



Linear-Time Computation of Generalized Minimal Absent Words for Multiple Strings

Kouta Okabe¹, Takuya Mieno², Yuto Nakashima³,
Shunsuke Inenaga³, and Hideo Bannai⁴

¹ Department of Information Science and Technology, Kyushu University, Fukuoka, Japan

² Department of Computer and Network Engineering, University of Electro-Communications, Tokyo, Japan
tmieno@uec.ac.jp

³ Department of Informatics, Kyushu University, Fukuoka, Japan
{nakashima.yuto.003, inenaga.shunsuke.380}@m.kyushu-u.ac.jp

⁴ M&D Data Science Center, Tokyo Medical and Dental University, Tokyo, Japan
hdbn.dsc@tmd.ac.jp

Abstract. A string w is called a *minimal absent word (MAW)* for a string S if w does not occur as a substring in S and all proper substrings of w occur in S . MAWs are well-studied combinatorial string objects that have potential applications in areas including bioinformatics, musicology, and data compression. In this paper, we generalize the notion of MAWs to a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of multiple strings. We first describe our solution to the case of $k = 2$ strings, and show how to compute the set \mathbf{M} of MAWs in optimal $O(n + |\mathbf{M}|)$ time and with $O(n)$ working space, where n denotes the total length of the strings in \mathcal{S} . We then move on to the general case of $k > 2$ strings, and show how to compute the set \mathbf{M} of MAWs in $O(n \lceil k / \log n \rceil + |\mathbf{M}|)$ time and with $O(n(k + \log n))$ bits of working space, in the word RAM model with machine word size $\omega = \log n$. The latter algorithm runs in optimal $O(n + |\mathbf{M}|)$ time for $k = O(\log n)$.

1 Introduction

A non-empty string w is said to be an *absent word* (a.k.a. *a forbidden word*) for a string S if w is *not* a substring of S . An absent word w for S is said to be a *minimal absent word (MAW)* for S if all proper substrings of w occur in S . For instance, for string $S = \text{bbaccbbaa}$ over an alphabet $\Sigma = \{\text{a, b, c, d}\}$, the set $\text{MAW}(S)$ of all MAWs for S is $\{\text{aaa, bbb, cccc, d, ab, ca, bc, aac, acb, cbb, accb, cbac, bbaa}\}$. MAWs are combinatorial string objects, and their interesting mathematical properties have extensively been studied in the literature (see [1, 7, 16, 17, 19, 23] and references therein). MAWs also enjoy several applications including phylogeny [11], data compression [3, 15, 18], musical information retrieval [14], and bioinformatics [2, 12, 22, 24].

It is known that the number $|\text{MAW}(S)|$ of MAWs for a string S of length n over an alphabet of size σ is $O(\sigma n)$ and that this bound is tight [17]. Crochemore et al. [17] gave an algorithm that computes $\text{MAW}(S)$ in $O(\sigma n)$ time with $O(n)$ working space. Fujishige et al. [20] showed an improved algorithm for computing $\text{MAW}(S)$ in optimal $O(n + |\text{MAW}(S)|)$ time with $O(n)$ working space, for an input string S of length n over an integer alphabet of polynomial size in n . Both of the two aforementioned algorithms utilize an $O(n)$ -size string data structure called the (*directed acyclic word graph*) DAWG [9], which recognizes the set of substrings of S , and can be built in $O(n \log \sigma)$ time for general ordered alphabets [9], and in $O(n)$ time for integer alphabets of polynomial size in n [20]. There also exist other efficient algorithms for computing MAWs with other string data structures such as suffix arrays and Burrows-Wheeler transforms [4, 8].

The aim of this paper is to extend the notion of MAWs to a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of multiple k strings. We are aware of a few related attempts in earlier work: Chairungsee and Crochemore [11] introduced a string similarity measure based on the symmetric difference $\text{MAW}(S_1) \Delta \text{MAW}(S_2)$ of the sets of MAWs for two strings S_1 and S_2 to compare. They introduced a length threshold $\ell \geq 1$, and described an approach for computing $(\text{MAW}(S_1) \Delta \text{MAW}(S_2)) \cap \Sigma^\ell$ with the two following steps: First, the tries of size $O(n\ell)$ each representing the substrings of S_1 and S_2 of length up to ℓ are built, where $n = |S_1| + |S_2|$. Then, two tries each representing $\text{MAW}(S_1) \cap \Sigma^\ell$ and $\text{MAW}(S_2) \cap \Sigma^\ell$ are built, which require $O(n\sigma)$ space. Finally, the length-bounded symmetric difference $(\text{MAW}(S_1) \Delta \text{MAW}(S_2)) \cap \Sigma^\ell$ is computed from $\text{MAW}(S_1) \cap \Sigma^\ell$ and $\text{MAW}(S_2) \cap \Sigma^\ell$, but the authors did not explicitly describe how this computation is done in their method. Overall, their algorithm requires $\Omega(n(\ell + \sigma))$ time and space [11]¹. Charalampapaulose et al. [12] tackled the same problem of computing the symmetric difference $\text{MAW}(S_1) \Delta \text{MAW}(S_2)$ (without length threshold ℓ), and proposed a solution that requires $O(\sigma n)$ time and space. Their method firstly computes $\text{MAW}(S_1)$ and $\text{MAW}(S_2)$ separately, and then removes the elements that are in $\text{MAW}(S_1) \cap \text{MAW}(S_2)$. Charalampopoulos, Crochemore, and Pissis [13] presented how to count the number $|\text{MAW}(S_1) \Delta \text{MAW}(S_2)|$ of elements in the symmetric difference $\text{MAW}(S_1) \Delta \text{MAW}(S_2)$ in $O(n)$ time in the case of integer alphabets of polynomial size in n , by avoiding to list the elements explicitly.

Let $\mathcal{S} = \{S_1, \dots, S_k\}$ be the input set of k strings, and $\mathbf{B} \in \{0, 1\}^k$ be a given bit vector of length k . Our problem is to list (generalized) MAWs w for \mathcal{S} and \mathbf{B} such that $w \in \text{MAW}(S_i)$ for every $\mathbf{B}[i] = 1$, and $w \notin \text{MAW}(S_i)$ for every $\mathbf{B}[i] = 0$. For $k = 2$, the aforementioned problem of computing $\text{MAW}(S_1) \Delta \text{MAW}(S_2)$ is equivalent to solving our problem for $\mathbf{B} = 01$ and $\mathbf{B} = 10$. In Sect. 4 and Sect. 5, we deal with the case with $k = 2$, and present an algorithm running in $O(n + |\mathbf{M}_{\mathbf{B}}|)$ time with $O(n)$ working space, where $\mathbf{M}_{\mathbf{B}}$ denotes the set of (generalized) MAWs to output for a given bit vector \mathbf{B} (Theorem 2). This immediately gives us an algorithm for listing the elements of the symmetric

¹ The claimed time bound for computing the trie is $O(n\sigma)$ (Theorem 1 of [11]). It seems that the authors regarded the length threshold ℓ as a constant.

difference $\text{MAW}(S_1) \Delta \text{MAW}(S_2)$ in optimal $O(n + |\text{MAW}(S_1) \Delta \text{MAW}(S_2)|)$ time (Corollary 1). In Sect. 6, we deal with the general case of $k > 2$, and extend our solution for $k = 2$ to the general case. Let n be the total length of the input k strings in \mathcal{S} . Our solution for general $k > 2$ works in $O(n \lceil k / \log n \rceil + |\mathbf{M}_{\mathbf{B}}|)$ time with $O(n(k + \log n))$ bits of working space on the word RAM model with machine word size $\omega = \log n$. Thus, for $k = O(\log n)$, our algorithm runs in optimal $O(n + |\mathbf{M}_{\mathbf{B}}|)$ time. All the bounds claimed in this paper are valid for linearly sortable alphabets, including integer alphabets of polynomial size in n .

As in the previous work [17, 20, 21], our key data structure is the DAWG for the input set \mathcal{S} of strings. The best-known algorithm for constructing the DAWG for a set of strings of total length n takes $O(n \log \sigma)$ time [10], thus it can require $O(n \log n)$ time for large alphabets. We describe how the DAWG for a given set \mathcal{S} of strings over an integer alphabet of polynomial size in n can be obtained in optimal $O(n)$ time (Theorem 1), which may be of independent interest.

2 Preliminaries

Strings. Let Σ be an ordered alphabet. An element of Σ is called a character. For characters $a, b \in \Sigma$, we write $a < b$ (or equivalently $b > a$) if a is lexicographically smaller than b . An element of Σ^* is called a string. The length of a string S is denoted by $|S|$. The empty string ε is the string of length 0. If $S = xyz$, then x , y , and z are called a *prefix*, *substring*, and *suffix* of S , respectively. They are called a *proper prefix*, *proper substring*, and *proper suffix* of S if $x \neq S$, $y \neq S$, and $z \neq S$, respectively. Let $\text{Substr}(S)$ denote the set of substrings of string S . For any $1 \leq i \leq |S|$, the i -th character of S is denoted by $S[i]$. For any $1 \leq i \leq j \leq |S|$, $S[i..j]$ denotes the substring of S starting at i and ending at j . For convenience, let $S[i..j] = \varepsilon$ for $0 \leq j < i \leq |S| + 1$. We say that a string w occurs in a string S iff w is a substring of S . Note that by definition the empty string ε is a substring of any string S and hence ε always occurs in S .

For a set \mathcal{S} of strings, let $\|\mathcal{S}\|$ denote the total length of the strings in \mathcal{S} , that is, $\|\mathcal{S}\| = \sum_{S \in \mathcal{S}} |S|$. Let $\text{Substr}(\mathcal{S})$ denote the set of substrings of the strings in \mathcal{S} , that is, $\text{Substr}(\mathcal{S}) = (\bigcup_{S \in \mathcal{S}} \{S[i..j] \mid 1 \leq i \leq j \leq |S|\}) \cup \{\varepsilon\}$.

Minimal Absent Words (MAWs). A string w is called an *absent word* for a string S if w does not occur in S . Let $\text{AW}(S) = \Sigma^* \setminus \text{Substr}(S)$ denote the set of absent words for a string S . An absent word $w \in \text{AW}(S)$ for string S is called a *minimal absent word* or *MAW* for S if any proper substring of w occurs in S . We denote by $\text{MAW}(S)$ the set of all MAWs for S . Let $\text{nonMAW}(S) = \text{AW}(S) \setminus \text{MAW}(S)$ be the set of absent words for S which are not MAWs. Note that, for strings w and S , it holds that $w \notin \text{MAW}(S)$ iff $w \in \text{Substr}(S) \cup \text{nonMAW}(S)$.

We extend the aforementioned notion of MAWs to a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of k strings for $k \geq 1$, as follows: Let \mathbf{B} be a bit-vector of length k , and let $\mathcal{S}_{\mathbf{B}}$ be a subset of \mathcal{S} such that $\mathcal{S}_{\mathbf{B}} = \{S_i \mid \mathbf{B}[i] = 1\}$. Let $\overline{\mathcal{S}_{\mathbf{B}}} = \{S_i \mid \mathbf{B}[i] = 0\} = \mathcal{S} \setminus \mathcal{S}_{\mathbf{B}}$. A string w is said to be a MAW for $\mathcal{S}_{\mathbf{B}}$ if (1) $w \in \bigcap_{S_i \in \mathcal{S}_{\mathbf{B}}} \text{MAW}(S_i)$ and (2) $w \notin \bigcup_{S_i \in \overline{\mathcal{S}_{\mathbf{B}}}} \text{MAW}(S_i)$. Condition (1) implies that w is a MAW for any string in $\mathcal{S}_{\mathbf{B}}$.

Condition (2) implies that w is *not* a MAW for any string in $\overline{\mathcal{S}_{\mathbf{B}}}$, which is equivalent to say that $w \in \bigcap_{S_i \in \overline{\mathcal{S}_{\mathbf{B}}}} (\text{Substr}(S_i) \cup \text{nonMAW}(S_i))$. Let $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ be the set of all MAWs for $\mathcal{S}_{\mathbf{B}}$. Here is some example: For string set $\mathcal{S} = \{\text{abaab}, \text{aacbba}\}$ over the alphabet $\Sigma = \{\mathbf{a}, \mathbf{b}, \mathbf{c}, \mathbf{d}\}$, $\text{MAW}(\mathcal{S}_{10}) = \{\text{aaba}, \text{bab}, \text{bb}, \mathbf{c}\}$, $\text{MAW}(\mathcal{S}_{01}) = \{\text{ab}, \text{baa}, \text{bac}, \text{bbb}, \text{bc}, \text{ca}, \text{cba}, \text{cc}\}$, and $\text{MAW}(\mathcal{S}_{11}) = \{\text{aaa}, \mathbf{d}\}$.

The problem we consider in this paper is the following:

Problem 1 (MAWs for multiple input strings). Given a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of k strings over an alphabet Σ and a bit vector \mathbf{B} of length k , compute $\text{MAW}(\mathcal{S}_{\mathbf{B}})$.

3 The DAWG Data Structure

We use the *directed acyclic word graph (DAWG)* [9] data structure for a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of k strings, which is a DFA of size $O(\|\mathcal{S}\|)$ that recognizes all suffixes of the strings in \mathcal{S} .

To give a formal definition of $\text{DAWG}(\mathcal{S})$, let $\text{End_Pos}_{\mathcal{S}}(w)$ denote the set of ending positions of all occurrences of a string w in the strings of \mathcal{S} , that is,

$$\text{End_Pos}_{\mathcal{S}}(w) = \{(i, j) \mid S_i[j - |w| + 1..j] = w, 1 \leq i \leq k, 1 \leq j \leq |S_i|\}.$$

We consider an equivalence relation $\equiv_{\mathcal{S}}$ of strings over Σ w.r.t. \mathcal{S} such that, for any two strings w and u , $w \equiv_{\mathcal{S}} u$ iff $\text{End_Pos}_{\mathcal{S}}(w) = \text{End_Pos}_{\mathcal{S}}(u)$. For any string $x \in \Sigma^*$, let $[x]_{\mathcal{S}}$ denote the equivalence class for x w.r.t. $\equiv_{\mathcal{S}}$. All the non-substrings $x \notin \text{Substr}(\mathcal{S})$ form a unique equivalence class, called the *degenerate class*.

Definition 1. *The DAWG of a set \mathcal{S} of strings, denoted $\text{DAWG}(\mathcal{S})$, is an edge-labeled DAG (V, E) such that*

$$\begin{aligned} V &= \{[x]_{\mathcal{S}} \mid x \in \text{Substr}(\mathcal{S})\}, \\ E &= \{([x]_{\mathcal{S}}, b, [xb]_{\mathcal{S}}) \mid x, xb \in \text{Substr}(\mathcal{S}), b \in \Sigma\}. \end{aligned}$$

We also define the set L of suffix links of $\text{DAWG}(\mathcal{S})$ by

$$L = \{([ax]_{\mathcal{S}}, a, [x]_{\mathcal{S}}) \mid x, ax \in \text{Substr}(\mathcal{S}), a \in \Sigma, [ax]_{\mathcal{S}} \neq [x]_{\mathcal{S}}\}.$$

Namely, two substrings x and y in $\text{Substr}(\mathcal{S})$ are represented by the same node of $\text{DAWG}(\mathcal{S})$ iff the ending positions of x and y in the strings of \mathcal{S} are equal. Note that $\text{DAWG}(\mathcal{S})$ does not contain the node for the degenerate class nor its in-coming edges. This is important for $\text{DAWG}(\mathcal{S})$ to have a total linear number of edges [9], and for our linear-time algorithm for listing all the MAWs for a given query.

For convenience, assume that each string S_i in $\mathcal{S} = \{S_1, \dots, S_k\}$ terminates with a unique end-marker $\#_i$ which does not occur elsewhere, where $\#_i \neq \#_j$ for $i \neq j$. Then $\text{DAWG}(\mathcal{S})$ has exactly k sink nodes, each of which recognizes all the non-empty suffixes of S_i . For each $1 \leq i \leq k$, the sink that recognizes the suffixes of S_i is labeled by i .

The DAWG for a single string T is the DAWG for a singleton $\{T\}$ and is denoted by $\text{DAWG}(T)$.

The state-of-the-art algorithm that builds $\text{DAWG}(\mathcal{S})$ is Blumer et al.’s online algorithm [9] which runs in $O(n \log \sigma)$ time with $O(n)$ space, where $n = \|\mathcal{S}\|$ is the total length of the strings in \mathcal{S} and σ is the alphabet size. Below we describe a faster construction of $\text{DAWG}(\mathcal{S})$ in the case of integer alphabets:

Theorem 1 (Linear-time DAWG construction for a set of strings).

For a given set $\mathcal{S} = \{S_1, \dots, S_k\}$ of k strings of total length n over an integer alphabet Σ of polynomial size in n , one can build the edge-sorted $\text{DAWG}(\mathcal{S})$ in $O(n)$ time and space.

Proof. We first create a concatenated string $T = S_1 \cdots S_k$ of total length n from the strings in \mathcal{S} . We build $\text{DAWG}(T)$ for the single string T in $O(n)$ time and space, using the algorithm of Fujishige et al. [20, 21], where the out-going edges of every node are lexicographically sorted. Our goal is to convert $G_T = \text{DAWG}(T)$ to $G_{\mathcal{S}} = \text{DAWG}(\mathcal{S})$. For a set P of integer pairs and a pair (a, b) of integers, let $P \oplus (a, b) = \{(p + a, q + b) \mid (p, q) \in P\}$. Our key observation is that, for any substrings $w \in \text{Substr}(\mathcal{S})$ that do not contain separators $\#_i$ except for their last positions, it holds that

$$\begin{aligned} &\text{End_Pos}_{\mathcal{S}}(w) \\ &= \text{End_Pos}_{S_1}(w) \cup \left(\bigcup_{2 \leq i \leq k} \text{End_Pos}_{S_i}(w) \oplus (i - 1, |S_1 \cdots S_{i-1}|) \right). \end{aligned} \tag{1}$$

Equation (1) implies that the substrings w of $T = S_1 \cdots S_k$ which are also substrings of \mathcal{S} are represented by essentially the same nodes in G_T and in $G_{\mathcal{S}}$, meaning that there is an injection from the nodes of $G_{\mathcal{S}}$ to the nodes of G_T .

What is left is how to remove the redundant nodes in G_T which represent the substrings y of T containing a separator $\#_i$ inside, which are thus not substrings of \mathcal{S} . Let us call the longest path of G_T that represents T as the *spine*. Since each $\#_i$ occurs exactly once in T , any substrings of T that contain $\#_i$ are represented by the spine of G_T . Thus, we can obtain $G_{\mathcal{S}}$ by removing the redundant nodes from the spine of G_T , but we ensure that for every i the suffixes of S_i ending with $\#_i$ are still represented in the graph. This can be achieved as follows: We process $i = k, \dots, 2$ in decreasing order. We first split the spine into two parts each spelling out $S_1 \cdots S_{k-1}$ and S_k . We remove the nodes in the S_k part which are not reachable from the source of the modified graph, together with their out-going edges and suffix links. This gives us $\text{DAWG}(\{S_1 \cdots S_{k-1}, S_k\})$. After processing $i = k$, we continue the same process for $i = k - 1$ with the remaining spine that spells out $S_1 \cdots S_{k-1}$. After processing $i = 2$, we obtain $G_{\mathcal{S}} = \text{DAWG}(\mathcal{S})$. See Fig. 1 for an example of our construction. It is trivial that all the redundant nodes can be removed in $O(n)$ time. \square

We remark that the order of concatenating the strings in \mathcal{S} does not affect the correctness nor the complexity of our algorithm.




Fig. 1. Illustration for our linear-time construction of $DAWG(S)$ for a set $S = \{abc\#_1, bbac\#_2, abca\#_3\}$ of strings. We first build $DAWG(T)$ for the concatenated string $T = abc\#_1bbac\#_2abca\#_3$. Then, we remove the redundant nodes in the spine of the DAWG for $i = 3$ and then for $i = 2$. This gives us $DAWG(S)$.

4 Algorithm Overview for $k = 2$

In what follows, we consider the case where our input set \mathcal{S} consists of two strings S_1 and S_2 which respectively terminate with special characters $\#_1$ and $\#_2$. We show how, given a bit vector $\mathbf{B} \in \{00, 01, 10, 11\}$ of length 2, we can compute $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ in $O(n + |\text{MAW}(\mathcal{S}_{\mathbf{B}})|)$ time and $O(n)$ working space, where $n = \|\mathcal{S}\|$.

We first build the edge-sorted $\text{DAWG}(\mathcal{S})$ for a given $\mathcal{S} = \{S_1, S_2\}$ in $O(n)$ time and space with Theorem 1. We label each node v of $\text{DAWG}(\mathcal{S})$ by $\#_i$ iff v represents a substring of S_i ($1 \leq i \leq 2$). Let $\text{label}(v) \in \{\#_1, \#_2, \#_1\#_2\}$ denote the label of node v . The labels of all nodes can be precomputed in $O(n)$ time.

Our algorithm is based on Fujishige et al.’s algorithm [20,21] for computing all the MAWs in the case of a single input string. As such, for each node x of $\text{DAWG}(\mathcal{S})$ we focus on the *shortest* string represented by x and denote it by au , where $a \in \Sigma$ and $u \in \Sigma^*$. We use the suffix link of the node x and its target node y whose *longest* member is u (namely, the first letter a of au is removed by following the suffix link from x to y). For ease of explanation, we identify the node x with the string au , and the node y with the string u .

Fujishige et al.’s algorithm compares the out-going edges of au and those of u one by one in the sorted order. Suppose au has an out-going edge labeled b . If u does not have an out-going edge labeled b , then their algorithm outputs aub as a MAW for the input string. Otherwise, it outputs nothing, and the cost is charged to the out-going edge of au labeled b . Each MAW aub in the output is encoded by a tuple (a, i, j) such that $w[i..j] = ub$, thus taking $O(1)$ space. This is how Fujishige et al.’s algorithm works in $O(n + |\text{MAW}(S)|)$ time and with $O(n)$ working space for a single string S .

However, in our case of multiple strings, depending on the label of nodes au , aub and ub , and depending on the value of the given bit vector \mathbf{B} , there may exist some edge comparisons that cannot be charged either to the output MAWs or to the out-going edges of node au . It is also possible that even if there is a node representing aub in $\text{DAWG}(\mathcal{S})$, still aub is a MAW for some string(s) in \mathcal{S} . To overcome these difficulties, we introduce *skip links* that permit us to avoid unwanted edge character comparisons.

5 Skip Links for $k = 2$

We use the same conventions for the nodes au , aub and u on $\text{DAWG}(\mathcal{S})$ as in the previous section, and also consider the node ub . We have three possible cases for the label of node au , where $\text{label}(au) = \#_1\#_2$, $\text{label}(au) = \#_1$, or $\text{label}(au) = \#_2$. In each of the three cases, there are some sub-cases for the labels of node aub and node ub . By inspection, we obtain all the possible cases that need to be considered, as shown in Fig. 2.

When $\mathbf{B} = 00$, then since $\text{MAW}(\mathcal{S}_{00}) = \Sigma^* \setminus (\text{MAW}(S_1) \cup \text{MAW}(S_2))$, there are no MAWs to output. In what follows, we describe our solutions to the cases with $\mathbf{B} \in \{10, 11\}$. We remark that the case with $\mathbf{B} = 01$ is symmetric to the case with $\mathbf{B} = 10$.

5.1 When $B = 10$

There are four cases in which we output aub as a MAW for $MAW(\mathcal{S}_{10})$ (see the table on the left of Fig. 2):

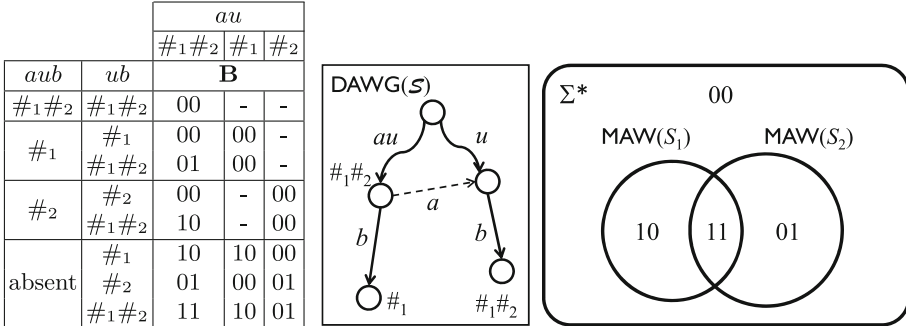


Fig. 2. Left: All possible cases of the labels of the nodes au , aub , and ub , and their corresponding bit vectors **B**. “absent” refers to the case where there is no out-going edge labeled b from node au . The cells with “-” refer to impossible combinations of node labels. Middle: Illustration for $DAWG(\mathcal{S})$ which shows the case where au is labeled $\#_1\#_2$, aub is labeled $\#_1$, and ub is labeled $\#_1\#_2$. In this case aub is a MAW in $MAW(\mathcal{S}_B)$ with $B = 01$ (see the left table). Right: The regions corresponding to the bit vectors $B \in \{00, 01, 10, 11\}$.

- (1) $label(au) = \#_1\#_2$, $label(aub) = \#_2$, and $label(ub) = \#_1\#_2$;
- (2) $label(au) = \#_1\#_2$, $aub \in AW(\mathcal{S})$, and $label(ub) = \#_1$;
- (3) $label(au) = \#_1$, $aub \in AW(\mathcal{S})$, and $label(ub) = \#_1$;
- (4) $label(au) = \#_1$, $aub \in AW(\mathcal{S})$, and $label(ub) = \#_1\#_2$.

When $label(au) = \#_1\#_2$. We create skip links that simultaneously manage Cases (1) and (2), both having $label(au) = \#_1\#_2$. We create a selected list $schar(u)$ of out-going edge labels of node u such that $schar(u) = \{b \mid label(ub) = \#_1\}$, where the elements are lexicographically sorted. Let $char(au)$ be the sorted list of all out-going edge labels of node au . For any list L of characters and any character $c \in \Sigma$, let $succ(c, L)$ denote the lexicographical successor of c in L . Our algorithm for $B = 10$ and $label(au) = \#_1\#_2$ is described in Algorithm 1.

When $label(au) = \#_1$. We create skip links that simultaneously manage Cases (3) and (4), both having $label(au) = \#_1$. We create another selected list $schar'(u)$ of out-going edge labels of node u such that $schar'(u) = \{b \mid label(ub) \in \{\#_1, \#_1\#_2\}\}$, where the elements are lexicographically sorted. We use the same $char(au)$ in the previous case. Our algorithm for $B = 10$ and $label(au) = \#_1$ is described in Algorithm 2.

Lemma 1 (Linear-time MAW computation for $B = 10$). *Given $B = 10$, one can compute $MAW(\mathcal{S}_{10})$ in $O(n + |MAW(\mathcal{S}_{10})|)$ time and $O(n)$ working space for integer alphabets of polynomial size in $n = \|\mathcal{S}\|$.*

Algorithm 1: Algorithm for $\mathbf{B} = 10$ and $\text{label}(au) = \#_1\#_2$

Input: A node au of $\text{DAWG}(\mathcal{S})$ such that $\text{label}(au) = \#_1\#_2$, $\mathbf{B} = 10$.

Output: A subset M of MAWs aub with $b \in \Sigma$.

```

1  $M \leftarrow \emptyset$ ;
2  $U \leftarrow \text{char}(au) \cup \{\$U\}$ ;           /*  $\$U$  is lex. largest in  $U$  */
3  $L \leftarrow \text{schar}(u) \cup \{\$L\}$ ;       /*  $\$L$  is lex. largest in  $L$  and  $\$L \prec \$u$  */
4  $\hat{b} \leftarrow U[1]$ ;  $b \leftarrow L[1]$ ;     /* start with lex. smallest characters */
5 while  $b \neq \$L$  do
6   if  $\hat{b} = b$  then
7     if  $\text{label}(aub) = \#_2$  and  $\text{label}(ub) = \#_1\#_2$  then
8        $M \leftarrow M \cup \{aub\}$ ;         /* output  $aub$  */
9        $\hat{b} \leftarrow \text{succ}(\hat{b}, U)$ ;       /* move to the next character in  $U$  */
10       $b \leftarrow \text{succ}(b, L)$ ;         /* move to the next character in  $L$  */
11   else if  $\hat{b} \succ b$  then
12      $M \leftarrow M \cup \{aub\}$ ;         /* output  $aub$  */
13      $b \leftarrow \text{succ}(b, L)$ ;         /* move to the next character in  $L$  */
14 return  $M$ ;

```

Proof. We run Algorithm 1 and Algorithm 2 for every node au of $\text{DAWG}(\mathcal{S})$.

In the preprocessing phase, we build the edge-sorted $\text{DAWG}(\mathcal{S})$ in $O(\|\mathcal{S}\|)$ time and space by Theorem 1. Since the out-going edges of every node are sorted, we can easily compute the sorted lists $\text{char}(au)$, $\text{schar}(u)$, $\text{schar}'(u)$, and $\text{schar}''(u)$ for all nodes in $O(n)$ total time.

Let us consider the complexity of the scanning phase of Algorithm 1. Each edge-label comparison that falls into “ $\hat{b} = b$ ” in line 6 of Algorithm 1 is associated either to the reported MAW aub if $\text{label}(aub) = \#_2$ and $\text{label}(ub) = \#_1\#_2$ (in line 7 and line 8), or to the out-going edge of node au labeled b otherwise. Each edge-label comparison that falls into “ $\hat{b} \succ b$ ” in line 11 is associated to the reported MAW aub in line 12. This ensures the desired time complexity for Algorithm 1. The complexity for Algorithm 2 is similar to show.

The correctness of Algorithm 1 and Algorithm 2 is immediate from the tables in Fig. 2 and 3. □

5.2 When $\mathbf{B} = 11$

There is a single case in which we output aub as a MAW for $\text{MAW}(\mathcal{S}_{11})$ (see Fig. 2): $\text{label}(au) = \#_1\#_2$, $aub \in \text{AW}(\mathcal{S})$, and $\text{label}(ub) = \#_1\#_2$.

Unwanted comparisons can occur here if $aub \in \text{AW}(\mathcal{S})$, and $\text{label}(ub) = \#_1$ or $\text{label}(ub) = \#_2$. To avoid such comparisons, we consider another carefully selected list $\text{schar}''(u)$ of out-going edge labels of node u such that $\text{schar}''(u) = \{b \mid \text{label}(ub) = \#_1\#_2\}$, where the elements are lexicographically sorted. We can use the same $\text{char}(au)$ in the previous subsection.

We can modify Algorithm 2 for $\mathbf{B} = 01$ with $\text{label}(au) = \#_1$ so that the modified algorithm computes MAWs for $\mathbf{B} = 11$, only by using $\text{schar}''(u)$ in place of $\text{schar}'(u)$. This leads us to the following lemma:

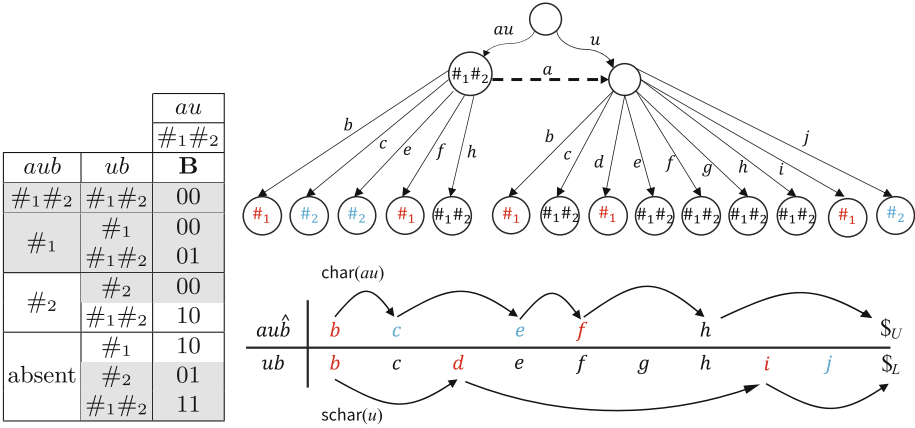


Fig. 3. Illustration for our algorithm for $\mathbf{B} = 10$ and $label(au) = \#_1\#_2$. The white cells in the table show the cases where we output elements of $MAW(\mathcal{S}_{10})$. We compare the labels of the selected out-going edges of node au and u which are connected by the skip links, in sorted order. In this diagram, aud and auj are output in line 12 and auc and $au e$ are output in line 12 of Algorithm 1 as elements of $MAW(\mathcal{S}_{10})$.

Lemma 2 (Linear-time MAW computation for $\mathbf{B} = 11$). *Given $\mathbf{B} = 11$, one can compute $MAW(\mathcal{S}_{11})$ in $O(n + |MAW(\mathcal{S}_{11})|)$ time and $O(n)$ working space for integer alphabets of polynomial size in $n = \|\mathcal{S}\|$.*

5.3 Our Main Result for $k = 2$

Finally we obtain the main result for a case of two strings with $k = 2$.

Theorem 2 (Linear-time MAW computation for a set of two strings). *Given a set $\mathcal{S} = \{S_1, S_2\}$ of two strings of total length n and a bit vector $\mathbf{B} \in \{01, 10, 11\}$, one can compute $MAW(\mathcal{S}_{\mathbf{B}})$ in $O(n + |MAW(\mathcal{S}_{\mathbf{B}})|)$ time and $O(n)$ working space for integer alphabets of polynomial size in n .*

The following corollary is immediate from Theorem 2.

Corollary 1. *Given a set $\mathcal{S} = \{S_1, S_2\}$ of two strings of total length n , one can compute $MAW(S_1) \cap MAW(S_2)$, $MAW(S_1) \cup MAW(S_2)$, and $MAW(S_1) \Delta MAW(S_2)$ in $O(n + |MAW(S_1) \cap MAW(S_2)|)$ time, $O(n + |MAW(S_1) \cup MAW(S_2)|)$ time, and $O(n + |MAW(S_1) \Delta MAW(S_2)|)$ time, respectively, using $O(n)$ working space, for integer alphabets of polynomial size in n .*

6 Algorithm for Arbitrary $k > 2$

In this section, we present our algorithm for computing $MAW(\mathcal{S}_{\mathbf{B}})$ in case where $\mathcal{S} = \{S_1, \dots, S_k\}$ contains $k > 2$ strings.

Algorithm 2: Algorithm for $\mathbf{B} = 10$ and $\text{label}(au) = \#_1$

Input: A node au of $\text{DAWG}(\mathcal{S})$ such that $\text{label}(au) = \#_1$, $\mathbf{B} = 10$.

Output: A subset M of MAWs aub with $b \in \Sigma$.

```

1  $M \leftarrow \emptyset$ ;
2  $U \leftarrow \text{char}(au) \cup \{\$U\}$ ;           /*  $\$U$  is lex. largest in  $U$  */
3  $L \leftarrow \text{schar}'(u) \cup \{\$L\}$ ;       /*  $\$L$  is lex. largest in  $L$  and  $\$L \prec \$u$  */
4  $\hat{b} \leftarrow U[1]$ ;  $b \leftarrow L[1]$ ;     /* start with lex. smallest characters */
5 while  $b \neq \$L$  do
6   if  $\hat{b} = b$  then
7      $\hat{b} \leftarrow \text{succ}(\hat{b}, U)$ ;         /* move to the next character in  $U$  */
8      $b \leftarrow \text{succ}(b, L)$ ;           /* move to the next character in  $L$  */
9   else if  $\hat{b} \succ b$  then
10     $M \leftarrow M \cup \{aub\}$ ;           /* output  $aub$  */
11     $b \leftarrow \text{succ}(b, L)$ ;         /* move to the next character in  $L$  */
12 return  $M$ ;
```

Let $\mathbf{B} \in \{0, 1\}^k \setminus \{0^k\}$ be an input bit vector of length $k > 2$. We redefine the labels of the nodes of $\text{DAWG}(\mathcal{S})$ such that $\text{label}(v)[i] = 1$ iff v is a substring of S_i for $1 \leq i \leq k$. Namely, $\text{label}(v)$ is now also a bit vector of length k .

Let $aub \in \Sigma^*$ ($a, b \in \Sigma$ and $u \in \Sigma^*$) be a candidate of an element of $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ as in the previous sections, where the suffix link of node au points to node u and node u has an out-going edge labeled b . Then, it follows from the definition of $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ that $aub \in \text{MAW}(\mathcal{S}_{\mathbf{B}})$ iff

- (A) $\text{label}(aub)[i] = 0$, $\text{label}(au)[i] = 1$, and $\text{label}(ub)[i] = 1$ (i.e. $aub \in \text{MAW}(S_i)$),
or
- (A') au has no out-going edge labeled b , $\text{label}(au)[i] = 1$, and $\text{label}(ub)[i] = 1$
(i.e. $aub \in \text{MAW}(S_i)$)

for all $1 \leq i \leq k$ with $\mathbf{B}[i] = 1$, and

- (B) $\text{label}(aub)[i] = 1$ (i.e. $aub \in \text{Substr}(S_i)$), or
- (C) $\text{label}(aub)[i] = 0$, and $\text{label}(au)[i] = 0$ or $\text{label}(ub)[i] = 0$
(i.e. $aub \in \text{nonMAW}(S_i)$), or
- (C') au has no out-going edge labeled b , and $\text{label}(au)[i] = 0$ or $\text{label}(ub)[i] = 0$
(i.e. $aub \in \text{nonMAW}(S_i)$)

for all $1 \leq i \leq k$ with $\mathbf{B}[i] = 0$.

For each node au in $\text{DAWG}(\mathcal{S})$ whose suffix link points to node u , we create a united single skip link $\text{schar}(ub)$ for the children ub of node u such that $b \in \text{schar}(ub)$ iff $\text{label}(ub)[i] = 1$ for every i with $\mathbf{B}[i] = 1$.

After the above preprocessing is finished, we proceed to the scanning phase of our algorithm. For each node au , we scan the skip links $\text{char}(aub)$ and $\text{schar}(ub)$ in parallel, analogously to the case with $k = 2$. Let $\hat{b} \in \text{char}(aub)$ and $b \in \text{schar}(ub)$. Our algorithm compares these characters in sorted order while keeping the invariant $\hat{b} \succeq b$ as in the case with $k = 2$.

When the comparison falls into the case “ $\hat{b} = b$ ”, then we output aub as an element of $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ if Case (A) is satisfied and if Case (B) or Case (C) is satisfied. When the comparison falls into the case “ $\hat{b} \succ b$ ”, then we output aub as an element of $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ if Cases (A’) and (C’) are both satisfied.

This already gives us an $O(nk)$ -time algorithm for computing $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ using $O(n(k + \log n))$ bits of working space, or alternatively $O(n\lceil k/\log n \rceil)$ words of working space in the word RAM model with machine word size $\omega = \log n$.

We can speed up checking Cases (A), (B), (C) for each node au by using bit masks of size $\omega = \log n$ each stored at nodes aub , au , and ub , from $O(k)$ time to $O(\lceil k/\log n \rceil)$ time. For Cases (A’) and (C’), it suffices for us to use only the bit masks stored at nodes au and ub , since node aub does not exist in these cases and we detect this as a result of “ $\hat{b} \succ b$ ” comparison.

Theorem 3 (Efficient MAW computation for a set of k strings). *Given a set $\mathcal{S} = \{S_1, \dots, S_k\}$ of k strings of total length n and a bit vector $\mathbf{B} \in \{0, 1\}^k \setminus \{0^k\}$, one can compute $\text{MAW}(\mathcal{S}_{\mathbf{B}})$ in $O(n\lceil k/\log n \rceil + |\text{MAW}(\mathcal{S}_{\mathbf{B}})|)$ time and $O(n(k + \log n))$ bits of working space (or alternatively $O(n\lceil k/\log n \rceil)$ words of working space), for integer alphabets of polynomial size in n .*

7 Discussions

Béal et al. [6] considered a different version of MAWs $\text{MAW}'(\mathcal{S})$ for a set \mathcal{S} of k strings, where a string $w = aub$ is a MAW for $\mathcal{S} = \{S_1, \dots, S_k\}$ if $aub \notin \text{Substr}(\mathcal{S})$, $au \in \text{Substr}(S_i)$ and $ub \in \text{Substr}(S_j)$ for some $1 \leq i, j \leq k$. They gave an $O(\sigma n)$ -time and space solution for computing $\text{MAW}'(\mathcal{S})$. This version of MAWs can be computed in optimal $O(n + |\text{MAW}'(\mathcal{S})|)$ time, independently of k , by running our algorithm without skip links. Ayad et al. [3] considered the problem of computing the same version of MAWs of length up to $\ell > 1$.

Independently to our work, the recent work by Béal and Crochemore [5] considered the following problem: Let \mathcal{T} and \mathcal{R} be sets of strings, where \mathcal{T} is called a target and \mathcal{R} is called a reference. A \mathcal{T} -specific string with respect to \mathcal{R} is a string u such that $u \in \text{Substr}(\mathcal{T})$, $u \notin \text{Substr}(\mathcal{R})$, $v \in \text{Substr}(\mathcal{R})$ for any proper substring v of u . By definition, a string u is a \mathcal{T} -specific string with respect to \mathcal{R} if and only if $u \in \text{MAW}(\mathcal{R}) \cap \text{Substr}(\mathcal{T})$. Béal and Crochemore [5] showed an algorithm for finding all \mathcal{T} -specific strings w.r.t. \mathcal{R} in $O(n\sigma)$ -time and $O(n)$ space, where n is the total length of the strings in \mathcal{T} and \mathcal{R} , assuming that the edges of the DAWG are represented by transition matrices (Proposition 2, [5]). Their algorithm also uses the DAWG built on \mathcal{T} and \mathcal{R} and marks its nodes in an appropriate way (Proposition 1, [5]). This marking technique is very similar to our skip links from Sect. 5 for the case of $k = 2$, and thus our algorithm can be extended to solve this problem in $O(n)$ time and space for integer alphabets.

Acknowledgments. This work was supported by JSPS KAKENHI Grant Numbers JP23H04381 (TM), JP21K17705, JP23H04386 (YN), JP22H03551 (SI), JP20H04141 (HB).

References

1. Akagi, T., et al.: Combinatorics of minimal absent words for a sliding window. *Theor. Comput. Sci.* **927**, 109–119 (2022). <https://doi.org/10.1016/j.tcs.2022.06.002>
2. Almirantis, Y., et al.: On avoided words, absent words, and their application to biological sequence analysis. *Algorithms Mol. Biol.* **12**(1), 5 (2017)
3. Ayad, L.A.K., Badkobeh, G., Fici, G., Héliou, A., Pissis, S.P.: Constructing anti-dictionaries of long texts in output-sensitive space. *Theory Comput. Syst.* **65**(5), 777–797 (2021)
4. Barton, C., Héliou, A., Mouchard, L., Pissis, S.P.: Linear-time computation of minimal absent words using suffix array. *BMC Bioinform.* **15**(1), 388 (2014)
5. Béal, M., Crochemore, M.: Fast detection of specific fragments against a set of sequences. In: Drewes, F., Volkov, M. (eds.) *Developments in Language Theory. DLT 2023*. LNCS, vol. 13911, pp. 51–60. Springer, Cham (2023). https://doi.org/10.1007/978-3-031-33264-7_5
6. Béal, M., Crochemore, M., Mignosi, F., Restivo, A., Sciortino, M.: Computing forbidden words of regular languages. *Fundam. Inform.* **56**(1–2), 121–135 (2003)
7. Béal, M.-P., Mignosi, F., Restivo, A.: Minimal forbidden words and symbolic dynamics. In: Puech, C., Reischuk, R. (eds.) *STACS 1996*. LNCS, vol. 1046, pp. 555–566. Springer, Heidelberg (1996). https://doi.org/10.1007/3-540-60922-9_45
8. Belazzougui, D., Cunial, F., Kärkkäinen, J., Mäkinen, V.: Versatile Succinct Representations of the Bidirectional Burrows-Wheeler Transform. In: Bodlaender, H.L., Italiano, G.F. (eds.) *ESA 2013*. LNCS, vol. 8125, pp. 133–144. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-40450-4_12
9. Blumer, A., Blumer, J., Haussler, D., Ehrenfeucht, A., Chen, M.T., Seiferas, J.I.: The smallest automaton recognizing the subwords of a text. *Theor. Comput. Sci.* **40**, 31–55 (1985)
10. Blumer, A., Blumer, J., Haussler, D., McConnell, R., Ehrenfeucht, A.: Complete inverted files for efficient text retrieval and analysis. *J. ACM* **34**(3), 578–595 (1987). <https://doi.org/10.1145/28869.28873>
11. Chairungsee, S., Crochemore, M.: Using minimal absent words to build phylogeny. *Theor. Comput. Sci.* **450**, 109–116 (2012)
12. Charalampopoulos, P., Crochemore, M., Fici, G., Mercas, R., Pissis, S.P.: Alignment-free sequence comparison using absent words. *Inf. Comput.* **262**, 57–68 (2018)
13. Charalampopoulos, P., Crochemore, M., Pissis, S.P.: On extended special factors of a word. In: Gagie, T., Moffat, A., Navarro, G., Cuadros-Vargas, E. (eds.) *SPIRE 2018*. LNCS, vol. 11147, pp. 131–138. Springer, Cham (2018). https://doi.org/10.1007/978-3-030-00479-8_11
14. Crawford, T., Badkobeh, G., Lewis, D.: Searching page-images of early music scanned with OMR: a scalable solution using minimal absent words. In: *ISMIR 2018*, pp. 233–239 (2018)
15. Crochemore, M., Mignosi, F., Restivo, A., Salemi, S.: Data compression using antidictionaries. *Proc. IEEE* **88**(11), 1756–1768 (2000)
16. Crochemore, M., Héliou, A., Kucherov, G., Mouchard, L., Pissis, S.P., Ramusat, Y.: Absent words in a sliding window with applications. *Inf. Comput.* **270**, 104461 (2020)
17. Crochemore, M., Mignosi, F., Restivo, A.: Automata and forbidden words. *Inf. Process. Lett.* **67**(3), 111–117 (1998)

18. Crochemore, M., Navarro, G.: Improved antidictionary based compression. In: 12th International Conference of the Chilean Computer Science Society, 2002. Proceedings, pp. 7–13. IEEE (2002)
19. Fici, G.: Minimal forbidden words and applications. Ph.D. thesis, Università di Palermo and Université Paris-Est Marne-la-Vallée (2006)
20. Fujishige, Y., Tsujimaru, Y., Inenaga, S., Bannai, H., Takeda, M.: Computing DAWGs and minimal absent words in linear time for integer alphabets. In: MFCS 2016, vol. 58, pp. 38:1–38:14 (2016)
21. Fujishige, Y., Tsujimaru, Y., Inenaga, S., Bannai, H., Takeda, M.: Linear-time computation of DAWGs, symmetric indexing structures, and MAWs for integer alphabets. *Theor. Comput. Sci.* (2023, to appear)
22. Koulouras, G., Frith, M.C.: Significant non-existence of sequences in genomes and proteomes. *Nucleic Acids Res.* **49**(6), 3139–3155 (2021)
23. Mieno, T., et al.: Minimal unique substrings and minimal absent words in a sliding window. In: Chatzigeorgiou, A., et al. (eds.) SOFSEM 2020. LNCS, vol. 12011, pp. 148–160. Springer, Cham (2020). https://doi.org/10.1007/978-3-030-38919-2_13
24. Pratas, D., Silva, J.M.: Persistent minimal sequences of SARS-CoV-2. *Bioinformatics* **36**(21), 5129–5132 (2020)