list in our contributions, where $n = \max\{|X|, |Y|\}$ and σ is the size of the alphabet Σ for X and Y :

- We introduce a labeled compact direct acyclic graph $\text{MDAG}(X, Y)$ (for *MCS DAG*) that represents the strings in $MCS(X, Y)$ in polynomial space $O(n^3\sigma)$ and can be built in polynomial time $O(n^3\sigma \log n)$. The previously mentioned conditional lower bound for LCS applies, since the longest path here is an LCS of X and Y .
- For any string P , we show how $\text{MDAG}(X, Y)$ can *search* the strings with prefix P from $MCS(X, Y)$ and report them in lexicographic order, in $O(|P| \log \sigma + occ)$ time, where occ is the number of reported strings.
- We can list all the strings from $MCS(X, Y)$ lexicographically in constant amortized time (CAT), namely, $O(|MCS(X, Y)|)$ time.
- By adding $O(n)$ -bits per node, we show how $\text{MDAG}(X, Y)$ with input i (where $1 \leq i \leq |MCS(X, Y)|$) can *select* the i th string S in lexicographic order from $MCS(X, Y)$, in $O(|S| \log \sigma)$ operations on $O(n)$ -bit integers; the inverse operation of *rank* for input S is supported in the same complexity, where i is returned.

In Section 3 we implement $\text{MDAG}(X, Y)$ as a direct acyclic graph (DAG) where the only zero indegree node is the source s and the only zero outdegree node is the target t . Each edge is labeled with a character from the alphabet, so that any two outgoing edges from the same node have different characters as labels. Moreover, each st -path corresponds to a string in $MCS(X, Y)$, obtained by concatenating the characters on the edge traversed along the st -path; vice versa, each string in $MCS(X, Y)$ is spelled out by an st -path. In order to define and build $\text{MDAG}(X, Y)$ we use some properties on MCSs previously introduced in [6] and described in Section 2, plus some new ideas to keep the size of MDAG polynomial. As a result, after MDAG has been built and its unary paths have been compacted, the aforementioned operations can be implemented in a simple way, as discussed in Section 4.

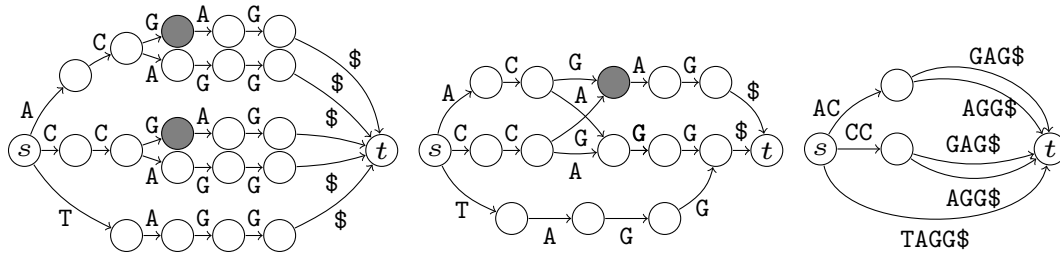


Figure 1 Left: edge-labeled DAG for strings $X = TCACAGAGA$ and $Y = ACCCGTAGG$, where each st -path corresponds to a string in $MCS(X, Y) = \{ACAGG, ACGAG, CCAGG, CCGAG, TAGG\}$. Center: MDAG for the same strings X and Y ; for instance, the two shaded nodes in the left graph were compacted into one, shaded in the right. Right: its compact MDAG, where $MCS(X, Y)$ is incrementally output as $push(AC)$, $push(AGG)$, $pop()$, $push(GAG)$, $pop()$, $pop()$, $push(CC)$, $push(AGG)$, $pop()$, $push(GAG)$, $pop()$, $pop()$, $push(TAGG)$, $pop()$, and constant-space references to string labels are the arguments of $push()$. Note that before each $pop()$ immediately following a $push()$, we find on the stack the string labels whose concatenation gives a string from $MCS(X, Y)$.

An example of MDAG and compact MDAG is shown in Figure 1, where $MCS(X, Y)$ is enumerated in CAT using constant-space references to the strings that label the edges in $MDAG(X, Y)$, by simply using a stack [23]. As formalized by Frank Ruskey [20, Section 1.7] in the “Don’t Count the Output Principle”, we account for the amount of data change that takes place, and the latter takes CAT per solution.

Related work. Maximal common subsequences were first introduced in [9], in the context of LCS approximation. The first algorithm for finding an MCS between two strings was presented in [21], and subsequently refined in [22]. The latter algorithm finds an MCS between two strings of length n in $O(n\sqrt{\log n/\log \log n})$. These algorithms can be also used to extend a given sequence to a maximal one in the same time, and to check whether a given subsequence is maximal in $O(n)$ time. In [11] the authors considered the problem of finding an MCS of $m > 2$ strings of total length n , and were able to solve this in $O(mn \log n)$ time and $O(n)$ space. As for MCS enumeration, [6] showed that this task can be performed in $O(n \log n)$ delay and quadratic space.

The automaton approach has been used in literature to deal with subsequence-related problems. The Directed Acyclic Subsequence Graph (DASG) introduced in [3] is an automaton which accepts all subsequences of a given string S . Given a set of strings, it can also be generalized to accepting subsequences of any string in the set. Later, the common subsequence automata (CSA) was introduced [7, 8, 24], which instead accepts *common* subsequences of a set of strings. Such automaton is similar to the common subsequence tree of [13], and it can also be used to find a longest common subsequence between two strings [17]. Binary decision diagrams such as ZDD [18] and SeqBDD [14] could potentially be employed to compactly store $MCS(X, Y)$ but their worst-case behaviour typically involves exponential construction time and space.

We stress that it is not straightforward to efficiently adapt all of these structures to the MCS problem: indeed, figuring out which of the accepted strings are subsequences of each other is not an immediate task. We therefore opted for a different approach, directly defining and constructing an MCS automaton, based on combinatorial properties specific to MCSs (during the review period of this manuscript, a pre-print on MCS automata appeared [10].)

Preliminaries. [*String notation*] A string S over an alphabet Σ (of size $\sigma = |\Sigma|$) is a concatenation of any number of its characters. The empty string is denoted with ε . For $i \in [0, |S| - 1]$, character $S[i]$ occurs at position i of string S , and position $next_S(c, i)$ denote the next occurrence of character c after position i in string S , if it exists (otherwise, $next_S(c, i) = |S| - 1$). The notation $S_{<i}$ indicates $S[0, i - 1]$, and $S_{\leq i}$ indicates $S[0, i]$. We say that S is a *subsequence* of a string X , denoted $S \subset X$, if there exist indices $0 \leq i_0 < \dots < i_{|S|-1} < |X|$ such that $X[i_k] = S[k]$ for all $k \in [0, |S| - 1]$. In this case we also say that X *contains* S . Given two strings X and Y , a string S is a common subsequence if $S \subset X$ and $S \subset Y$. Furthermore, S is a *Maximal Common Subsequence*, denoted $S \in MCS(X, Y)$, if it is a common subsequence which is *inclusion-maximal*: there is no string $T \neq S$ such that $S \subset T \subset X, Y$. [*Graph notation*] Given a directed graph $G = (V, E)$, we define the in-neighbors of a node $v \in V$ as $N^-(v) = \{u \in V \mid (u, v) \in E\}$, and the out-neighbors as $N^+(v) = \{z \in V \mid (v, z) \in E\}$; the in-degree is $d^-(v) = |N^-(v)|$ and the out-degree is $d^+(v) = |N^+(v)|$. A subgraph of graph $G = (V, E)$ is a graph $H = (W, F)$ such that $W \subseteq V$ and $F \subseteq E$. A *path* P in a graph $G = (V, E)$ is a sequence of distinct adjacent nodes: $P = u_1 \dots u_k$ where $u_i \in V$ for all i , $u_i \neq u_j$ for all $i \neq j$, and $(u_i, u_{i+1}) \in E$ for all i . P is called an $u_1 u_k$ -path. A path where the first and last nodes coincide is called a *cycle*. A directed graph with no cycles is called a directed acyclic graph, or *DAG*.

2 Structure of MCSs

In this section, we show known properties of MCSs which will be at the base of MDAG. Firstly, we show how to naturally represent all the strings in $MCS(X, Y)$ for any two given strings X and Y , using a finite-state automaton A corresponding to a DAG with a single source s and a single target t , where each edge is labeled with a character of the alphabet Σ ; we show that st -paths in A have a one-to-one correspondence with the strings in $MCS(X, Y)$. Secondly, we recall and contextualize some useful properties of MCSs proven in [6].

2.1 Modeling MCS as an Exponentially Large DAG

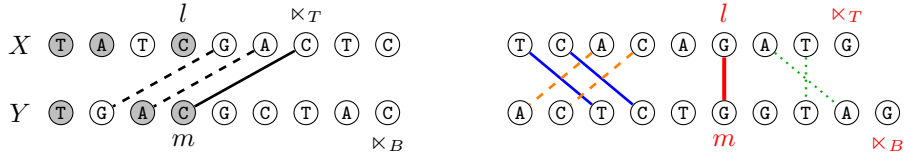
We here define a deterministic finite-state automaton A that accepts the strings in $MCS(X, Y)$. It is more convenient to define it directly as an equivalent DAG, using the extended alphabet $\Sigma' = \Sigma \cup \$$, where the dollar sign is a new special character to identify the end of an MCS.

► **Definition 1.** *The edge-labeled DAG $A = (V \cup \{s, t\}, E, \ell(\cdot))$ is defined for two input strings X and Y to store their set $MCS(X, Y)$, as follows:*

- *Every node in $V \cup \{s\}$ corresponds to a distinct prefix of a string in $MCS(X, Y)$, with s called source node corresponding to the empty string prefix.*
- *For any two nodes $u, u' \in V$, let P, P' be their corresponding two prefixes, and let $c \in \Sigma$ be a character. Then, $P' = P c$ iff $(u, u') \in E$ with label $\ell(u, u') = c$.*
- *For any node $u \in V$, if its corresponding prefix is a whole string $P \in MCS(X, Y)$, then $(u, t) \in E$ with label $\ell(u, t) = \$$.*

Given DAG A , we have a bijection between the strings in $MCS(X, Y)$ and the labeled st -paths (see the left of Figure 1 for an example). Indeed, it is immediate that any such path corresponds to a string in $MCS(X, Y)$ due to the way edges are placed. Vice versa, by definition, we cannot have two out-going edges from the same node that share the same label. Therefore, each MCS has exactly one corresponding st -path. Note that the property of the outgoing labels also implies that $d^+(u) \leq \sigma$ for all $u \in V$, and therefore $|E| \leq \sigma|V|$. As the number of nodes in A is $\Omega(|MCS(X, Y)|)$, its size is exponential in the worst case.

21:6 A Compact DAG for Storing and Searching Maximal Common Subsequences



■ **Figure 2** Left: Swings (\times_T, \times_B) for valid prefix TAC in strings $X = \text{TATCGACTC}$ and $Y = \text{TGACGCTAC}$. Consider swing \times_T : $\text{TAC} \notin \text{MCS}(X_{\leq \times_T}, Y_{\leq m})$ since TGAC is a common subsequence (dashed). Note how the bottom swing \times_B is not given by the next occurrence of **C** after m , since TAC is still maximal for strings TATC and TGACGC . Right: Two prefixes, TCG (solid blue) and ACG (dashed orange), both ending at solid red positions (l, m) , and having the same swings (\times_T, \times_B) . The valid extensions are the same (dotted green): **A** and **T**.

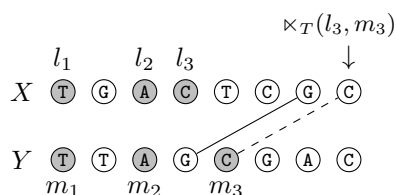
It follows from the definition that $P \in \text{MCS}(X_{<i}, Y_{<j}) \iff i \leq \times_T(P)$ and $j \leq \times_B(P)$. Given the swings, condition 1 can thus be checked in constant time [6]. Note that the definition of top swing necessarily implies $X[i] = X[l]$.¹ The symmetric holds for the bottom swing.

We briefly also recall how to incrementally compute the swings of a prefix, first described in [6], as it will be useful in our proofs. We only describe the procedure for the top swing as the bottom swing is symmetrical:

- If P is composed of a single character c , with first occurrence in X at position l and in Y at position m , then it suffices to compute, for every character d in $Y_{<m}$, the first occurrence of d in $X_{>l}$, and take the minimum of these. The swing then corresponds to the first occurrence of c after such minimum.
- Let $P = p_1 \cdots p_N$ be a valid prefix with $N > 1$, and let l_1, \dots, l_N (resp. m_1, \dots, m_N) be the positions of X (resp. Y) such that $X_{\leq l_i}$ (resp. $Y_{\leq m_i}$) is the shortest prefix containing $p_1 \cdots p_i$. The *personal top swing* $\times_T(l_N, m_N)$ of the last position is the top swing of character p_N when seen as a prefix over the strings $X_{>l_{N-1}}, Y_{>m_{N-1}}$, instead of over the whole strings (and thus computed as above). The personal bottom swing is defined analogously. In other words, the personal swing of a character expresses the change necessary to have an insertion between itself and the previous character of the prefix (see Figure 3 for an example). The top swing of P is the minimum between the personal swing of (l_N, m_N) , and the first occurrence of $Y[m_N] = p_N$ after the top swing of prefix $p_1 \cdots p_{N-1}$. This second swing indicates the change required for an insertion in the previous part of the prefix.

Updating the swings when a character is added can be done in $O(\sigma)$ time, provided that we can find in constant time the next occurrence of a character c after a given position in the strings; more importantly, this update does *not* need knowledge of the whole prefix, but just the positions its final character (l_N, m_N) and their current swings.

¹ Consider any $h > l$ such that $P \notin \text{MCS}(X_{\leq h}, Y_{\leq m})$. Since $P = p_1 \cdots p_s \in \text{MCS}(X_{\leq l}, Y_{\leq m})$ and we are only extending string X , P can only become non-maximal if an insertion occurs between p_i and p_{i+1} , for $i < s$. If h' is the last occurrence of $p_s = X[l]$ before h , it is $P \notin \text{MCS}(X_{\leq h'}, Y_{\leq m})$, as the substring of X between h' and h does not contain a suffix of P as subsequence. The minimum index $i > l$ such that $P \notin \text{MCS}(X_{\leq i}, Y_{\leq m})$ must satisfy $X[i] = X[l]$.



■ **Figure 3** Personal top swing for edge (l_3, m_3) , corresponding to the last character of the leftmost mapping of prefix TAC. Indeed, such swing allows for insertion of character G after (l_2, m_2) .

3 Polynomial-Size MDAG

The construction of DAG A satisfying the conditions of Definition 1 would require exponential time and space: the number of nodes of A is between $\Omega(|MCS(X, Y)|)$ and $O(n|MCS(X, Y)|)$, so it can be exponential in n . In this section, we show how to obtain MDAG, where we still have a bijection between st -paths and MCS, but which can instead always be constructed in $O(n^3\sigma \log n)$ time and $O(n^3\sigma)$ space, as per Theorem 13. Intuitively, the relevant information discussed for A are the quadruples (l, m, t, b) , where $X_{\leq l}$ and $Y_{\leq m}$ are some prefixes and pair t, b is some swing: these quadruples are the candidates for being nodes in MDAG.

The formal definition of MDAG is based on an equivalence relation over the nodes of A , given in Section 3.1. Afterwards, in Section 3.2, we describe an algorithm for *directly* constructing MDAG. We present the complexity bounds for the construction in Section 3.3.

3.1 Equivalence Relation for Defining MDAG

Let A be a DAG as defined in Definition 1. Our construction algorithm for MDAG is based on the concepts from Section 2.2. The idea is to use the characterization of valid extensions to identify the out-neighbors of a given node of DAG A . We identify an equivalence relation over the prefixes of $MCS(X, Y)$, and thus on the nodes of A , that allow us to always build MDAG in polynomial time and space. We begin with the following lemma:

► **Lemma 4.** *Given any valid prefix P , let $X_{\leq l}$ and $Y_{\leq m}$ be the shortest prefixes containing P , and $\times(P) = (t, b)$ be its swing. Consider another valid prefix $P' \neq P$ with the same shortest prefixes $X_{\leq l}, Y_{\leq m}$ and swing $\times(P') = \times(P)$ as P . Then, the set of valid extensions is the same for both P and P' , and for each valid extension $c \in \Sigma$, the swings of Pc are the same as the ones of $P'c$.*

Proof. The definition of $Ext_{l,m}$ only depends on the value of l and m , therefore such set is the same for both P and P' . Since the swings for P and P' are equal, the set of valid extensions is necessarily the same. Let now $c \in \Sigma$ be a valid extension for P and P' . Let $X_{\leq l_c}$ and $Y_{\leq m_c}$ respectively be the shortest prefixes of X and Y containing Pc . These are also the shortest prefixes containing $P'c$: the shortest prefixes containing P and P' were the same, and l_c is simply the first occurrence of c after l , analogously for m_c and m . The swings of Pc are given by the minimum of the swings of P , and the personal swing $\times(l_N, m_N)$ obtained by adding the new character c . The latter personal swing is the same for both Pc and $P'c$, since we are considering the same positions l_c, m_c . Since the previous swings were also equal, this means that the swings of Pc and $P'c$ are indeed the same. ◀

Lemma 4 has an implication for prefixes P and P' that share the *same* swing: if M_1, \dots, M_N are strings extending as $PM_i \in MCS(X, Y)$, and M'_1, \dots, M'_M extending as $P'M'_i \in MCS(X, Y)$, then they are equal: $\{M_i \mid i = 1, \dots, N\} = \{M'_i \mid i = 1, \dots, M\}$.

Given a node u of A , let P be the corresponding prefix. We assign u the quadruple of parameters $ID(u) = \langle l, m, \times_T(P), \times_B(P) \rangle$, where l and m are such that $X_{\leq l}$ and $Y_{\leq m}$ are the shortest prefixes containing P . By Lemma 4, this tuple completely identifies the valid extensions of P , which means that it completely identifies the neighbors of node u .

► **Corollary 5.** *Let $u \neq u'$ with $ID(u) = ID(u')$. Then, for each $v \in N^+(u)$ there exists exactly one $v' \in N^+(u')$ such that $ID(v) = ID(v')$ and $\ell(u, v) = \ell(u', v')$.*

Therefore, we can define the following *equivalence relation* on the nodes of A : $u \sim u'$ if and only if $ID(u) = ID(u')$. We can then identify a class of equivalent nodes in the DAG, choosing one representative for it. Because of Corollary 5, this does not change the set of labeled st -paths of the DAG: the nodes that are identified as one have the same labelled out-edges, leading to the same out-neighbors. Our data structure MDAG is then defined as the DAG resulting from this identification:

► **Definition 6.** *Data structure MDAG is a node- and edge-labelled DAG built as follows:*

1. Start from DAG A (Definition 1). For each node u , consider its (unique) corresponding prefix P , and let $X_{\leq l}$ and $Y_{\leq m}$ be the shortest prefixes of X and Y containing P , and $\times(P) = (t, b)$. Assign to node u the node-label $ID(u) = \langle l, m, t, b \rangle$.
 2. Merge every pair of nodes $u \neq u'$ with the same label $ID(u) = ID(u')$ into one node.
- An example of such DAG is shown in the right of Figure 1.

It is possible to further compress the MDAG, as detailed in the next section, by compacting nodes of out-degree 1. This will not change its worst-case size, but will impact the efficiency of our method.

3.2 Direct and Incremental Construction of MDAG

We build MDAG directly, without the intermediate DAG A . We apply the incremental procedure below, in a DFS fashion, using the node ID s to avoid repeated computation. At any moment we have built a node- and edge-labeled DAG $H = (W, F)$. A node $u \in W$ corresponds to a set of prefixes P_1, \dots, P_k , given by the concatenation of the edge-labels of all su -paths using edges of F . All prefixes P_i share the same ending positions of the shortest prefixes of X and Y that contain them, and the corresponding swings; these four values form the label $ID(u)$ assigned to u .

Every recursive call $BUILDDAG(u)$ takes as input a node u which belongs to the current DAG H , and expands DAG H accordingly as follows:

1. Let $ID(u) = \langle l, m, t, b \rangle$. First, compute set $Ext_{l,m}$, and use it to compute the valid extensions: select characters c that have a corresponding pair $(i, j) \in Ext_{l,m}$, with $i \leq t$ and $j \leq b$.
2. For such character c , compute the positions (l_c, m_c) such that $X_{\leq l_c}$ and $Y_{\leq m_c}$ are the shortest prefixes containing Pc , and update the swings t_c, b_c .
3. Now, check if the DAG H generated so far already has a node with label $\langle l_c, m_c, t_c, b_c \rangle$:
 - a. If such a node $v \in W$ exists, then simply add edge (u, v) with label c to the edges F of H , without recursing. Indeed, a recursive call for v has been previously performed.
 - b. Otherwise, add node v to W , with $ID(v) = \langle l_c, m_c, t_c, b_c \rangle$, and add edge (u, v) to F , with label c . Then, perform the recursive call $BUILDDAG(v)$.

► **Corollary 7 (Correctness).** *$BUILDDAG(s)$ correctly builds $MDAG(X, Y)$ starting from $H = (\{s\}, \emptyset)$.*

Proof. We show that, at every step, H is a subgraph of DAG A where nodes with equal $ID()$ values have been identified. This is true at the beginning, when $H = \{s\}$. If this holds at the beginning of a recursive call for a node u , then it must hold at the end. Indeed, we use valid extensions to identify neighbors of a given node, which is also the definition of neighbors for a node of A . For each valid extension c , we compute the label $\langle l_c, m_c, t_c, b_c \rangle$ of the corresponding node. Now, if there already exists a node v with such label, we add an edge between u and v , with label c . This correctly identifies the two nodes with equal labels as being the same node, as per operation 2. Otherwise, the new node must be added, with the correct label $\langle l_c, m_c, t_c, b_c \rangle$ as per operation 1, since it represents a new prefix. ◀

Once we have built MDAG, we can easily compute compact MDAG in linear time in the size of MDAG, in the following way. Let us proceed in topological order of the nodes; when a node v with $N^+(v) = \{w\}$ (i.e. out-degree 1) is encountered, we perform the following:

1. For each $u \in N^-(v)$, remove edge (u, v) and add edge (u, w) with label $\ell(u, v)\ell(v, w)$.
 2. After all in-neighbors have been processed, remove node v and edge (v, w) from MDAG.
- Clearly, the minimum out-degree of is 2, and nodes s and t are never removed.

Complexity. Let us now study the time and space complexity of the procedure. As mentioned before, finding $Ext_{l,m}$ requires $O(\sigma \log n)$ time (see [6] for details). Checking whether each c corresponding to an element of $Ext_{l,m}$ satisfies the swings' condition requires constant time per character. For each such c , computing positions (l_c, m_c) can be done in constant time, using appropriate data structures as outlined next. In order to attain our goal, we only need to ensure constant-time queries for the next occurrence of character c after a given position i . To this end, let us keep two bit-vectors for each $c \in \Sigma$, one for X and one for Y , indicating the positions in which c occurs in the strings. By equipping these vectors with rank and select data structures, which employ $O(n\sigma)$ space, we can find in constant time the next occurrence of any character after a given position [19]. Performing this operation in constant time also allows us to update the swings in $O(\sigma)$ time, as we have explained in Section 2.2. All operations described so far are performed exactly once per node. Therefore, the total time required for Steps 1 and 2 is $O(|V|\sigma \log n)$.

Let us now consider Step 3. We need to check if a node v belongs to the current DAG, and add it if it does not. To be able to efficiently perform these operations, let us keep and dynamically update a bit matrix for pairs l, m , where 1 occurs if that pair currently corresponds to at least one node. Then, each cell filled with a one has an associated balanced binary search tree, which indexes the pairs of swings (t, b) such that there currently is a node $u \in W$ with $ID(u) = \langle l, m, t, b \rangle$. These pairs are ordered according to the total order: $(t, b) < (t', b')$ if and only if $t < t'$ or $t = t'$ and $b < b'$. Lookup and insertions in such data structure require $O(\log n)$ each, and the total space employed is $O(|V|)$. Now, note that we perform a membership check, with subsequent possible insertion, exactly once per edge of the DAG. Recalling that $|E| \leq \sigma|V|$, the total time required for Step 3 is again $O(|E| \log n) = O(|V|\sigma \log n)$. Thus, summing the time required for Steps 1, 2 and 3 yields $O(|V|\sigma \log n)$ total time for the algorithm.

To be able to give final complexity bounds, we therefore need bounds on the size of the DAG we constructed. Trivially, we can bound $|V| = O(n^4)$, since no two nodes share the same ID , and the number of different ID s is bounded by n^4 . Therefore, we surely have a polynomial-time and space algorithm for building a DAG. We can actually do better than this: thanks to some properties of the swings, in the next section we show that $|V| = O(n^3)$, leading to the complexity bounds given in Theorem 13.

3.3 Cubic Size of the MCS DAG

To conclude the proof of Theorem 13, we need to study the size of $\text{MDAG}(X, Y)$ as constructed in Section 3.2. We prove a monotonicity property of the swing values, which will allow us to show that the number of nodes of $\text{MDAG}(X, Y)$ is bounded by $O(n^3)$.

In Section 2.2, we saw that the top swing of a valid prefix P is defined as $\times_T(P) = \min\{i > l \mid P \notin \text{MCS}(X_{\leq i}, Y_{\leq m})\}$, where $X_{\leq l}$ and $Y_{\leq m}$ are the shortest prefixes of respectively X and Y containing P . In other words, if we start from strings $X_{\leq l}, Y_{\leq m}$ (where P is obviously maximal), it is the minimum extension of string X that ensures at least one insertion in P . Symmetrical definition holds for bottom swings, by switching the two strings.

Recall that, as seen in Section 2.2, if $\times_T(P) = t$ then $X[t] = Y[m] = X[l]$: the swings' positions are occurrences of the last character of the prefix. We also note the following, which follows from the definition of swings:

► **Remark 8.** Let $P = p_1 \cdots p_N$ a valid prefix with swings $\langle t, b \rangle$. Let $X_{\leq l_i}$ and $Y_{\leq m_i}$ be the shortest prefixes respectively of X and Y that contain $p_1 \cdots p_i$. Then, there is at least one match between $Y[m_{N-1}, m_N]$ and $X[l_N, t]$. More specifically, there can either be a match between $Y[m_{N-1}, m_N]$ and $X[l_N, t]$, or between $Y[m_{N-1}]$ and $X[l, t]$ which will lead to an insertion in a previous part of the prefix.

We now present some new swing properties. This lemma proves that, when two prefixes are extended with a valid extension that occurs at the same pair of positions, then the relative order of the swings remains unchanged during the extension:

► **Lemma 9.** *Let u and u' be two nodes of MDAG, with $ID(u) = \langle x, y, \times_T(P), \times_B(P) \rangle$ and $ID(u') = \langle x', y', \times_T(P'), \times_B(P') \rangle$ such that $x \neq x'$ or $y \neq y'$, where P (resp. P') is any prefix associated to u (resp. u'). Assume that we have $v \in N^+(u)$ with $ID(v) = \langle l, m, \times_T(Pc), \times_B(Pc) \rangle$, and $v' \in N^+(u')$ with $ID(v') = \langle l, m, \times_T(P'c), \times_B(P'c) \rangle$ (i.e. same positions l, m corresponding to character c). Then, the swings change monotonically:*

- $\times_T(P) < \times_T(P') \Rightarrow \times_T(Pc) \leq \times_T(P'c)$
- $\times_B(P) < \times_B(P') \Rightarrow \times_B(Pc) \leq \times_B(P'c)$

Proof. We prove the result for top swings. Let $\times_T(P) = t$ and $\times_T(P') = t'$, with $t < t'$. We note that for (l, m) to be a valid extension for both prefixes, we must have $l \leq t < t'$ (Swing condition 1 for valid extensions' characterization in Section 2.2). By the incremental computation of swings, the swings of Pc and $P'c$ are computed by taking the minimum between the personal swing $\times_T^{(l, m)}$ of the new positions, and the next occurrence of the corresponding character after the top swings t, t' . More specifically, $\times_T(Pc) = \min\{\times_T^{(l, m)}, \text{next}_X(c, t)\}$ and $\times_T(P'c) = \min\{\times_T^{(l, m)}, \text{next}_X(c, t')\}$. Since the first component of the minimum is the same, it suffices to prove that $\text{next}_X(c, t) \leq \text{next}_X(c, t')$ to conclude $\times_T(Pc) \leq \times_T(P'c)$. Indeed, since $t < t'$ are positions in the same string, the next occurrence of a given character after t cannot be strictly bigger than the next occurrence of the same character after t' . Thus, we have proved the claim for top swings; bottom swings are symmetrical. ◀

The next corollary shows that the opposite implication holds for strict inequalities:

► **Corollary 10.** *Under the same hypotheses of Lemma 9, we have*

- $\times_T(Pc) < \times_T(P'c) \Rightarrow \times_T(P) < \times_T(P')$
- $\times_B(Pc) < \times_B(P'c) \Rightarrow \times_B(P) < \times_B(P')$

Proof. We reverse the proof of Lemma 9. Consider top swings, and recall $\times_T(Pc) = \min\{\times_T^{(l,m)}, \text{next}_X(c, t)\}$ and $\times_T(P'c) = \min\{\times_T^{(l,m)}, \text{next}_X(c, t')\}$, where $t = \times_T(P)$ and $t' = \times_T(P')$. Since the first part of the minimum is the same, and $\times_T(Pc) < \times_T(P'c)$, we have two options

1. $\text{next}_X(c, t) \leq \times_T^{(l,m)} \leq \text{next}_X(c, t')$, where at most one inequality can be an equality; or
2. $\text{next}_X(c, t) < \text{next}_X(c, t') \leq \times_T^{(l,m)}$.

In any case, we have $\text{next}_X(c, t) < \text{next}_X(c, t')$. Since t and t' are positions in the same string, this relationship between the next occurrence of the same character immediately also implies $t < t'$, which concludes the proof. \blacktriangleleft

We are now ready to prove our main result:

▶ Theorem 11. *For any two nodes $u \neq u'$ of MDAG, let $ID(u) = \langle l, m, t, b \rangle$ and $ID(u') = \langle l, m, t', b' \rangle$. Then swing pairs for the same l, m do not dominate each other; namely, if $t > t'$, then $b \leq b'$.*

Proof. Consider any $v \in N^-(u)$, and $v' \in N^-(u')$. Let $ID(v) = \langle x, y, t_v, b_v \rangle$ and $ID(v') = \langle x', y', t_{v'}, b_{v'} \rangle$. Let us first assume that $x \neq x'$ or $y \neq y'$, i.e. they are not the same pair of positions. If we look at positions (x, y) and (x', y') we must have either $x \leq x'$ and $y > y'$, or $x > x'$ and $y \leq y'$. Indeed, assume by contradiction that $x \leq x'$ and $y \leq y'$. Since both of these nodes have a valid extension corresponding to positions (l, m) , we must also have $x \leq x' < l < t_v, t_{v'}$ and $y \leq y' < m < b_v, b_{v'}$. Then, (l, m) would not be a valid extension for v : the corresponding prefix is not maximal until the positions given by the swings, since we have an insertion corresponding to the character occurring at positions (x', y') .

We now show that $t > t'$ implies $y > y'$. By Remark 8, t is the smallest value such that a match occurs between $X[l+1, t]$ and $Y[y, m-1]$. That is, there is no $\tau < t$ such that $X[l+1, \tau]$ and $Y[y, m-1]$ have a match. Let $t > t'$, and assume by contradiction that $y \leq y'$. Then, $Y[y', m-1] \subseteq Y[y, m-1]$. By definition of t' , there is a match between $X[l+1, t']$ and $Y[y', m-1] \subseteq Y[y, m-1]$. This is a contradiction on the minimality of t : there is a smaller $\tau = t' < t$ which yields a match. Now, since we cannot have both $y' \leq y$ and $x' \leq x$, we must have $x \leq x'$. By a symmetrical reasoning, we show that the bottom swings must satisfy $b' \leq b$. Indeed, recall that b is the minimum value for which a match occurs between $X[x, l-1]$ and $Y[m+1, b]$, and assume by contradiction that $b > b'$. Since we have $X[x', l-1] \subseteq X[x, l-1]$, we have a match between $X[x, l-1]$ and $Y[m+1, b']$ for a smaller value $b' < b$: contradiction.

Let us now consider the case where the in-neighbors v and v' have the same pair of positions in their IDs: $x = x'$ and $y = y'$. Let us inductively consider $w_{i+1} \in N^-(w_i)$ and $w'_{i+1} \in N^-(w'_i)$, where $w_0 = v$ and $w'_0 = v'$. We stop at the first j such that $ID(w_j) = \langle x_j, y_j, t_j, b_j \rangle$ and $ID(w'_j) = \langle x'_j, y'_j, t'_j, b'_j \rangle$ with $x_j \neq x'_j$ or $y_j \neq y'_j$. Such pair satisfies the conditions of the first part of the proof, since $x_{j+1} = x'_{j+1}$ and $y_{j+1} = y'_{j+1}$ by hypothesis. Nodes w_j and w'_j are obtained by going backwards in the DAG for j steps, starting respectively from nodes u and u' . Since j is the first index such that the corresponding positions for extensions differ, we have $\ell(w_{k+1}, w_k) = \ell(w'_{k+1}, w'_k)$ for all $k = 1, \dots, j-1$. By iterating Corollary 10, we thus have that $t > t'$ implies $t_j > t'_j$. By the first part of the proof, we therefore have $b_j \leq b'_j$. By Lemma 9, this propagates to the end of the path in the MCS DAG, to also yield $b \leq b'$. \blacktriangleleft

From Theorem 11, we can derive the following result, which proves that the number of swings for a fixed pair of positions (l, m) is linear:

► **Corollary 12.** *For any given choice of l, m , there are just $O(n)$ nodes of $\text{MDAG}(X, Y)$ having the form $ID() = \langle l, m, \cdot, \cdot \rangle$.*

Proof. Let us fix l, m , and consider the swings' set $S_{l,m} \subseteq \{0, \dots, n-1\} \times \{0, \dots, n-1\}$, where $(a, b) \in S_{l,m}$ if and only there exists u such that $ID(u) = \langle l, m, a, b \rangle$. By Theorem 11, if two pairs (a, b) and (c, d) belong to $S_{l,m}$, then it cannot be $a \leq c$ and $c \leq d$ (or vice versa). So one pair cannot dominate the other.

We observe that the size of $|S_{l,m}|$ is the size of the classical Pareto frontier: for an arbitrary set of points in the $\{0, \dots, n-1\} \times \{0, \dots, n-1\}$ grid, the number of points in a Pareto frontier is less than $2n$. The observation is folklore: each point in the frontier either increases the x -coordinate or decreases the y -coordinate (possibly both). Hence, there cannot be more points on the frontiers as the sum of the n possible x -coordinates plus the n possible y -coordinates. Hence, $|S_{l,m}| = O(n)$. ◀

Therefore, the number of nodes of the MDAG is cubic, as we have $O(n^2)$ choices for l, m , and every such choice gives at most a linear amount of swings (t, b) . Furthermore, it is immediate by construction of MDAG that the out-degree of every node is at most σ (the characters that are valid extensions the given prefix). This gives $O(n^3\sigma)$ nodes and edges in MDAG. When obtaining the compact MDAG, the number of nodes and edges do not increase, and a suitable representation of the string labels gives the space occupancy of $O(n^3\sigma)$ memory words stated in Theorem 13.

4 Efficient Operations on MDAG

We describe here how to support some operations on compact MDAG(X, Y), which has source s , target t , and no unary nodes. We assume that each node u stores the number $p(u)$ of ut -paths. As compact MDAG(X, Y) is a DAG of $O(n^3\sigma)$ edges, we can compute $p(u)$ for each node u in total $O(n^3\sigma)$ time by running a DFS, as $p(u)$ is the sum of the $p(v)$'s for the out-neighbors v 's of u .

4.1 CAT Enumeration of MCSs

The strings in $\text{MCS}(X, Y)$ can be listed in lexicographic order by enumerating the (labeled) st -paths in compact MDAG, which can be done in Constant Amortized Time with a simple DFS algorithm where the out-neighbors of each node are visited in increasing order of the labels of their outgoing edges. Although folklore, for completeness we sketch the DFS algorithm here.

Consider a node u , where initially $u = s$, and denote the set of all ut -paths in $G = \text{compact MDAG}(X, Y)$ by $\mathcal{P}_{u,t}(G)$. The central idea is that any ut -path starts with u , followed by an element of $\mathcal{P}_{v,t}(G \setminus u)$, i.e., a path in $G \setminus u$ (i.e. u and its incident edges removed) from an out-neighbor v of u to t . Since G is a DAG, we can go a step further: a path from u cannot reach u again, therefore it is *not* even necessary to remove u (and its incident edges) from the DAG. We can thus represent $\mathcal{P}_{u,t}(G)$ as the following disjoint union: $\mathcal{P}_{u,t}(G) = \bigcup_{v \in N(u)} \{\ell(u, v)P \mid P \in \mathcal{P}_{v,t}(G)\}$, where $\ell(u, v)$ denotes the character(s) labeling edge u, v . From this, it is immediate that enumeration can be performed by keeping a current path prefix which we expand by traversing the DAG, backtracking every time computation terminates for all children of a node.

This recursive procedure runs in constant amortized time. Consider its calls and observe that they form a recursion tree with the following properties:

- Each leaf in the recursion tree is a distinct st -path (and all st -paths are in these leaves without any duplication).
- Each internal node in the recursion tree always generates at least two children, due to the path compression in G : hence the number of internal nodes is not greater than the number of leaves.
- In each node we spend just constant time per recursive call. In each leaf we spend constant time, not accounting for explicitly printing its whole st -path.

Letting N be the number of st -paths, there are N leaves and at most N internal nodes in the recursion tree. Each call takes constant time to generate its node in the recursion tree, which means $O(1)$ time per node/leaf. This gives a total of $O(N)$ time for listing all the st -paths when starting from G , recalling that G has no unary node (except possibly t). This provides a CAT enumeration as we can charge $O(1)$ time per solution to cover this cost.

4.2 Searching, Selecting, and Ranking

We observe that each node u has at most σ out-neighbors and the edges towards them are distinct (each must start with a different character of Σ): this constitutes a “lexicographical partition” of the paths from u to t since, after a common prefix, a path starting with a larger character will lead to a lexicographically larger string.

Searching a string P as a prefix of strings from $MCS(X, Y)$ traverses compact MDAG starting from s and matching the characters in P along the labels for the edges in the path. Either the search fails before reaching the end of P , or it succeeds and leads to a node u (or to an arc with endpoint u). At this point, we can run the CAT enumeration (Section 4.1) starting from u to list all the ut -paths and so all the extensions of P to strings in $MCS(X, Y)$.

Selecting the i th string in lexicographic order from $MCS(X, Y)$ is similar but uses the $O(n)$ -bit information $p(v)$ in each traversed node v , that is, $p(v)$ is the number of vt -paths and requires $O(n)$ bits as the number of MCS can be $O(2^n)$. The select operation scans the outgoing edges in order of their labels, and checks the corresponding $p(v)$: whenever the edges scanned have the largest partial sum i , it follows the current edge e and subtracts from i the partial sum of the edges examined before e . It stops at node t . Let S be the string thus found by concatenating the labels on the traced path from s to t . The number of traversed nodes is at most $|S| + 1$, and in each node we will consider up to σ edges, for a total cost of $O(|S| \log \sigma)$, as this can be further optimized by storing, for each edge, the sum of the annotations of the previous edges, and using binary search on these $O(\sigma)$ sums. Ranking S is like searching S and using the partial sums as mentioned above to get a total sum of i when t is reached.

► **Theorem 13.** *Given two strings X and Y of length n on an alphabet of size σ , compact MDAG(X, Y) stores $MCS(X, Y)$ in space $O(n^3\sigma)$ and can be built in $O(n^3\sigma \log n)$ time. It supports the following operations:*

- List all strings in $MCS(X, Y)$ in $O(|MCS(X, Y)|)$ time, i.e., Constant Amortized Time.
- For a given string P , report just the strings from $MCS(X, Y)$ that have prefix P , in $O(|P| \log \sigma + occ)$ time, where occ is the number of reported strings.

The $MCS(X, Y)$ can be augmented with $O(n)$ bits per node, so that two further operations are supported, where each operation handles $O(n)$ -bit integers:

- For any integer $1 \leq i \leq |MCS(X, Y)|$, select the i th string S in lexicographic order from $MCS(X, Y)$, in $O(|S| \log \sigma)$ operations.
- For any string $S \in MCS(X, Y)$, return its rank i among the strings in lexicographic order from $MCS(X, Y)$, in $O(|S| \log \sigma)$ operations.

It is worth noting that (compact) MDAG can be stored in a succinct way, for example, using some recent methods introduced for automata [5].

5 Conclusions

In this paper we considered the problem of storing and searching the Maximal Common Subsequences of two input strings, as they may reveal further common structures in sequence analysis with respect to the LCSs. Our main contribution is that of reducing time and space from exponential bounds, using the current state of the art, to polynomial bounds, using our new compact DAG, which efficiently supports CAT (constant amortized time) enumeration, counting, and random access to the i -th element (i.e., rank and select operations). As future work, we plan to improve the time and space bounds for rank and select in Theorem 13 by storing the number of paths in each node in a more refined way, so as the space does not increase asymptotically, and the time bounds reduce from $O(|S| \log \sigma)$ operations on $O(n)$ -bit integers to $O(|S| \log \sigma + n)$ time, borrowing ideas from [2].

References

- 1 A. Abboud, A. Backurs, and V. V. Williams. Tight hardness results for LCS and other sequence similarity measures. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 59–78, October 2015. doi:10.1109/FOCS.2015.14.
- 2 Amihood Amir, Gianni Franceschini, Roberto Grossi, Tsvi Kopelowitz, Moshe Lewenstein, and Noa Lewenstein. Managing unbounded-length keys in comparison-driven data structures with applications to online indexing. *SIAM J. Comput.*, 43(4):1396–1416, 2014. doi:10.1137/110836377.
- 3 Ricardo A Baeza-Yates. Searching subsequences. *Theoretical Computer Science*, 78(2):363–376, 1991.
- 4 L. Bergroth, H. Hakonen, and T. Raita. A survey of longest common subsequence algorithms. In *Proceedings Seventh International Symposium on String Processing and Information Retrieval. SPIRE 2000*, pages 39–48, September 2000. doi:10.1109/SPIRE.2000.878178.
- 5 Sankardeep Chakraborty, Roberto Grossi, Kunihiko Sadakane, and Srinivasa Rao Satti. Succinct representation for (non)deterministic finite automata. *J. Comput. Syst. Sci.*, 131:1–12, 2023. doi:10.1016/j.jcss.2022.07.002.
- 6 Alessio Conte, Roberto Grossi, Giulia Punzi, and Takeaki Uno. Enumeration of maximal common subsequences between two strings. *Algorithmica*, pages 1–27, 2022.
- 7 Maxime Crochemore, Bořivoj Melichar, and Zdeněk Troníček. Directed acyclic subsequence graph – Overview. *Journal of Discrete Algorithms*, 1(3-4):255–280, 2003.
- 8 Maxime Crochemore and Zdeněk Troníček. Directed acyclic subsequence graph for multiple texts. *Rapport IGM*, pages 99–13, 1999.
- 9 C. B. Fraser, R. W. Irving, and M. Middendorf. Maximal common subsequences and minimal common supersequences. *Information and Computation*, 124(2):145–153, 1996. doi:10.1006/inco.1996.0011.
- 10 Miyuji Hirota and Yoshifumi Sakai. Efficient algorithms for enumerating maximal common subsequences of two strings. *CoRR*, abs/2307.10552, 2023. doi:10.48550/arXiv.2307.10552.
- 11 Miyuji Hirota and Yoshifumi Sakai. A fast algorithm for finding a maximal common subsequence of multiple strings. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, page 2022DML0002, 2023.
- 12 D. S. Hirschberg. Algorithms for the longest common subsequence problem. *J. ACM*, 24(4):664–675, October 1977. doi:10.1145/322033.322044.
- 13 W. J. Hsu and M. W. Du. Computing a longest common subsequence for a set of strings. *BIT Numerical Mathematics*, 24(1):45–59, 1984.

- 14 Elsa Loekito, James Bailey, and Jian Pei. A binary decision diagram based approach for mining frequent subsequences. *Knowl. Inf. Syst.*, 24(2):235–268, 2010. doi:10.1007/s10115-009-0252-9.
- 15 David Maier. The complexity of some problems on subsequences and supersequences. *Journal of the ACM (JACM)*, 25(2):322–336, 1978.
- 16 W. J. Masek and M. S. Paterson. A faster algorithm computing string edit distances. *Journal of Computer and System Sciences*, 20(1):18–31, 1980. doi:10.1016/0022-0000(80)90002-1.
- 17 Bořivoj Melichar and Tomáš Polcar. The longest common subsequence problem a finite automata approach. In *Implementation and Application of Automata: 8th International Conference, CIAA 2003 Santa Barbara, CA, USA, July 16–18, 2003 Proceedings*, pages 294–296. Springer, 2003.
- 18 Shin-ichi Minato. Zero-suppressed bdds for set manipulation in combinatorial problems. In *Proceedings of the 30th International Design Automation Conference, DAC '93*, pages 272–277, New York, NY, USA, 1993. Association for Computing Machinery. doi:10.1145/157485.164890.
- 19 R. Raman, V. Raman, and S. R. Satti. Succinct indexable dictionaries with applications to encoding k -ary trees, prefix sums and multisets. *ACM Trans. Algorithms*, 3(4):43–es, November 2007. doi:10.1145/1290672.1290680.
- 20 Frank Ruskey. Combinatorial generation. *Preliminary working draft. University of Victoria, Victoria, BC, Canada*, 11:20, 2003.
- 21 Yoshifumi Sakai. Maximal common subsequence algorithms. In Gonzalo Navarro, David Sankoff, and Binhai Zhu, editors, *Annual Symposium on Combinatorial Pattern Matching (CPM 2018)*, volume 105 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 1:1–1:10, Dagstuhl, Germany, 2018. Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.CPM.2018.1.
- 22 Yoshifumi Sakai. Maximal common subsequence algorithms. *Theoretical Computer Science*, 793:132–139, 2019. doi:10.1016/j.tcs.2019.06.020.
- 23 Etsuji Tomita, Akira Tanaka, and Haruhisa Takahashi. The worst-case time complexity for generating all maximal cliques and computational experiments. *Theoretical Computer Science*, 363(1):28–42, 2006. Computing and Combinatorics.
- 24 Zdeněk Troníček. Common subsequence automaton. In *International Conference on Implementation and Application of Automata*, pages 270–275. Springer, 2002.
- 25 R. A. Wagner and M. J. Fischer. The string-to-string correction problem. *J. ACM*, 21(1):168–173, January 1974. doi:10.1145/321796.321811.