

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and
Information Systems

School of Computing and Information Systems

5-2022

Optimal in-place suffix sorting

Zhize LI

Singapore Management University, zhizeli@smu.edu.sg

Jian LI

Hongwei HUO

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

LI, Zhize; LI, Jian; and HUO, Hongwei. Optimal in-place suffix sorting. (2022). *Information and Computation*. 285, 1-25.

Available at: https://ink.library.smu.edu.sg/sis_research/8690

This Journal Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.



Optimal in-place suffix sorting [☆]

Zhize Li ^{a,*}, Jian Li ^{a,*}, Hongwei Huo ^b

^a Institute for Interdisciplinary Information Sciences, Tsinghua University, China

^b Computer Science Department, Xidian University, China



ARTICLE INFO

Article history:

Received 30 July 2021

Received in revised form 29 October 2021

Accepted 7 November 2021

Available online 10 November 2021

Keywords:

Suffix sorting

Suffix array

In-place

Optimal time

ABSTRACT

The suffix array is a fundamental data structure for many applications that involve string searching and data compression. We obtain the first in-place suffix array construction algorithms that are optimal both in time and space for (read-only) integer alphabets. We make the following contributions:

1. For integer alphabets, we obtain the first linear time suffix sorting algorithm which uses only $O(1)$ workspace. The input string may be modified during the execution of the algorithm, but should be restored upon termination of the algorithm.
2. We strengthen the first result by providing the first in-place linear time algorithm for read-only integer alphabets with $|\Sigma| = O(n)$ (i.e., the input string cannot be modified). This algorithm settles the open problem posed by Franceschini and Muthukrishnan in ICALP 2007.
3. Besides, for the read-only general alphabets (i.e., only comparisons are allowed), we present an optimal in-place $O(n \log n)$ time suffix sorting algorithm.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

In SODA 1990, suffix arrays were introduced by Manber and Myers [3] as a space-saving alternative to suffix trees [4,5]. Since then, it has been used as a fundamental data structure for many applications in string processing, data compression, text indexing, information retrieval and computational biology [6–10]. Particularly, the suffix arrays are often used to compute the Burrows-Wheeler transform [11] and Lempel-Ziv factorization [12]. Comparing with suffix trees, suffix arrays use much less space in practice. Abouelhoda et al. [13] showed that any problem which can be computed using suffix trees can also be solved using suffix arrays with the same asymptotic time complexity, which makes suffix arrays very attractive both in theory and in practice. Hence, suffix arrays have been studied extensively over the last 20 years (see e.g., [14–20]). We refer the readers to the surveys [21,22] for many suffix sorting algorithms.

In 1990, Manber and Myers [3] obtained the first $O(n \log n)$ time suffix sorting algorithm over general alphabets. In 2003, Ko and Aluru [15], Kärkkäinen and Sanders [14] and Kim et al. [23] independently obtained the first linear time algorithm for suffix sorting over integer alphabets. Clearly, these algorithms are optimal in terms of asymptotic time complexity.

[☆] This paper extends a preliminary version appeared in the Proceedings of the 25th International Symposium on String Processing and Information Retrieval (SPIRE 2018) [1] and a one-page summary appeared in the Proceedings of the 28th IEEE Data Compression Conference (DCC 2018) [2].

* Corresponding authors at: Tsinghua University, Beijing 100084, China.

E-mail addresses: zhizeli.thu@gmail.com (Z. Li), lijian83@mail.tsinghua.edu.cn (J. Li), hwhuo@mail.xidian.edu.cn (H. Huo).

However, in many applications, the computational bottleneck is the *space* as we need the space-saving suffix arrays instead of suffix trees, and significant efforts have been made in developing *lightweight* (in terms of space usage) suffix sorting algorithms for the last decade (see e.g., [24,25,15,26,27,17,28,18–20]). In particular, the ultimate goal in this line of work is to obtain *in-place algorithms* (i.e., $O(1)$ additional space), which are also asymptotically optimal in time.

1.1. Problem setting

Problem: Given a string $T = T[0 \dots n - 1]$ with n characters, we need to construct the *suffix array* (SA) which contains the *indices* of all sorted suffixes of T (see Definition 1 for the formal definition of SA).

We consider the following three popular settings. Note that the constant alphabets (e.g., ASCII code) is a special case of integer alphabets, and the (read-only) integer alphabets is commonly used in practice. We measure the space usage of an algorithm in the unit of *words* same as [17,20], i.e., the standard RAM model is used. A word contains $\lceil \log n \rceil$ bits. One standard arithmetic or bitwise boolean operation on word-sized operands costs $O(1)$ time.

1. Integer alphabets: Each $T[i] \in [1, |\Sigma|]$ where the cardinality of the alphabets is $|\Sigma| \leq n$ and each $T[i]$ is stored in a word. The input string T may be modified by the algorithm, but should be restored upon termination of the algorithm.
2. Read-only integer alphabets: Each $T[i] \in [1, |\Sigma|]$ where $|\Sigma| = O(n)$. Moreover, the input string T is *read-only*. Each $T[i]$ can be read in $O(1)$ time. Here we allow $|\Sigma| = O(n)$ rather than $|\Sigma| \leq n$ in Case 1.
3. Read-only general alphabets: The only operations allowed on the characters of T are comparisons. The input string T is read-only and we assume that each comparison takes $O(1)$ time. We *cannot* write the input space, make bit operations, even copy an input character $T[i]$ to the work space. Clearly, $\Omega(n \log n)$ time is a lower bound for suffix sorting in this case, as it generalizes comparison-based sorting.

The *workspace* used by an algorithm is the total space needed by the algorithm, excluding the space required by the input string T and the output suffix array SA. An algorithm which uses $O(1)$ words workspace to construct SA is called an *in-place algorithm*. See Tables 1¹ and 2 for an overview.

1.2. Related work and our contributions

1.2.1. Integer alphabets

In this case, we allow the algorithm to modify the string T during the execution of the algorithm. We also describe how to restore T (if needed) at the end of the algorithm in Appendix C. Chan et al. [29] denote this model as the *restore model* in their paper. We list several previous results and our new result in Table 1. Earlier algorithms that require more than $O(n)$ words workspace (see Table 1) do not need to modify the string T as they can afford to create a new array with n words to store the input.

Nong et al. [30,19] obtained the first nearly linear time algorithm that used sublinear workspace. Recently, Nong [20] obtained a linear time algorithm which used $|\Sigma|$ words workspace without modifying the string T . We improve their results as in the following theorem.

Theorem 1. *There is an in-place linear time algorithm for suffix sorting over integer alphabets with $|\Sigma| \leq n$.*

Our algorithm is based on the induced sorting framework developed in [15] (which is also used in several previous algorithms [17,21,18,30,19,20]). We develop a few elementary, yet effective tricks to further reduce the space usage to constant. The proposed algorithm and the new tricks are also useful for the read-only integer and general alphabets.

1.2.2. Read-only integer alphabets

Now, we consider the more difficult case where the input string T is read-only. This is the main contribution of this paper. There are many existing algorithms for this case. See Table 1 for an overview. In ICALP 2007, Franceschini and Muthukrishnan [17] posed an open problem for designing an in-place algorithm that takes $o(n \log n)$ time or ultimately $O(n)$ time for (read-only) integer alphabets with $|\Sigma| \leq n$ (in fact, they did not specify whether the input string T is read-only or not). The current best result along this line is provided by Nong [20], which used $|\Sigma|$ words workspace (Nong's algorithm is in-place if $|\Sigma| = O(1)$, i.e., constant alphabets). Note that in the worst case $|\Sigma|$ can be as large as $O(n)$.

In this paper, we settle down this open problem by providing the first optimal linear time in-place algorithm, as in the following theorem. Note that our result is in fact slightly stronger since we allow $|\Sigma| = O(n)$ instead of $|\Sigma| \leq n$ mentioned in the open problem [17].

Theorem 2 (Main Theorem). *There is an in-place linear time algorithm for suffix sorting over integer alphabets, even if the input string T is read-only and the size of the alphabet $|\Sigma|$ is $O(n)$.*

¹ Some previous algorithms state the space usages in terms of bits. We convert them into words.

Table 1
Time and workspace of suffix sorting alg. for (read-only) integer alphabets Σ .

Time	Workspace (words)	Algorithms
$O(n^2)$	$O(n)$	[31]
$O(n \log^2 n)$	$O(n)$	[32]
$O(n\sqrt{ \Sigma \log(n/ \Sigma)})$	$O(n)$	[33]
$O(n \log n)$	$O(n)$	[3,34]
$O(n \log \log n)$	$O(n)$	[35]
$O(n)$	$O(n)$	[23,14,15]
$O(n \log \log \Sigma)$	$O(n \log \Sigma / \log n)$	[26]
$O(vn)$	$O(n/\sqrt{v}) \quad v \in [1, \sqrt{n}]$	[16]
$O(n)$	$n + n/\log n + O(1)$	[18]
$O(n^2 \log n)$	$cn + O(1) \quad c < 1$	[24,27]
$O(\frac{1}{\epsilon}n)^*$	$n^\epsilon + n/\log n + O(1)^a$	[30,19]
$O(n^2 \log n)$	$ \Sigma + O(1)$	[36]
$O(n \log \Sigma)$	$ \Sigma + O(1)$	[28]
$O(n)$	$ \Sigma + O(1)$	[20]
$O(n)$	$O(1)$	This paper

T is read-only in all algorithms except in the fifth to last row (marked with *).

^a Nong et al. [30,19] assumed that the word size is 32 bits and any integer can fit into one word. The result listed here is under the standard assumption that a word contains $\lceil \log n \rceil$ bits. It is not hard to verify that the bucket array B in their algorithm requires n^ϵ words. They also need an n bits array (or equivalently $n/\log n$ words).

Table 2
Time and workspace of suffix sorting algorithms for read-only general alphabets.

Time	Workspace (words)	Algorithms
$O(n \log n)$	$O(n)$	[3,34]
$O(vn + n \log n)$	$O(v + n/\sqrt{v}) \quad v \in [2, n]$	[25]
$O(vn + n \log n)$	$O(n/\sqrt{v}) \quad v \in [1, \sqrt{n}]$	[16]
$O(n \log n)$	$O(1)$	[17]
$O(n \log n)$	$O(1)$	This paper

1.2.3. Read-only general alphabets

Now, we consider the case where the only operations allowed on the characters of string T (read-only) are comparisons. See Table 2 for an overview of the results. In 2002, Manzini and Ferragina [24] posed an open problem, which asked whether there exists an $O(n \log n)$ time algorithm using $o(n)$ workspace. In 2007, Franceschini and Muthukrishnan [17] obtained the first in-place algorithm that runs in optimal $O(n \log n)$ time. Their conference paper is somewhat complicated and densely-argued.

We also give an optimal in-place algorithm which achieves the same result, as in the following theorem. In addition, our algorithm does not make any bit operations while theirs uses bit operations heavily. Our algorithm is also arguably simpler.

Theorem 3. *There is an in-place $O(n \log n)$ time algorithm for suffix sorting over general alphabets, even if the input string T is read-only and only comparisons between characters are allowed.*

1.3. Difficulties and our approach

1.3.1. Difficulties

Typically, the suffix sorting algorithms are recursive algorithms. The size of the recursive (reduced) sub-problem is usually less than half of the current problem. See e.g., [15,16,21,17–20]. However, all previous algorithms require extra arrays, e.g., *bucket array* (which needs $|\Sigma|$ words at the top recursive level and $n/2$ words at the deep recursive levels), *type array* (which needs $n/\log n$ words) and/or other auxiliary arrays (which need up to $O(n)$ words), to construct the reduced problems and use the results of the reduced problems to sort the original suffixes.²

In particular, Nong et al. [18] made a breakthrough by providing the SA-IS algorithm which only required one bucket array (which needs $\max\{|\Sigma|, n/2\}$ words) and one type array ($n/\log n$ words). Note that the bucket array and type array are reused for each recursive level.

Currently, the best result is provided by Nong [20]. However, Nong’s algorithm requires the bucket array at the top recursive level, but not required at the deep levels. Hence, it needs $|\Sigma|$ words instead of $\max\{|\Sigma|, n/2\}$ words. Note that $|\Sigma|$ can be $O(n)$ in the worst case for integer alphabets. For the type array, Nong used this bucket array to implicitly indicate the type information.

² The definitions of bucket array and type array can be found in Section 2.

Thus, the main technical difficulty is to remove the workspace for the bucket array at the top recursive level since there is no extra space to use. Note that it is non-trivial since T is read-only and SA needs to store the final order of all suffixes. Besides, the previous sorting steps or tricks may not work if one removes the bucket array. For example, Nong [20] used the bucket array to indicate the type information. If the bucket array is removed, then one would need the type array.

1.3.2. Our approach

We briefly describe our optimal in-place linear time suffix sorting algorithms that overcome these difficulties. We provide an *interior counter trick* which can implicitly represent the dynamic *LF/RF-entry* information (see Section 2 for the definition) in SA. Besides, we provide a *pointer data structure* which can represent the bucket heads/tails in SA. Combining these two techniques, we can remove the workspace needed by the bucket array entirely. Note that it is non-trivial for the top recursive level which is the most difficult part, since the pointer data structure needs *nonconstant* workspace and we only have $O(1)$ extra workspace. In our algorithm, we divide the sorting step into two stages to address this issue. In order to remove the type array, we provide some useful properties and observations which allows us to retrieve the type information efficiently. For the general alphabets case, we provide simple sorting steps and extend the interior counter trick to obtain an optimal in-place $O(n \log n)$ time suffix sorting algorithm.

Organization: The remaining of the paper is organized as follows. Section 2 covers the preliminary knowledge. In Section 3, 4 and 5, we describe the framework and the details of our optimal in-place suffix sorting algorithms for the integer alphabets, read-only integer alphabets and read-only general alphabets, respectively. Finally, we conclude in Section 6.

2. Preliminaries

Given a string $T = T[0 \dots n-1]$ with n characters, the suffixes of T are $T[i \dots n-1]$ for all $i \in [0, n-1]$, where $T[i \dots j]$ denotes the substring $T[i]T[i+1] \dots T[j]$ in T . To simplify the argument, we assume that the final character $T[n-1]$ is a sentinel which is lexicographically smaller than any other characters in Σ . Without loss of generality, we assume $T[n-1] = 0$.³ Any two suffixes in T must be different since their lengths are different, and their lexicographical order can be determined by comparing their characters one by one until we see a difference due to the existence of the sentinel.

Definition 1. The suffix array SA contains the indices of all suffixes of T which are sorted in lexicographical order, i.e., $\text{suf}(\text{SA}[i]) < \text{suf}(\text{SA}[j])$ for all $i < j$, where $\text{suf}(i)$ denotes the suffix $T[i \dots n-1]$.

For example, if $T = "1220"$, then all suffixes are $\{1220, 220, 20, 0\}$ and $\text{SA} = [3, 0, 2, 1]$. Note that SA always uses n words no matter what the alphabets Σ are, since it contains the permutation of $\{0, \dots, n-1\}$, where n is the length of T .

A suffix $\text{suf}(i)$ is said to be *S-suffix* (S-type suffix) if $\text{suf}(i) < \text{suf}(i+1)$. Otherwise, it is *L-suffix* (L-type suffix) [15]. The last suffix $\text{suf}(n-1)$ containing only the single character 0 (the sentinel) is defined to be an S-suffix. Equivalently, the $\text{suf}(i)$ is S-suffix if and only if (1) $i = n-1$; or (2) $T[i] < T[i+1]$; or (3) $T[i] = T[i+1]$ and $\text{suf}(i+1)$ is S-suffix. Obviously, the types can be computed by a linear scan of T (from $T[n-1]$ to $T[0]$). We further define the type of a character $T[i]$ to *S-type* (or *L-type* resp.) if $\text{suf}(i)$ is S-suffix (or L-suffix resp.). A substring $T[i \dots j]$ is called an *S-substring* if both $T[i]$ and $T[j]$ are S-type, and there is no other S-type characters between them, or $i = j = n-1$ (the single sentinel). We can define *L-substring* similarly.

Example: We use the following running example for the integer alphabets case throughout the paper. Consider a string $T[0 \dots 12] = "2113311331210"$.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S

E.g., $T[2]$ is S-type since $T[2] < T[3]$. All S-substrings are $\{11, 1331, 11, 1331, 1210, 0\}$. \square

Obviously, the indices of all suffixes, which begin with the same character, must appear consecutively in SA. We denote a subarray in SA for these suffixes with the same beginning character as a *bucket*, where the *head* and the *tail* of a bucket refer to the first and the last index of the bucket in SA respectively. Moreover, we define the first common character as its *bucket character*. We often use the bucket character to index the bucket. For example, if the bucket character is $T[i]$, we refer to this bucket as bucket $T[i]$. Sometimes we say that we place suffix $\text{suf}(i)$ of T into SA, it always means that we place its corresponding index i into SA since $\text{suf}(i)$ is a substring.

The *induced sorting* technique, developed by Ko and Aluru [15], is responsible for many recent advances of suffix sorting algorithms [21,17,18,30,19,20], and is also crucial to us. It can be used to induce the lexicographical order of L-suffixes from the sorted S-suffixes. Before introducing the induce sorting technique, we need the following useful property with respect to L-suffixes and S-suffixes (the proof simply follows from the definition of L- and S-suffix).

³ Some previous papers use $\$$ to denote the sentinel. We use 0 here since we consider the integer alphabets.

Property 1 ([15]). In any bucket, S-suffixes always appear after the L-suffixes in SA, i.e., if an S-suffix and an L-suffix begin with the same character, the L-suffix is always smaller than the S-suffix.

Now, we briefly introduce the standard induced sorting technique which needs the *bucket array* and *type array* explicitly. The bucket array contains $|\Sigma|$ integers and each denotes the position of a bucket head/tail in SA. The type array contains n bits and each entry denotes an L/S-type information for T (i.e., 0/1 for L/S-type).

Inducing the order of L-suffixes from the sorted S-suffixes: Assume that all indices of the sorted S-suffixes are already in their correct positions in SA (i.e., in the tail of their corresponding buckets in SA). Now, we define some new notations (e.g., *LF/RF-entry*) to simplify the representation. We scan SA from left to right (i.e., from SA[0] to SA[n - 1]). We maintain an *LF-pointer* (leftmost free pointer) for each bucket which points to the leftmost free entry (called the *LF-entry*) of the bucket. The LF-pointers initially point to the head of their corresponding buckets. When we scan SA[i], let $j = SA[i] - 1$. If $\text{suf}(j)$ is an L-suffix (indicated by the type array), we place the index j of $\text{suf}(j)$ into the LF-entry of bucket $T[j]$, and then let the LF-pointer of this bucket $T[j]$ point to the next free entry. The LF-pointers are maintained in the bucket array. If $\text{suf}(j)$ is an S-suffix, we do nothing (since all S-suffixes are already sorted in the correct positions). We give a running example in Appendix A.1.

Sorting all S-suffixes from the sorted L-suffixes is completely symmetrical: we scan SA from right to left, maintaining an *RF-pointer* (rightmost free pointer) for each bucket which points to the *RF-entry* (rightmost free entry) of the bucket.

Lemma 1 ([15]). Suppose all S-suffixes (or L-suffixes resp.) of T are already sorted. Then using induced sorting, all L-suffixes (or S-suffixes resp.) can be sorted correctly.

The idea of induced sorting is that the lexicographical order between $\text{suf}(i)$ and $\text{suf}(j)$ is decided by the order of $\text{suf}(i+1)$ and $\text{suf}(j+1)$ if $\text{suf}(i)$ and $\text{suf}(j)$ are in the same bucket (i.e., $T[i] = T[j]$). We only need to specify the correct order of these L-suffixes in the same buckets since we always place the L-suffixes in their corresponding buckets. Consider two L-suffixes $\text{suf}(i)$ and $\text{suf}(j)$ in the same bucket. We have $\text{suf}(i+1) < \text{suf}(j+1)$ and $\text{suf}(j+1) < \text{suf}(i+1)$ by the definition of L-suffix. Since we scan SA from left to right, $\text{suf}(i+1)$ and $\text{suf}(j+1)$ must appear earlier than $\text{suf}(i)$ and $\text{suf}(j)$. Hence the correctness of induced sorting is not hard to prove by induction.

Inducing the order of L-suffixes from the sorted LMS-suffixes: A suffix $\text{suf}(i)$ is called an *LMS-suffix* (leftmost S-type) if $T[i]$ is S-type and $T[i - 1]$ is L-type, for $i \geq 1$. Nong et al. [18] observed that we can sort all L-suffixes from the sorted LMS-suffixes (instead of S-suffixes) if they are stored in the tail of their corresponding buckets in SA. Roughly speaking, the idea is that in the induced sorting, only LMS-suffixes are useful for sorting L-suffixes. One difference from the standard induced sorting is that we may scan some empty entries in SA. However, the empty entries can be ignored and all L-suffixes can still be sorted correctly. We provide a running example in Appendix A.2.

Lemma 2 ([18]). Suppose all LMS-suffixes of T are already sorted and stored in the tail of their buckets. Then using induced sorting, all L-suffixes can be sorted correctly.

After we sort all L-suffixes from the sorted LMS-suffixes, we can induce the order of all S-suffixes from the sorted L-suffixes by Lemma 1, and sort all suffixes. Now, we introduce how to sort the LMS-suffixes. First, we define some notations. A character $T[i]$ of T is called *LMS-character* if $\text{suf}(i)$ is LMS-suffix. A substring $T[i \dots j]$ is called an *LMS-substring* if both $T[i]$ and $T[j]$ are LMS-characters, and there is no other LMS-character between them, or $i = j = n - 1$ (the single sentinel). Similarly, we can define *LML-suffix* (leftmost L-type) and *LML-substring*.

Sort the LMS-suffixes: If we know the lexicographical order of all LMS-substrings, then we can use their ranks to construct the reduced problem T_1 . Sorting the suffixes of T_1 is equivalent to sorting the LMS-suffixes of T (see the following example). Nong et al. [18] showed that we can use the same induced sorting step to sort all LMS-substrings from the sorted LMS-characters of T . We briefly sketch their idea. We refer the readers to [18] for the details. We define the *LMS-prefix* of a suffix $\text{suf}(i)$ to be $T[i \dots j]$, where $j > i$ is the smallest position in $\text{suf}(i)$ such that $T[j]$ is an LMS character (e.g., the LMS-prefix of $\text{suf}(4)$ is "31"). Suppose all LMS-characters are stored in the tail of their corresponding buckets in SA. First, we sort all LMS-prefix of L-suffixes from the sorted LMS-characters, using one scan of induced sorting from left to right (the same as induce the order of L-suffixes from LMS-suffixes). Then we sort all LMS-prefix of S-suffixes from the sorted LMS-prefix of L-suffixes (in the same way as inducing the order of S-suffixes from L-suffixes). After this, we have sorted all LMS-substrings since all LMS-substrings are LMS-prefix of S-suffixes by the definition of LMS-prefix. The correctness proof follows the same argument as in the standard setting.

Example: Continue the running example:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
LMS		*				*				*			*

Note that the LMS-substrings are $\{11331, 11331, 1210, 0\}$. Their ranks in lexicographical order are $\{1, 1, 2, 0\}$. Thus, the reduced problem is $T_1 = 1120$. The order of the suffixes of T_1 is the same as the order of corresponding LMS-suffixes of T . The suffix array of the reduced problem T_1 is solved recursively. \square

Note that in this preliminary section, the induced sorting steps are not *in-place* since they require explicit storage for the bucket and type arrays.

3. Suffix sorting for integer alphabets

3.1. Framework

To avoid confusion, we recall that an LMS-character is a single character, an LMS-substring is a substring which begins with an LMS-character and ends with an LMS-character (and there is no other LMS-character between them) or is the single sentinel, and an LMS-suffix is a suffix of T which begins with an LMS-character. Our optimal in-place suffix sorting algorithm for integer alphabets consists of the following steps.

1. (Section 3.2) Rename T .
2. (Section 3.3) Sort all LMS-characters of T .
3. (Section 3.4) Induced sort all LMS-substrings from the sorted LMS-characters.
4. (Section 3.5) Construct the reduced problem T_1 from the sorted LMS-substrings.
5. (Section 3.6) Sort the LMS-suffixes by solving T_1 recursively.
6. (Section 3.7) Induced sort all suffixes of T from the sorted LMS-suffixes.

In a high level, the framework is similar to several other previous algorithms based on induced sorting [15,17,21,18,30,19,20], and in particular to [18]. Our algorithm differs in the detailed implementation of the above steps to obtain the first in-place algorithm (i.e., remove the bucket array and type array). We describe the details of the above steps in the following sections. We also describe how to restore T (if needed) at the end of the algorithm in Appendix C.

3.2. Rename T

In this section, we rename each L-type character of T to be the index of its bucket head and each S-type character of T to be the index of its bucket tail (Nong et al. [19] has a similar renaming step). The correctness of the step is shown in the following Lemma 3.

Lemma 3. *The renaming step does not change the lexicographical order of all suffixes of T .*

Proof. For any two suffixes, beginning with the same character, the L-suffix is smaller than the S-suffix (Property 1). Hence, the renaming step does not change the relative orders of all suffixes. \square

Example: We illustrate the renaming process in our running example.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(12)	(11	1	5	9	2	6)	(10	0)	(4	8	3	7)
Bucket	(0)	(1	1	1	1	1	1)	(2	2)	(3	3	3	3)

After renaming, we get T' as following:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T'	7	6	6	9	9	6	6	9	9	6	7	1	0

For instance, $T'[0] = 7$ since $T[0]$ is L-type and the head of bucket $T[0]$ (i.e., bucket 2) is 7, and $T'[1] = 6$ since $T[1]$ is S-type and the tail of bucket $T[1]$ is 6. Note that the heads of bucket 0, 1, 2, 3 are 0, 1, 7, 9, respectively. \square

Now, we describe how to implement this step using linear time and $O(1)$ workspace. This step is very simple and similar to the counting sort (see e.g., [37, Ch. 8]). We first rename all characters to be the index of its bucket head, and then rename all S-type characters to be the index of its bucket tail.

1. First we scan T once to compute the number of times each character occurs in T and store them in SA. Then we perform a *prefix sum computation* to determine the starting position of each character (i.e., bucket head) in SA. Finally we scan T once again to rename each character as the index of its bucket head.

2. Now we need to let the S-type characters of T to be the index of its bucket tail. We scan T to compute the number of times each character occurs in T and store them in SA. Then, we scan T once again from right to left. For each S-type $T[i]$, we let it be the index of its bucket tail (i.e., just add the number of characters belonging to this bucket to $T[i]$ and minus 1). Note that if we scan T from right to left, for each $T[i]$, we can know its type is L-type or S-type in $O(1)$ time. There are two cases: 1) if $T[i] \neq T[i + 1]$, we can know its type immediately by definition; 2) if $T[i] = T[i + 1]$ then its type is the same as the type of $T[i + 1]$. We only need to maintain *one* boolean variable which represents the type of previous scanned character $T[i + 1]$.

3.3. Sort all LMS-characters

Now, we sort all LMS-characters of T , i.e., place the indices of the LMS-characters in the tail of their corresponding buckets in SA. Note that we do not have extra space to store the LF/RF-pointers/counters for each bucket to indicate the position of the free entries in the process. For this purpose, we develop a simple trick, called *interior counter trick*, which allows us to carefully use the space in SA to store the information of both the indices and the pointers. The implementation details are described below. In the steps, we use three special symbols which are Unique, Empty and Multi.⁴

Step 1. Initializing SA: First we clear SA (i.e., $SA[i] = \text{Empty}$, for all $i \in [0, n - 1]$). Then we scan T from right to left. For every $T[i]$ which is an LMS-character (this can be easily decided in constant time), do the following:

- (1) If $SA[T[i]] = \text{Empty}$, let $SA[T[i]] = \text{Unique}$ (meaning it is the unique LMS-character in this bucket). Note that after the renaming step, $T[i]$ is the index of its bucket tail.
- (2) If $SA[T[i]] = \text{Unique}$, let $SA[T[i]] = \text{Multi}$ (meaning the number of LMS-characters in this bucket is at least 2).
- (3) Otherwise, do nothing.

Step 2. Placing all indices of LMS-characters into SA: We scan T from right to left. For every $T[i]$ which is an LMS-character, we distinguish the following cases:

- (1) $SA[T[i]] = \text{Unique}$: In this case, we let $SA[T[i]] = i$ (i.e., $T[i]$ is the unique LMS-character in its bucket, and we just put its index into its bucket).
- (2) $SA[T[i]] = \text{Multi}$ and $SA[T[i] - 1] = \text{Empty}$: In this case, $T[i]$ is the first (i.e. largest index, since we scan T from right to left) LMS-character in its bucket. So if $SA[T[i] - 2] = \text{Empty}$, we let $SA[T[i] - 2] = i$ and $SA[T[i] - 1] = 1$ (i.e., we use $SA[T[i] - 1]$ as the counter for the number of LMS-characters which has been added to this bucket so far). Otherwise, $SA[T[i] - 2] \neq \text{Empty}$ (i.e., $SA[T[i] - 2]$ is in a different bucket, which implies that this bucket has only two LMS-characters). Then we let $SA[T[i]] = i$ and $SA[T[i] - 1] = \text{Empty}$ (We do not need a counter in this case and the last LMS-character belonging to this bucket will be dealt with in the later process).
- (3) $SA[T[i]] = \text{Multi}$ and $SA[T[i] - 1] \neq \text{Empty}$: In this case, $SA[T[i] - 1]$ is maintained as the counter. Let $c = SA[T[i] - 1]$. We check whether the position ($SA[T[i] - c - 2]$), i.e. $c + 2$ positions before its tail, is Empty or not. If $SA[T[i] - c - 2] = \text{Empty}$, let $SA[T[i] - c - 2] = i$ and increase $SA[T[i] - 1]$ by one (i.e., update the counter number). Otherwise $SA[T[i] - c - 2] \neq \text{Empty}$ (i.e., reaching another bucket), we need to shift these c indices to the right by two positions (i.e., move $SA[T[i] - c - 1 \dots T[i] - 2]$ to $SA[T[i] - c + 1 \dots T[i]]$), and let $SA[T[i] - c] = i$ and $SA[T[i] - c - 1] = \text{Empty}$. After this, only one LMS-character needs to be added into this bucket in the later process.
- (4) $SA[T[i]]$ is an index: From case (2) and (3), we know the current $T[i]$ must be the last LMS-character in its bucket. So we scan SA from right to left, starting with $SA[T[i]]$, to find the first position j such that $SA[j] = \text{Empty}$. Then we let $SA[j] = i$. Now, we have filled the entire bucket. However, we note that not every bucket is fully filled as we have only processed LMS-characters so far.

After the above Step 1 and 2, there may be still some special symbols Multi and the counters (because the bucket is not fully filled, so we have not shifted these indices to the right in the bucket). We need to free these position. We scan SA once more from right to left. If $SA[i] = \text{Multi}$, we shift the indices of LMS-characters in this bucket to the right by two positions (i.e., $SA[i - c - 1 \dots i - 2]$ to $SA[i - c + 1 \dots i]$) and let $SA[i - c - 1] = SA[i - c] = \text{Empty}$, where $c = SA[i - 1]$ denotes the counter.

Lemma 4. *The indices of the LMS-characters can be placed in the tail of their corresponding buckets in SA using linear time and $O(1)$ workspace.*

Proof. We only need to show that the Step 2, i.e., placing all indices of LMS-characters into SA, takes $O(n)$ time. For each scanned $T[i]$, it takes $O(1)$ time except when the $T[i]$ is the last two LMS-characters of its bucket. In this case, we need to

⁴ We use at most five special symbols in this paper. The special symbol is only used to simplify the argument and we do not have to impose any additional assumption to accommodate these symbols (including the read-only cases in Section 4 and 5). These special symbols can be handled using an extra $O(1)$ workspace. We defer the details to Appendix B.

shift the indices in this bucket when dealing with the penultimate LMS-character in its bucket, and scan the bucket when dealing with the last one. It takes $O(n)$ time since every bucket only needs to be shifted and scanned once. The space usage of this step is obvious. \square

Example: Continue our example (U, E and M denote Unique, Empty and Multi, respectively):

Step 1. Initializing SA:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	7	6	6	9	9	6	6	9	9	6	7	1	0
LMS		*				*				*			*
SA	E	E	E	E	E	E	E	E	E	E	E	E	E

After initialization:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(<u>U</u>)	(E)	(E	E	E	E	(<u>M</u>)	(E	E)	(E	E	E	E)

Step 2. Placing all indices of LMS-characters into SA:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(E)	(E	E	E	E	M)	(E	E)	(E	E	E	E)
SA	(12)	(E)	(E	E	9	1	M)	(E	E)	(E	E	E	E)
SA	(12)	(E)	(E	5	9	2	M)	(E	E)	(E	E	E	E)
SA	(12)	(E)	(1	5	9	3	M)	(E	E)	(E	E	E	E)
SA	(12)	(E)	(E	E	1	5	9)	(E	E)	(E	E	E	E)

In the last row, we remove all Multi symbols and counters. \square

3.4. Induced sort all LMS-substrings from the sorted LMS-characters

In this section, we sort all LMS-substrings from the sorted LMS-characters using induced sorting. Since all LMS-substrings are LMS-prefix of S-suffixes (Recall that the LMS-prefix of $\text{suf}(i)$ is $T[i \dots j]$, where $j > i$ is the smallest position in $\text{suf}(i)$ such that $T[j]$ is an LMS character) and sorting the LMS-prefix of all suffixes from the sorted LMS-characters is the same as sorting all suffixes from the sorted LMS-suffixes (see the preliminary Section 2).

Now, we describe the details. We divide this step into two parts.

- (1) First, we sort the LMS-prefix of all suffixes from the sorted LMS-characters. Since this part is the same as sorting all suffixes from the sorted LMS-suffixes, we will describe the details in Section 3.7.
- (2) Then, we place the indices of all sorted LMS-substrings in $\text{SA}[n - n_1 \dots n - 1]$, where n_1 denotes the number of LMS-characters. Note that the number of LMS-characters, LMS-suffixes, and LMS-substrings are the same. Moreover, $n_1 \leq \frac{n}{2}$ since any two LMS-characters are not adjacent.

The first part is described and proved in Section 3.7. Here, we only need to explain the second part how to place the indices of all sorted LMS-substrings in $\text{SA}[n - n_1 \dots n - 1]$. Now we give the following Lemma 5 to show that this step can be done in linear time using $O(1)$ workspace. For its proof, we need the Observation 1 and Lemma 6.

Lemma 5. *The indices of all sorted LMS-substrings can be placed in $\text{SA}[n - n_1 \dots n - 1]$ using linear time and $O(1)$ workspace.*

Proof. In Step (2), we scan SA from right to left to place the indices of all LMS-substrings at the end of SA. We only need to explain how to distinguish whether $T[\text{SA}[i]]$ is LMS-character or not when we scanning $\text{SA}[i]$. Note that if we can tell if $T[\text{SA}[i]]$ is S-type or not, we can also tell if $T[\text{SA}[i]]$ is an LMS-character or not since $T[\text{SA}[i]]$ is an LMS-character if and only if $T[\text{SA}[i]]$ is S-type and $T[\text{SA}[i] - 1] > T[\text{SA}[i]]$. In the scanning process, when we reach a new bucket, we can see whether this bucket contains S-type characters or not from Observation 1. Furthermore, if we know the number of S-type characters in this bucket (Lemma 6), we are done with this step since all S-type suffixes appear after the L-suffixes in any bucket (Property 1). This step costs $O(n)$ time overall since each character is scanned at most twice. \square

Example: Continue our example:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
LMS		*				*				*			*
SA	(12)	(E)	(E	E	1	5	9)	(E	E)	(E	E	E	E)

- (1) Sorting the LMS-prefixes of all suffixes (see Section 3.7):

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(11)	(1	5	9	2	6)	(10	0)	(4	8	3	7)

(2) Placing the indices of all sorted LMS-substrings in $SA[n - n_1 \dots n - 1]$:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12	
SA	E	E	E	E	E	E	E	E	E	<u>12</u>	<u>1</u>	<u>5</u>	<u>9</u>	□

Observation 1. For any bucket in SA, let t be its bucket tail. Then $T[SA[t]]$ is S-type if and only if $T[SA[t]] < T[SA[t + 1]]$. Similarly, $T[SA[h]]$ is L-type if and only if $T[SA[h]] > T[SA[h + 1]]$, where h is the bucket head.

Lemma 6. If a bucket contains S-type characters, then one can scan this bucket once to compute the number of S-type characters in this bucket using $O(1)$ workspace.

Proof. We scan this bucket from its tail to its head. For the current scanning entry $SA[i]$, there are two possibilities: 1). If $T[SA[i]] \geq T[SA[i + 1]]$, do nothing; 2). Otherwise, let j be the smallest index such that $T[k] = T[SA[i]]$ for any $k \in [j, SA[i]]$. Then we increase num by $SA[i] - j + 1$. Here variable num counts the number of S-type characters in this bucket and initially is 0. □

3.5. Construct the reduced problem T_1

In this section, we construct the smaller recursive problem T_1 . We rename the sorted LMS-substrings (obtained from the previous step) using their ranks to obtain T_1 . Note that this step is not difficult and similar to the previous algorithms (e.g., [18,20]). The order of any two LMS-substrings is determined by comparing their corresponding characters from left to right; for each pair of characters, compare their lexicographical values first and next their types. If the two characters have the same lexicographical value then the S-type is larger than L-type (same as the Definition 2.3 (Substring Order) in [18]).

Now, we spell out the details for this step. Initially, all LMS-substrings are sorted in $SA[n - n_1 \dots n - 1]$. First, we let the rank of the smallest LMS-substring corresponding to $SA[n - n_1]$ (i.e., the LMS-substring which begins from index $SA[n - n_1]$) be 0 (it must be the sentinel). Then, we scan $SA[n - n_1 + 1 \dots n - 1]$ from left to right to compute the rank for each LMS-substring. When scanning $SA[i]$, we compare the LMS-substring corresponding to $SA[i]$ and that corresponding to $SA[i - 1]$. If they are the same, $SA[i]$ gets the same rank as $SA[i - 1]$. Otherwise, the rank of $SA[i]$ is the rank of $SA[i - 1]$ plus 1. Since we have no extra space, we need to store the ranks in SA as well. In particular, the rank of $SA[i]$ is stored in $SA[\lfloor \frac{SA[i]}{2} \rfloor]$. There is no conflict since any two LMS-characters are not adjacent. Finally, we shift nonempty entries in $SA[0 \dots n - n_1 - 1]$ to the head of SA, so that the ranks occupy a consecutive segment of the space. Now, we have obtained the reduced problem T_1 which is stored in $SA[0 \dots n_1 - 1]$. In other words, $SA[i]$ ($i \in [0, n_1 - 1]$) stores the new name of the i -th LMS-substring with respect to its appearance in the input string T .

Example: Continue our example:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	7	6	6	9	9	6	6	9	9	6	7	1	0
LMS	*				*				*			*	
SA	E	E	E	E	E	E	E	E	E	12	1	5	9

After scanning $SA[n - n_1 \dots n - 1]$ (which stores the sorted LMS-substrings):

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	<u>1</u>	E	<u>1</u>	E	<u>2</u>	E	<u>0</u>	E	E	12	1	5	9

Finally, we get T_1 stored in $SA[0 \dots n_1 - 1]$ by shifting nonempty items in $SA[0 \dots n - n_1 - 1]$ to the head of SA.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	<u>1</u>	<u>1</u>	<u>2</u>	<u>0</u>	E	E	E	E	E	12	1	5	9

Note that $T_1 = "1120"$ which corresponds to the LMS-substrings {"66996", "66996", "6710", "0"}. □

First, we give an observation which helps us to identify the S-type and L-type characters of T . Then we obtain the following lemma which implies that T_1 can be obtained in linear time.

Observation 2. For any index i of T , let $j \in [i + 1, n - 1]$ be the smallest index such that $T[j] < T[j + 1]$ (So $T[j]$ is S-type). Furthermore let $k \in [i + 1, j]$ be the smallest index such that $T[l] = T[j]$ for any $k \leq l \leq j$. Then $T[k]$ is the first S-type character after index i . Moreover, all characters between $T[i]$ and $T[k]$ are L-type, and characters between $T[k]$ and $T[j]$ are S-type.

Lemma 7. T_1 can be obtained using $O(n)$ time and $O(1)$ workspace.

Proof. For the workspace part, it is obvious since we do not use any extra space beyond SA in the above step. For the time part, we only need to explain the running time of the comparison process. When we compare SA[i] and SA[i - 1], we can know the length of these two LMS-substrings (indicated by SA[i] and SA[i - 1]) from the Observation 2. Note that each character of T is scanned at most twice since it is only scanned when identifying its length and its adjacent predecessor LMS-substring. Thus the comparison process takes O(n) time because the total length of all LMS-substrings is less than 2n. □

3.6. Sort the LMS-suffixes by solving T₁ recursively

In this section, we sort all LMS-suffixes and place their indices in the tail of their corresponding buckets in SA. This can be done as follows:

1. We first solve T₁ recursively. From Section 3.5, T₁ is stored in SA[0...n₁ - 1]. We define SA₁ to be SA[n - n₁...n - 1] and use SA₁ to store the output of the subproblem T₁.
2. Now, we put all indices of LMS-suffixes in SA. First we move SA₁ to SA[0...n₁ - 1] (i.e., move SA[n - n₁...n - 1] to SA[0...n₁ - 1]). Then we scan T from right to left. For every LMS-character T[i], place i (i.e., index of suf(i)) in the tail of SA.
3. For notational convenience, we define LMS[0...n₁] \triangleq SA[n - n₁...n - 1]. Now, we obtain the sorted order of all LMS-suffixes of the original string T by letting SA[i] = LMS[SA[i]] for all i \in [0, n₁ - 1].
4. Finally, we scan SA[0...n₁ - 1] once more from right to left, and move the indices of LMS-suffixes in the same bucket to the tail of its bucket and clear other entries. This is easy to do since each S-type T[i] (after the renaming step in Section 3.2) has pointed to the tail of its bucket.

Example: Continue our example:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	1	1	2	0	E	E	E	E	E	E	E	E	E

Step 1. Solve T₁ recursively:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	1	1	2	0	E	E	E	E	E	<u>3</u>	<u>0</u>	<u>1</u>	<u>2</u>

Step 2. After move SA₁ to SA[0...n₁ - 1] and put all LMS-suffixes in SA:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	<u>3</u>	<u>0</u>	<u>1</u>	<u>2</u>	E	E	E	E	E	<u>1</u>	<u>5</u>	<u>9</u>	<u>12</u>

Step 3. Get all sorted LMS-suffixes:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	<u>12</u>	<u>1</u>	<u>5</u>	<u>9</u>	E	E	E	E	E	1	5	9	12

Step 4. Move the indices of LMS-suffixes in the same bucket to the tail of its bucket and clear other entries:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(<u>12</u>)	(E)	(E	E	<u>1</u>	<u>5</u>	<u>9</u>)	(E	E)	(E	E	E	E)

Lemma 8. All LMS-suffixes can be sorted by solving the reduced problem T₁ recursively and placed in the tail of their corresponding buckets in SA using O(n) time and O(1) workspace.

Proof. The time and space used in this step are easy to verify. We only show the correctness of this step. Each character of T₁ corresponds to an LMS-substring of T and this character is the rank of the corresponding sorted LMS-substring. Hence, the lexicographical order of LMS-suffixes of T is the same as the order of suffixes in T₁. □

3.7. Induced sort all suffixes of T

Now, we sort all suffixes of T from the sorted LMS-suffixes using induced sorting with our interior counter trick (Note that this step is the same as what we did in Section 3.4. Now, we describe the details here). First, we induced-sort the order of all L-suffixes from LMS-suffixes. Then we induce the order of S-suffixes from the L-suffixes. Now, we show how to carry out these steps with the desired optimal time and space.

Step 1. Induced sort all L-suffixes from the sorted LMS-suffixes: Some details of this step is similar to our previous step in Section 3.3 where we introduced our interior counter trick. We divide this step into two parts as follows:

- (1) First initialize SA: We scan T from right to left. For every T[i] which is L-type, do the following:

- (i) If $SA[T[i]] = \text{Empty}$, let $SA[T[i]] = \text{Unique}$ (unique L-type character in this bucket).
 - (ii) If $SA[T[i]] = \text{Unique}$, let $SA[T[i]] = \text{Multi}$ (the number of L-type characters in this bucket is at least 2).
 - (iii) Otherwise do nothing.
- (2) Then we scan SA from left to right to sort all the L-suffixes.
- (i) If $SA[i] = \text{Empty}$, do nothing.
 - (ii) If $SA[i]$ is an index, we let $j = SA[i] - 1$. Then, if $\text{suf}(j)$ is L-suffix (this can be identified in constant time from the following Lemma 9), we place $\text{suf}(j)$ into the LF-entry (recall that LF-entry denotes the leftmost free entry in its bucket) of its bucket and increase the counter by one.
 - (iii) If $SA[i] = \text{Multi}$, which means $SA[i]$ is the head of its bucket, and this bucket has at least two L-suffixes which are not sorted, we use $SA[i]$ and $SA[i + 1]$ as the bucket head (the symbol Multi) and the counter of this bucket, respectively. Then we skip these two entries and continue to scan $SA[i + 2]$.

Now, all L-suffixes have been sorted. Note that we still need to scan SA once more to free these positions occupied by Multi and counters. After this, the indices of all L-suffixes are in their final positions in SA.

Step 2. Remove LMS-suffixes from SA: We can use a trick similar to the previous Step 2 in Section 3.3, i.e., placing the indices of LMS-characters into SA. The difference is that instead of placing the actual LMS-characters, we place the Empty symbol instead. Also note that we do not delete the sentinel since it must be in the final position. Now, SA contains only all L-suffixes and the sentinel, and all of them are in their final positions in SA.

Step 3. Induced sort all S-suffixes from the sorted L-suffixes: Now, this step is completely symmetrical to the above Step 1 (Sort all L-suffixes using induced sorting). We use S-type and RF-entry instead of L-type and LF-entry, and we do not repeat the details here.

Example: Continue our example:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	7	6	6	9	9	6	6	9	9	6	7	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(12)	(E)	(E E	1 5	9)	(E E)	(E E	E E	E E	E E	E E	E E	E E

Step 1. Induced-sort all L-suffixes from the sorted LMS-suffixes:

(1) After initialization:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(U)	(E E	1 5	9)	(M E)	(M E)	(M E	E E	E E	E E	E E	E E

(2) Scan SA from left to right to sort all L-suffixes:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(11)	(E E	1 5	9)	(M E)	(M E)	(M E	E E	E E	E E	E E	E E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(M E	E E	E E	E E	E E	E E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(M	1	4	E	E	E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(M	2	4	8	E	E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(M	2	4	8	E	E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(4	8	3	E	E	E
SA	(12)	(11)	(E E	1 5	9)	(10	0)	(4	8	3	7)	E	E

Note that the third to last line is the case (iii). So we skip these two entries (i.e., 'M' and '2').

Step 2. Remove LMS-Suffixes from SA:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(11)	(E E	E E	E)	(10	0)	(4	8	3	7)	E	E

Step 3. Induced-sort all S-suffixes from the sorted L-suffixes:

(1) After initialization:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(11)	(E E	E E	M)	(10	0)	(4	8	3	7)	E	E

(2) Scan SA from right to left to sort all S-suffixes:

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
SA	(12)	(11)	(E	E	<u>6</u>	<u>1</u>	<u>M</u>	(10	0)	(4	8	3	<u>7</u>)
SA	(12)	(11)	(E	<u>2</u>	6	<u>2</u>	<u>M</u>	(10	0)	(4	8	<u>3</u>	7)
SA	(12)	(11)	<u>9</u>	2	6	<u>3</u>	<u>M</u>	<u>(10</u>	0)	(4	8	3	7)
SA	(12)	(11)	(9	2	<u>6</u>	3	M)	(10	0)	(4	8	3	7)
SA	(12)	(11)	<u>(E</u>	<u>5</u>	<u>9</u>	<u>2</u>	<u>6</u>)	(10	0)	(4	8	3	7)
SA	(12)	(11)	<u>(1</u>	5	9	<u>2</u>	6)	(10	0)	(4	8	3	7)
SA	(12)	(11)	(1	5	9	2	6)	(10	0)	(4	8	3	7)

In order to show the time used in this step, we need the following useful lemma to identify the L/S-type information in the induced-sorting step.

Lemma 9. When we scan SA from left to right to induced-sort L-suffixes, for each scanned SA[i], we want to identify the type of $\text{suf}(\text{SA}[i] - 1)$. If $T[\text{SA}[i] - 1] \neq T[\text{SA}[i]]$, the type of $\text{suf}(\text{SA}[i] - 1)$ can be obtained immediately. Otherwise $T[\text{SA}[i] - 1] = T[\text{SA}[i]]$ (this case $\text{suf}(\text{SA}[i] - 1)$ belongs to the current scanned bucket $T[\text{SA}[i]]$), the type of $\text{suf}(\text{SA}[i] - 1)$ is indicated by the type of the current bucket $T[\text{SA}[i]]$.

Proof. The type of $\text{suf}(\text{SA}[i] - 1)$ in the first case $T[\text{SA}[i] - 1] \neq T[\text{SA}[i]]$ directly follows from the definition of L/S-suffix in Section 2. For the second case $T[\text{SA}[i] - 1] = T[\text{SA}[i]]$, we know that the suffix $\text{suf}(\text{SA}[i] - 1)$ we are currently considering belongs to the current scanned bucket $T[\text{SA}[i]]$. Moreover, we know that all suffixes belonging to the same bucket have the same type after the renaming step in Section 3.2. Thus it is sufficient to identify the type of the current bucket $T[\text{SA}[i]]$. According to Observation 1, we can identify the type whenever we reach a new bucket in SA, e.g., check the bucket head h to see whether $T[\text{SA}[h]] > T[\text{SA}[h] + 1]$, if yes, then $T[\text{SA}[h]]$ is L-type and thus all suffixes belonging to this bucket is also L-type. \square

Lemma 10. Given all sorted LMS-suffixes of T , all suffixes can be sorted correctly using $O(n)$ time and $O(1)$ workspace according to the induced sorting steps.

Proof. For the correctness: we can sort all L-suffixes correctly from the sorted LMS-suffixes using induced sorting step from Lemma 2 and we can sort all S-suffixes correctly from the sorted L-suffixes using induced sorting step from Lemma 1. \square

Now, we obtain the following theorem for our optimal in-place algorithm.

Theorem 4. Our Algorithm takes $O(n)$ time and $O(1)$ workspace to compute the suffix array of string T over integer alphabets with $|\Sigma| \leq n$.

Proof. The time complexity simply follows from the recursion $T(n) = T(n/2) + O(n) = O(n)$. For the workspace, every step in our algorithm uses $O(1)$ workspace and different steps can reuse the $O(1)$ workspace too. In the recursive subproblem, we can also reuse the $O(1)$ workspace.⁵ \square

4. Suffix sorting for read-only integer alphabets

4.1. Framework

First, we define some notations. Let n_L and n_S denote the number of L-suffixes and S-suffixes, respectively. Let n_1 denote the length of the reduced problem T_1 , i.e., n_1 equals to the number of LMS-suffixes (Case 1) or LML-suffixes (Case 2). Now, we describe the framework of our algorithm as follows:

1. If $n_L \leq n_S$ (i.e., the number of L-suffixes is not larger than that of S-suffixes), then:

(1) (Section 4.2) Sort all LMS-characters of T .

We use *counting sort* to sort all LMS-characters of T in $\text{SA}[n - n_1 \dots n - 1]$. In the counting sort step, we use $\text{SA}[0 \dots n/2 - 1]$ as the temporary space (counting array). After this step, all indices of the sorted LMS-characters are stored in $\text{SA}[n - n_1 \dots n - 1]$. Note that this step is different from our previous algorithm in Section 3 since we sort and store the LMS-characters in the end of SA instead of their corresponding buckets, because T is read-only now.

(2) (Section 4.4) Induced sort all LMS-substrings from the sorted LMS-characters.

This induced-sorting step is the same as Step (4) below where we induced-sort all suffixes from the sorted LMS-suffixes. Thus, we only describe the details of this step in Section 4.4. After this step, all indices of the sorted LMS-substrings are stored in $\text{SA}[n - n_1 \dots n - 1]$.

⁵ If one worries the $O(\log n)$ workspace in the recursion, one can use the highest bits in SA (i.e., n bits) to store them since the size of the reduced sub-problem is not larger than $n/2$.

- (3) (Section 4.3) Construct and solve the reduced problem T_1 from the sorted LMS-substrings.

We construct the reduced problem T_1 using the ranks of all sorted LMS-substrings which are stored in $\text{SA}[n - n_1 \dots n - 1]$, where the ranks of LMS-substrings correspond to the lexicographical order of the sorted LMS-substrings (this construction step is the same as that in Section 3.5). Then we get the reduced problem T_1 in $\text{SA}[0 \dots n_1 - 1]$ and solve T_1 recursively to obtain the sorted LMS-suffixes. In the recursive step, we use $\text{SA}_1 = \text{SA}[n - n_1 \dots n - 1]$ as the output space for T_1 . After this step, all indices of the sorted LMS-suffixes are stored in $\text{SA}[n - n_1 \dots n - 1]$.

- (4) (Section 4.4) Induced sort all suffixes of T from the sorted LMS-suffixes (T_1).

We induced-sort all suffixes of T from the sorted LMS-suffixes which are stored in $\text{SA}[n - n_1 \dots n - 1]$. In the sorting step, we use the *interior counter trick*, as in Section 3, which can implicitly represent the dynamic LF/RF-entry information in SA. Besides, we provide a *pointer data structure* which can represent the bucket heads/tails in SA. Combining these two techniques, we can remove the workspace needed by the bucket array. Due to the space required by our pointer data structure is nonconstant, we divide this induced sorting step into two stages to address this issue. Note that this step is the main technical part of our optimal in-place algorithm. After this step, all indices of the suffixes of T are sorted and stored in $\text{SA}[0 \dots n - 1]$.

2. Otherwise, execute the above steps switching the role of LMS with LML.

Without loss of generality, we assume that $n_L \leq n_S$. Note that we compare the number of L-suffixes and S-suffix at the beginning since we need half of the space of SA to construct our pointer data structure for induced-sorting the L-suffixes (from the sorted LMS-suffixes) and S-suffixes (from the sorted L-suffixes) in Step (4). The in-place implementation of this induced sorting step is the main technical part of our optimal in-place algorithm. Note that the unused space is enough since the number of LMS-suffixes (i.e., n_1) and L-suffixes (i.e., n_L) both are less than or equal to $n/2$, where $n_1 \leq n/2$ since any two LMS-characters are not adjacent by the definition of LMS-characters, and $n_L \leq n/2$ since $n_L \leq n_S$. Note that for previous algorithms (e.g., [18,20]), they do not need the comparison at the beginning since they use the bucket array (which needs $|\Sigma|$ words workspace) in the induced sorting step. Here, we construct the pointer data structure and combine our interior counter trick to reduce the workspace.

Now, we describe the details of our in-place algorithm in the following sections.

4.2. Sort all LMS-characters of T

In this section, we sort all LMS-characters of T and place their indices in $\text{SA}[n - n_1 \dots n - 1]$. Recall that n_1 denotes the number of LMS-characters.

Now, we describe the details. Since $|\Sigma| = O(n)$, we can assume that $|\Sigma| \leq dn$ for some constant d . We divide the LMS-characters of T into $2d$ partitions and sort each partition one by one. The partition i contains the LMS-characters which belong to $\left[\frac{i|\Sigma|}{2d}, \frac{(i+1)|\Sigma|}{2d} - 1 \right]$, for $0 \leq i < 2d$. We use m_i to denote the number of LMS-characters in partition i . Then for each partition i , we use the standard *counting sort* (see e.g., [37, Chap. 8]) to sort these m_i LMS-characters (the LMS-characters can be identified by scanning T once from right to left). Concretely, we use $\text{SA}[0 \dots n/2 - 1]$ as the temporary counting array, and use $\text{SA}[n/2 + \sum_{j=0}^{i-1} m_j + 1 \dots n/2 + \sum_{j=0}^i m_j]$ as the output array. After this counting sort step, the indices of these m_i sorted LMS-characters have been placed in $\text{SA}[n/2 + \sum_{j=0}^{i-1} m_j + 1 \dots n/2 + \sum_{j=0}^i m_j]$.

Note that we can use the counting sort step for each partition. Because the gap of each partition is $\frac{|\Sigma|}{2d} \leq \frac{dn}{2d} = \frac{n}{2}$, the space of $\text{SA}[0 \dots n/2 - 1]$ is enough for the temporary counting array (its size equals to the gap) of counting sort step. It is not hard to see that the sorting step takes $O(n)$ time and uses $O(1)$ workspace since we only make $2d$ times of counting sort steps (each step takes linear time).

After sorting all $2d$ partitions, all indices of the sorted LMS-characters are placed in $\text{SA}[n/2 + 1, n/2 + \sum_{j=0}^{2d-1} m_j]$ (i.e., $\text{SA}[n/2 + 1, n/2 + n_1]$). Then we move them to $\text{SA}[n - n_1 \dots n - 1]$, which can be easily done in linear time and $O(1)$ workspace.

4.3. Construct and solve the reduced problem T_1 from sorted LMS-substrings

Construct the reduced problem T_1 : We construct the reduced problem T_1 using the ranks of all sorted LMS-substrings which are stored in $\text{SA}[n - n_1 \dots n - 1]$ from the Step (2) (see the framework in Section 4.1), where the ranks of LMS-substrings are corresponding to the lexicographical order of the sorted LMS-substrings.

Note that this construction step is exactly the same as that in Section 3.5. Thus we omit the details and recall the same lemma as follows.

Lemma 11. T_1 can be constructed using $O(n)$ time and $O(1)$ workspace.

Solve T_1 recursively: Now, we sort all LMS-suffixes by solving T_1 recursively and place their indices in the tail of SA (i.e., $\text{SA}[n - n_1 \dots n - 1]$). This step is carried out as follows (similar to the Section 3.6 except Step 4 below):

1. We first solve T_1 recursively. Recall that T_1 is stored in $\text{SA}[0 \dots n_1 - 1]$. We define SA_1 to be $\text{SA}[n - n_1 \dots n - 1]$ and use SA_1 to store the output of the subproblem T_1 .
2. Now, we put all indices of LMS-suffixes in SA. First we move SA_1 to $\text{SA}[0 \dots n_1 - 1]$ (i.e., move $\text{SA}[n - n_1 \dots n - 1]$ to $\text{SA}[0 \dots n_1 - 1]$). Then we scan T from right to left. For every LMS-character $T[i]$, place i (i.e., index of $\text{suf}(i)$) in the tail of SA.
3. For notational convenience, we define $\text{LMS}[0 \dots n_1] \triangleq \text{SA}[n - n_1 \dots n - 1]$. Now, we obtain the sorted order of all LMS-suffixes of the original string T by letting $\text{SA}[i] = \text{LMS}[\text{SA}[i]]$ for all $i \in [0, n_1 - 1]$.
4. Finally, we finish this step by moving $\text{SA}[0 \dots n_1 - 1]$ to $\text{SA}[n - n_1 \dots n - 1]$. Now, all indices of the sorted LMS-suffixes are stored in $\text{SA}[n - n_1 \dots n - 1]$.

Lemma 12. All LMS-suffixes can be sorted by solving the reduced problem T_1 recursively and placed in the tail of SA using $O(n)$ time and $O(1)$ workspace.

Proof. The time and space used in this step are easy to verify. We only show the correctness of this step. Each character of T_1 corresponds to an LMS-substring of T and this character is the rank of the corresponding sorted LMS-substring. Hence, the lexicographical order of LMS-suffixes of T is the same as the order of suffixes in T_1 . \square

4.4. Induced sort all suffixes of T from the sorted LMS-suffixes

In this section, we show how to induced-sort all suffixes from the sorted LMS-suffixes. All indices of the sorted LMS-suffixes have been placed in $\text{SA}[n - n_1 \dots n - 1]$ from the previous step (see Lemma 12). Note that this step is the main technical part of our in-place algorithm. We develop two techniques (*interior counter trick* and *pointer data structure*) for implementing this induced sorting in-place.

Let $\text{SA}_L = \text{SA}[0 \dots n_L - 1]$ and $\text{SA}_S = \text{SA}[n_L \dots n - 1]$. Recall that n_S and n_L denote the number of S-suffixes and L-suffixes, respectively. Also note that $n_L + n_S = n$. First, we sort all n_L L-suffixes from the sorted LMS-suffixes which are stored in $\text{SA}[n - n_1 \dots n - 1]$ and store the sorted L-suffixes in SA_L . Then, we sort all n_S S-suffixes from the sorted L-suffixes and store the sorted S-suffixes in SA_S . Finally, we merge the sorted L-suffixes (stored in SA_L) and S-suffixes (stored in SA_S) to sort all suffixes in $\text{SA}[0 \dots n - 1]$.

Note that sorting the n_L L-suffixes from the sorted LMS-suffixes is totally symmetrical as sorting the n_S S-suffixes from the sorted L-suffixes, as stated in Section 2. Thus, we only need to show the details of how to sort all n_S S-suffixes from the sorted L-suffixes which has already been stored in SA_L , and store the sorted S-suffixes in SA_S . Before we show the details, we briefly recall the original induced sorting step here (we have introduced it in Section 2 and provided a running example in Appendix A.1.). Now, we recall the symmetrical case here, i.e., inducing the order of S-suffixes from the sorted L-suffixes.

Inducing the order of S-suffixes from the sorted L-suffixes: We scan SA from right to left (i.e., from $\text{SA}[n - 1]$ to $\text{SA}[0]$). When we scan $\text{SA}[i]$, let $j = \text{SA}[i] - 1$. If $T[j]$ is S-type, i.e., $\text{suf}(j)$ is an S-suffix (indicated by the type array), we place the index of $\text{suf}(j)$ (i.e. j) into the RF-entry of bucket $T[j]$, and then let the RF-pointer of this bucket $T[j]$ point to the next free entry. If $T[j]$ is L-type, we do nothing (since all L-suffixes are sorted in the correct positions). The RF-pointers are maintained by the bucket array.

In order to obtain the in-place algorithm, we need to show how to remove the workspace needed by the bucket array and type array in the induced sorting step. Briefly speaking, the purpose of the pointer data structure is to indicate the bucket tails of S-suffixes, and the purpose of the interior counter trick is to maintain the RF-pointers of the buckets dynamically. Thus, for a query of RF-entry for $\text{suf}(j)$ in bucket $T[j]$, we know the tail of the bucket $T[j]$ from the pointer data structure in constant time (Lemma 16), then we use the interior counter trick to indicate the RF-entry in this bucket (Lemma 13). Instead of using the type array, we use the Lemma 14 to identify the L/S-suffixes in the induced sorting step.

Now, we describe the details. First, we introduce our interior counter trick, extending the one from Section 3. Here, we assume that the tail of the bucket of any S-suffix is known (which is indicated by the Lemma 16).

Interior counter trick: Note that the buckets of the S-suffixes we discussed in this section are in $\text{SA}_S = \text{SA}[n_L \dots n - 1]$, since we already have placed the sorted L-suffixes in $\text{SA}_L = \text{SA}[0 \dots n_L - 1]$. Thus, we only need to sort all S-suffixes to their corresponding buckets in SA_S and the buckets only contains S-suffixes now.

Here we only describe the details of interior counter trick for one bucket since other buckets are the same. Note that we assume that the tail of the bucket of any S-suffix is known (Lemma 16). To simplify the representation, we assume the bucket from index 0 to index $m - 1$ of SA_S , where m is the size of this bucket (i.e. the number of S-suffixes in this bucket is m). We only describe the case where $m > 3$ since other cases with $m \leq 3$ are similar and simpler. We define five special symbols B_H (head of the bucket), B_T (tail of the bucket), E (Empty), R_1 (one remaining S-suffix) and R_2 (two remaining S-suffixes).⁶

First, we use three special symbols to initialize this bucket, i.e., let $\text{SA}_S[0] = B_H$, $\text{SA}_S[m - 2] = E$ and $\text{SA}_S[m - 1] = B_T$. Let S_i denote the index of the i -th S-suffix which needs to be placed into the RF-entry of this bucket. Now, we describe

⁶ Recall that the special symbol is only used to simplify the argument. See Appendix B for the details.

how to place the indices of these m S-suffixes into the RF-entry of this bucket one by one. We distinguish the following four cases:

- (1) If $SA_S[m - 1] = B_T$, and $SA_S[m - 2] = E$ or $SA_S[m - SA_S[m - 2] - 3] \neq B_H$: In this case, we place the index of the current S-suffix (i.e., S_i) into the RF-entry of this bucket, where $1 \leq i \leq m - 3$. Concretely, we know the position of the tail of this bucket in SA_S , i.e., $m - 1$ according to the assumption. Then, we use $SA_S[m - 2]$ as the counter to denote the number of the indices of S-suffixes has been placed so far. Note that the RF-entry of this bucket is pointed by this counter (i.e. RF-pointer). Thus, we can place the index of the current S-suffix (S_i) into the RF-entry of this bucket in constant time, and then update the counter $SA_S[m - 2]$.
- (2) If $SA_S[m - 1] = B_T$ and $SA_S[m - SA_S[m - 2] - 3] = B_H$: In this case, we place the index of the third to last S-suffix (i.e. S_{m-2}) into the RF-entry of this bucket. Concretely, we shift the previous $m - 3$ S-suffixes which stored in $SA_S[1 \dots m - 3]$ to $SA_S[2 \dots m - 2]$. Then, we place S_{m-2} into $SA_S[1]$ and let $SA_S[m - 1] = R_2$. This step takes $O(m)$ time since we shift $m - 3$ S-suffixes.
- (3) If $SA_S[m - 1] = R_2$: In this case, we place the index of the second to last S-suffix (i.e. S_{m-1}) into the RF-entry of this bucket. We shift the previous $m - 2$ S-suffixes which stored in $SA_S[1 \dots m - 2]$ to $SA_S[2 \dots m - 1]$. Then, we place S_{m-1} into $SA_S[1]$ and let $SA_S[0] = R_1$. This step takes $O(m)$ time since we shift $m - 2$ S-suffixes.
- (4) Otherwise: In this case, we place the index of the last S-suffix (i.e. S_m) into the RF-entry of this bucket. First, we know the tail of the bucket indicated by our pointer data structure in constant time. Then, we search the entries before the tail one by one until that we find the special symbol R_1 . We let this entry to be S_m . This step takes $O(m)$ time since we search $m - 1$ S-suffixes.

In order to demonstrate these four cases more clearly, we also provide a demonstration as follows:

Index	0	1	...	$m - 3$	$m - 2$	$m - 1$
SA_S	$SA_S[0]$	$SA_S[1]$...	$SA_S[m - 3]$	$SA_S[m - 2]$	$SA_S[m - 1]$
After initialization :						
SA_S	<u>B_H</u>	$SA_S[1]$...	$SA_S[m - 3]$	<u>E</u>	<u>B_T</u>
Case (1) :						
SA_S	B_H	$SA_S[1]$...	<u>S_1</u>	<u>1</u>	B_T
\vdots					\vdots	
SA_S	B_H	<u>S_{m-3}</u>	...	S_1	<u>$m - 3$</u>	B_T
Case (2) :						
SA_S	B_H	<u>S_{m-2}</u>	...	<u>S_2</u>	<u>S_1</u>	<u>R_2</u>
Case (3) :						
SA_S	<u>R_1</u>	<u>S_{m-1}</u>	...	<u>S_3</u>	<u>S_2</u>	<u>S_1</u>
Case (4) :						
SA_S	<u>S_m</u>	S_{m-1}	...	S_3	S_2	S_1

Note that this step uses $O(1)$ workspace since there is no bucket array and type array, and the space needed by our interior counter trick and pointer data structure is in SA_S . The purpose of the interior counter trick is to dynamic maintain the RF-pointers of the buckets. E.g., for a query of RF-entry for $suf(j)$ in bucket $T[j]$, first we know the tail of the bucket $T[j]$ by the assumption, then we use the interior counter trick to indicate the RF-entry in this bucket. We have the following lemma.

Lemma 13. *If the tail of the bucket of any S-suffix is known, one can sort the S-suffixes from the sorted L-suffixes using the induced sorting step with the interior counter trick in linear time and $O(1)$ workspace.*

Note that in the induced sorting step, one uses the type array to identify whether the $suf(j)$ is S-suffix or not. Instead of using the type array, we use the following Lemma 14 to identify the type of the L- or S-suffix in the induced-sorting step.

Lemma 14. *If $T[j] \neq T[SA[i]]$, the type of $suf(j)$ can be obtained immediately, where $j = SA[i] - 1$. Otherwise $T[j] = T[SA[i]]$, if all S-suffixes of T that belong to bucket $T[SA[i]]$ are not already sorted, then the $suf(j)$ is S-suffix.*

Proof. During the induced-sorting step, when we scan $SA[i]$, let $j = SA[i] - 1$ and we need to identify whether $suf(j)$ is an S-suffix or not. First, if $T[j] \neq T[SA[i]]$, then the type of $suf(j)$ can be obtained immediately. Otherwise $T[j] = T[SA[i]]$ (this case $suf(j)$ belongs to the current scanning bucket $T[SA[i]]$), clearly if all S-suffixes of T that belong to bucket $T[SA[i]]$ are not already sorted, then the $suf(j)$ is S-suffix. It is not hard to decide whether all S-suffixes in the current scanning bucket $T[SA[i]]$ are already sorted or not. One can use an extra variable to denote how many S-suffixes remain in the current scanning bucket $T[SA[i]]$. When we begin to scan a new bucket, we scan this bucket once to initialize this variable. Note

that we can do this initialization, since there are special symbols in SA_S for each bucket after we initialize the buckets in our interior counter trick. \square

Now, there is only one thing left: how to know the tails of the bucket of S-suffixes in the induced sorting step. The purpose of the pointer data structure is to indicate the tails of the bucket of S-suffixes. However, the pointer data structure requires c_p words, where the value of c_p will be specified later. Thus, we need to divide this induced-sorting step into two stages. The first stage (Section 4.4.1) is to sort the first $n_S - c_p$ S-suffixes (i.e. the largest $n_S - c_p$ S-suffixes), where our pointer data structure still exists. The second stage (Section 4.4.2) is to sort the last c_p S-suffixes, where there is no space for the pointer data structure.

4.4.1. The first stage

Pointer data structure: Now, we construct our pointer data structure which supports to find the tails of the buckets in constant time. We store the pointer data structure in the tail of SA_S , recall that $SA_S = SA[n_L \dots n - 1]$. Now, we describe the details. We divide the S-suffixes of T into $4d$ parts according to their first characters, and construct the pointer data structure for each part respectively. The $4d$ parts are divided by $T[j] \in \left[\frac{i\Sigma}{4d} + 1, \frac{(i+1)\Sigma}{4d} \right]$, for $0 \leq i < 4d$. Let D_i denote the pointer data structure of the i -th part. We only show the details how we construct the pointer data structure D_0 as follows, since constructing D_i is similar for $0 < i < 4d$ (one only need to shift $T[j]$ with $\frac{i\Sigma}{4d}$).

- (1) First, we let $SA_S[i] = 1$ for all $i \in [1, \frac{\Sigma}{4d}]$. Then we scan T from right to left. For every S-type $T[i] \in [1, \frac{\Sigma}{4d}]$, we increase $SA_S[T[i]]$ by one.
- (2) Then we scan $SA_S[1 \dots \frac{\Sigma}{4d}]$ from left to right. We use a variable sum to count the sum, first initialize $sum = -1$. For each $SA_S[i]$ which is being scanned, first let $sum = sum + SA_S[i]$, then let $SA_S[i] = sum$. Now, for any S-suffix $suf(i)$ satisfying $T[i] \in [1, \frac{\Sigma}{4d}]$, $SA_S[T[i]] - T[i]$ must indicate the tail of bucket $T[i]$ in SA_S . Since we want every entry in $SA_S[1 \dots \frac{\Sigma}{4d}]$ to be distinct, we initialize $SA_S[i] = 1$ for all $i \in [1, \frac{\Sigma}{4d}]$ in Step (1). Hence the tail of bucket $T[i]$ is $SA_S[T[i]] - T[i]$.
- (3) Finally, we construct D_0 for $SA_S[1 \dots \frac{\Sigma}{4d}]$ according to Lemma 15. D_0 uses at most $c(n + \frac{\Sigma}{4d})/\log n$ words space. We store D_0 in the tail of SA_S (i.e., $SA_S[n_S - c(n + \frac{\Sigma}{4d})/\log n \dots n_S - 1]$). D_0 supports to find the tail of the bucket of any S-suffix $suf(i)$ satisfying $T[i] \in [1, \frac{\Sigma}{4d}]$ in constant time.

Lemma 15. For a strictly increasing sequence of m non-negative integers $0 \leq a_0 < a_1 \dots < a_{m-1} \leq n$ such that $m \leq n$ and $n > 1024$, one can construct a data structure using linear time (i.e., $O(n)$ time) and at most $cn/\log n$ words for $1 < c < 2$ such that each query to the i -th smallest integer a_i can be answered in constant time.

Proof. We first construct a bit array $B[0 \dots n]$ such that $B[i] = 1$ if index i is an element in the increasing sequence; $B[i] = 0$ otherwise. Using the bitmap B , we convert the query of the i -th smallest integer a_i on the increasing sequence into the $select_1(i)$ query of the index of the i -th 1 on B . Then we can use the methods [38,39] to answer any $select_1(i)$ query in constant time using an extra data structure of $O(n/\log \log n)$ bits space with $O(n)$ construction time. Particularly, the extra data structure can be represented in $3n/\log \log n + n^{\frac{1}{4}}(\frac{1}{4} \log n \log \log n + \log \log n)$ bits [39]. Converting bits to words, we can see that the data structure uses at most $cn/\log n$ words for $1 < c < 2$ when $n > 1024$. \square

After this step, the pointer data structure (i.e. D_i for all $0 \leq i < 4d$) is stored in $SA_S[n_S - c_p \dots n_S - 1]$, where $c_p = \lceil 4d \cdot (c(n + \frac{\Sigma}{4d})/\log n) \rceil \leq \lceil 5dcn/\log n \rceil$. Recall that the sorted n_L L-suffixes are stored in SA_L and $n_S \geq n_L$. Thus the unused space (i.e., $SA_S = SA[n_L \dots n - 1]$) is at least half of the space of SA . It is enough to construct and store the pointer data structure, i.e., $\frac{\Sigma}{4d} \leq \frac{dn}{4d} = \frac{n}{4}$ for constructing it and $\frac{n}{4}$ for storing it (which uses c_p words). Note that we assume that $n > 4c_p$, otherwise it is easy to solve since n is constant. Hence the pointer data structure is constructed in linear time and supports to find the tail of the bucket of any S-suffix in constant time.

Lemma 16. We can construct the pointer data structure in linear time, and this pointer data structure uses at most c_p words and can support to find the bucket tail of any S-suffix in constant time.

Proof. We only need to specify the query time constant. We use $4d$ values, i.e., m_1, \dots, m_{4d} , which denote the number of S-suffixes in each interval, respectively (They can be obtained from the variable sum which is computed in the final stage of the Step (2)). Now, if we want to find the bucket tail of an S-suffix $suf(i)$, we first compare $T[i]$ with $\frac{i\Sigma}{4d}$ (for $0 \leq i < 4d$) to see which pointer data structure $T[i]$ belongs to. Assume that it belongs to D_j . Then we do a $select(T[i] - (j - 1)\frac{\Sigma}{4d})$ query on D_j , and combine the $select$ result with the corresponding m_k ($k < j$) to identify the tail of bucket $T[i]$. All the above operations can be done in constant time. \square

Now according to Lemma 13, 14 and 16, we can sort the first largest $n_S - c_p$ S-suffixes from the sorted L-suffixes which stored in SA_L using the induced sorting step with the interior counter trick and the pointer data structure which stored in $SA[n - c_p \dots n - 1]$, and we store the indices of the sorted $n_S - c_p$ S-suffixes in $SA_S[0 \dots n_S - c_p - 1]$.

4.4.2. The second stage

Now, we only need to show how to sort the last c_p S-suffixes which is occupied by our pointer data structure. First, we specify how to identify whether an S-suffix belongs to the first $n_S - c_p$ S-suffixes or not. We can scan T once to find the character of $n_S - c_p$ largest S-suffix using the pointer data structure and let ch denote this character. If the tail of the bucket ch is exactly $c_p - 1$ in SA_S , then we compare the beginning character of the S-suffix with ch to identify which part it belongs to. Otherwise, the $n_S - c_p$ largest S-suffix belongs to bucket ch , and it will be stored in $SA_S[c_p - 1]$. We only need two variables to indicate whether the S-suffix belongs to the first $n_S - c_p$ S-suffixes or not. One is the number to denote the tail of bucket ch , and the other is also an integer number which denotes the gap between the tail of bucket ch and $c_p - 1$.

Now, we describe the details to sort the last c_p S-suffixes. First, we move these largest $n_S - c_p$ S-suffixes to the tail of SA_S , i.e., $SA_S[c_p, n_S - 1]$. Then we scan the T from right to left to place the smallest c_p S-suffixes into $SA_S[0, c_p - 1]$. Now, we use merge sort with the in-place linear time merging algorithm [40] to sort these c_p S-suffixes, the sorting key for each S-suffix is its beginning character. After this sorting step, these c_p S-suffixes have been placed in their corresponding buckets in $SA_S[0, c_p - 1]$. Note that we can use the same sorting step (which we used for sorting the first $n_S - c_p$ S-suffixes) to sort the last c_p S-suffixes without the pointer data structure. Now, the key point is that we can use the binary search (instead of the pointer data structure) to find the tails of the bucket for these c_p S-suffixes, since $c_p \leq \lceil 5dcn / \log n \rceil$ is small enough (i.e. $c_p \log n = O(n)$) to maintain that the time complexity of our algorithm is $O(n)$.

Using the binary search to extend interior counter trick is not very difficult, one can see the details in Section 5.5 which we induced sort all L-suffixes from the sorted S-suffixes for the general alphabets (Note that the optimal time is $O(n \log n)$ for the general alphabets case. Thus, we directly use the binary search to extend interior counter trick and do not use the pointer data structure in that case).

After this step, all n_S S-suffixes are sorted in SA_S . Now we have all sorted L-suffixes in SA_L (i.e., $SA[0 \dots n_L - 1]$) and all sorted S-suffixes in SA_S (i.e., $SA[n_L \dots n - 1]$). We use the stable, in-place, linear time merging algorithm [40] to merge the ordered SA_L and SA_S (the merging key for $SA[i]$ is $T[SA[i]]$, i.e., the first character of $\text{suf}(SA[i])$). After this merging step, all suffixes of T have been sorted in $SA[0 \dots n - 1]$.

Finally, we obtain the following theorem for our optimal in-place algorithm.

Theorem 5 (Main Theorem). *Our Algorithm takes $O(n)$ time and $O(1)$ workspace to compute the suffix array of string T over integer alphabets Σ , where T is read-only and $|\Sigma| = O(n)$.*

5. Suffix sorting for read-only general alphabets

5.1. Framework

The framework of our algorithm for read-only general alphabets (i.e., the only operations allowed on the characters of T (read-only) are comparisons) is described as follows:

1. If $n_S \leq n_L$ (i.e., the number of S-suffixes is not larger than that of L-suffixes), then
 - (1) (Section 5.2) Sort all S-substrings of T using mergesort directly.
We use mergesort to sort all S-substring of T in $SA[n - n_S \dots n - 1]$. In the merging step of mergesort, we use $SA[0 \dots n_S - 1]$ as the temporary space. After this step, all S-substrings should be in the lexicographical order stored in $SA[n - n_S \dots n - 1]$.
 - (2) (Section 5.3) Construct the reduced problem T_1 from the sorted S-substrings.
We construct the reduced problem T_1 using the ranks of all sorted S-substrings which are stored in $SA[n - n_S \dots n - 1]$. The ranks of S-substrings are corresponding to the lexicographical order of the sorted S-substrings. After this step, we get the reduced problem T_1 in $SA[0 \dots n_S - 1]$.
 - (3) (Section 5.4) Sort the S-suffixes by solving T_1 recursively.
We sort $T_1 = SA[0 \dots n_S - 1]$ recursively. In the recursive step, we use $SA_1 = SA[n - n_S \dots n - 1]$ as the output space for T_1 . Then we use the suffix array of T_1 (i.e. SA_1) to place all indices of the sorted S-suffixes of T into $SA[n - n_S \dots n - 1]$.
 - (4) (Section 5.5) Induced sort all suffixes from the sorted S-suffixes.
First, we place all indices of S-suffixes in their final positions in SA by using mergesort together with a stable, in-place, linear time merging algorithm [40]. Then we extend our *interior counter trick* to sort all L-suffixes from the sorted S-suffixes. Finally, all indices of the sorted suffixes of T are stored in $SA[0 \dots n - 1]$.
2. Otherwise, execute the above steps switching the roles of L and S .

Note that our algorithm is very simple since the optimal time complexity in this case is $O(n \log n)$ instead of $O(n)$. Moreover, our algorithm does not make any bit operations rather than previous algorithms, e.g., [17].

Without loss of generality, we assume that $n_S \leq n_L$. Now, we describe the details of our in-place algorithm in the following sections.

5.2. Sort all S-substrings of T

In this section, we sort all S-substrings of T as follows:

1. First, we scan T from right to left and place all indices of S-type characters into $SA[n - n_S \dots n - 1]$. Note that $n_S \leq n/2$ since we assume that $n_S \leq n_L$.
2. Then, we sort $SA[n - n_S \dots n - 1]$ using mergesort (the sorting key for $SA[i]$ is the S-substring of T which begins at $T[SA[i]]$). We use $SA[0 \dots n_S - 1]$ as the temporary space for mergesort. To compare two keys (i.e., two S-substrings) in mergesort, we simply do the straightforward character-wise comparisons.

After the above two steps, all the S-substring have been sorted in $SA[n - n_S \dots n - 1]$. We have the following lemma.

Lemma 17. *We can sort all S-substrings using $O(n \log n)$ time and $O(1)$ workspace.*

Proof. Step 1 does not need any extra space and costs linear time, because we can compute the type of each character in $O(1)$ time during the right-to-left scan of T . Moreover, we know mergesort needs linear workspace. Hence, it is sufficient to use $SA[0 \dots n_S - 1]$ as the workspace for mergesort. For Step 2, it suffices to show that the time spent for comparison process in any recursive level of mergesort (there are $O(\log n)$ recursive levels) can be bounded by $O(n)$. For the comparison process in the merging step of two sorted sequences, each comparison of two S-substrings can determine a smaller one and add it to the merged sequence, and the comparison time is bounded by the shorter one. The length of the S-substrings can be obtained according to Observation 2. Recall that each character of T is scanned at most twice since it only be scanned when identifying the length of its adjacent predecessor S-substring and itself. Thus the time complexity for each recursive level of mergesort can be bounded by $O(n)$ because the total length of all S-substrings is less than $2n$. In sum, all S-substrings can be sorted in $O(n \log n)$ time via mergesort. \square

5.3. Construct the reduced problem T_1 from the sorted S-substrings

In this section, we construct the reduced problem T_1 by renaming the sorted S-substrings. After Section 5.2, all S-substrings have been sorted in $SA[n - n_S \dots n - 1]$. The construction of T_1 consists of the following two steps:

1. We rename the S-substrings by their ranks. First let the rank of $SA[n - n_S]$ be 0. We scan $SA[n - n_S + 1 \dots n - 1]$ from left to right. When scanning $SA[i]$, we compare S-substring beginning with $T[SA[i]]$ and S-substring beginning with $T[SA[i - 1]]$. If they are different, let the rank of $SA[i]$ be the rank of $SA[i - 1]$ plus one. Otherwise, the rank of $SA[i]$ is the same as that of $SA[i - 1]$. We store the rank of $SA[i]$ in $SA[i - n + n_S]$.
2. Next, we use the heapsort to sort $SA[n - n_S \dots n - 1]$ (the sorting key for $SA[i]$ is $SA[i]$ itself). When we exchange two entries (say, $SA[i]$ and $SA[j]$, $i, j \in [n - n_S \dots n - 1]$) in $SA[n - n_S \dots n - 1]$ during heapsort, we also exchange the corresponding two entries (i.e., $SA[i - n + n_S]$ and $SA[j - n + n_S]$) in $SA[0 \dots n_S - 1]$. Note that we use heapsort here since it is in-place, so we do not need any extra space.

After the above two steps, we get the reduced problem T_1 in $SA[0 \dots n_S - 1]$.

Lemma 18. *T_1 can be constructed in $O(n \log n)$ time and $O(1)$ workspace.*

Proof. In Step 1, each S-substring beginning with $T[SA[i]]$ is compared with S-substring beginning with $T[SA[i + 1]]$. So each S-substring only participates in two comparisons. Now the argument is similar to a comparison process in a recursive level of mergesort in Lemma 17, thus it costs linear time. Obviously, Step 2 takes $O(n \log n)$ time and $O(1)$ workspace. \square

5.4. Sort the S-suffixes by solving T_1 recursively

In this section, we solve $T_1 = SA[0 \dots n_S - 1]$ recursively to obtain the order of all S-suffixes in $SA[n - n_S \dots n - 1]$. For the recursive step, we use $SA_1 = SA[n - n_S \dots n - 1]$ as the output space for T_1 . After the recursive call, SA_1 stores the suffix array of T_1 . We need to restore their names back to the indices of S-suffixes in T they represented. This step can be done as follows.

1. First, we scan T from right to left. We maintain a counter sum for the number of S-type characters we have scanned so far. Initially sum is 0. If $T[i]$ is S-type, we increase sum by 1 and place $suf(i)$ into $SA[n_S - sum]$ (i.e., let $SA[n_S - sum] \leftarrow i$). Now $SA[0 \dots n_S - 1]$ stores the indices of all S-suffixes of T .

2. Then for $i \in [n - n_S, n - 1]$, let $SA[i] \leftarrow SA[SA[i]]$.

Now, we have obtained all S-suffixes in the lexicographical order in $SA[n - n_S \dots n - 1]$.

5.5. Induced sort all suffixes of T

From Section 5.4, we have obtained the sorted S-suffixes in $SA[n - n_S \dots n - 1]$. Now, we sort all suffixes from these sorted S-suffixes.

Preprocessing: First, we scan T from right to left to place all indices of L-suffixes into $SA[0 \dots n - n_S - 1]$. Then, we sort $SA[0 \dots n - 1]$ (the sorting key of $SA[i]$ is $T[SA[i]]$ i.e., the first character of $\text{suf}(SA[i])$) using the mergesort, with the merging step implemented by the stable, in-place, linear time merging algorithm developed by Salowe and Steiger [40]. After this sorting step, we show some useful observations.

Observation 3. All suffixes of T have been sorted by their first characters in SA, i.e., in their corresponding buckets.

Observation 4. All indices of L-suffixes beginning with the same character in SA are in increasing order, due to the stability of the above sorting algorithm.

Lemma 19. All S-suffixes are already in their final positions in SA.

Proof. Before the sorting step, all sorted S-suffixes are in $SA[n - n_S \dots n - 1]$ and all L-suffixes are in $SA[0 \dots n - n_S - 1]$. Because the merging step is stable, the S-suffixes are behind the L-suffixes in the same bucket and hence are already in their final positions in SA from Observation 3 and Property 1. \square

Induced Sorting: Now, we induce the order of all L-suffixes from the sorted S-suffixes (which are already in their final position in SA by Lemma 19) using induced sorting. Now, we extend the interior counter trick in Section 4.4 to handle the read-only general alphabets. We use five special symbols B_H (Head of L-suffixes), B_T (Tail of L-suffixes), E (Empty), R_1 (one remaining L-suffix) and R_2 (two remaining L-suffixes).⁷ We do the following two steps to sort all L-suffixes:

Step 1. Initializing SA: Firstly, we initialize all buckets in SA by placing some special symbols in each bucket in order to inform us the number of L-suffixes in the bucket. Concretely, we scan T from right to left. For each scanning character $T[i]$ which is L-type, if bucket $T[i]$ has not been initialized, we need to initialize bucket $T[i]$ (we will show how to identify the bucket is initialized or not in the end of this step). Before to initialize bucket $T[i]$, we first need to obtain the value N_L , which is the number of L-suffixes in this bucket. Let l denote the head of bucket $T[i]$ in SA (i.e. l is the smallest index in SA such that $T[SA[l]] = T[i]$) and r denote the tail of bucket $T[i]$ in SA (i.e. r is the largest index in SA such that $T[SA[r]] = T[i]$). Furthermore, we let r_L denote the tail of L-suffixes in this bucket (i.e., r_L is the largest index in SA such that $T[SA[r_L]] = T[i]$ and $T[SA[r_L]]$ is L-type). Note that $N_L = r_L - l + 1$. Hence, it suffices to compute l and r_L . The following steps compute l and r_L , respectively.

- (i) We can find l by searching $T[i]$ in SA (the search key for $SA[i]$ is $T[SA[i]]$) using *binary search*. This uses $O(\log n)$ time from Observation 3.
- (ii) For r_L , since the bucket $T[i]$ has not been initialized, $\text{suf}(i)$ is the first L-suffix in its bucket being scanned. From Observation 4, $\text{suf}(i)$ must be stored in $SA[r_L]$ (i.e., $SA[r_L] = i$) since we scan T from right to left. Hence, we can scan this bucket from l to r to find r_L which satisfies $SA[r_L] = i$.

After this, we have obtained the value of N_L . Now, we initialize the bucket $T[i]$ as follows:

- (1) If $N_L = 1$, we do nothing (there is only one L-suffix in this bucket, and obviously it is in the final position).
- (2) If $N_L = 2$, let $SA[l + 1] = B_T$ (recall that l is the head of bucket $T[i]$ and r is the bucket tail, i.e., $SA[l \dots r]$ is the bucket $T[i]$. Moreover, r_L is the tail of L-suffixes in this bucket)

Index	l	$l + 1(r_L)$	$l + 2$	\dots	r
Type	L	L	S	\dots	S
SA	$SA[l]$	B_T	$SA[l + 2]$	\dots	$SA[r]$

⁷ Recall that the special symbol is only used to simplify the argument. See Appendix B for the details. The only difference is that we need to use the in-place linear time merging algorithm [40] to place all indices of the suffixes of T in their corresponding buckets in SA (the merging key is the beginning character of the suffix), then we scan the SA once to know the bucket heads/tails for the five integers (instead of scanning the string T for the (read-only) integer alphabets).

(3) If $N_L = 3$, let $SA[l + 1] = B_H$ and $SA[l + 2] = B_T$.

Index	l	$l + 1$	$l + 2(r_L)$	$l + 3$...	r
Type	L	L	L	S	...	S
SA	$SA[l]$	B_H	B_T	$SA[l + 3]$...	$SA[r]$

1. If $N_L > 3$, let $SA[l + 1] = B_H$, $SA[l + 2] = E$ and $SA[l + N_L - 1] = B_T$.

Index	l	$l + 1$	$l + 2$	$l + 3$...	$l + N_L - 1(r_L)$	$l + N_L$...	r
Type	L	L	L	L	...	L	S	...	S
SA	$SA[l]$	B_H	E	$SA[l + 3]$...	B_T	$SA[l + N_L]$...	$SA[r]$

Note that we can find out whether the bucket $T[i]$ is already initialized or not in $O(\log n)$ time. We do a binary search find l , then check $SA[l + 1]$ is B_H or B_T or others. If the bucket $T[i]$ has been initialized, $SA[l + 1]$ is B_H or B_T . Otherwise, it has not been initialized yet.⁸ It is not hard to see that this initialization step uses $O(n \log n)$ time and $O(1)$ workspace.

Step 2. Sort all L-suffixes using induced sorting: We scan SA from left to right to sort all L-suffixes. The step is similar to sorting all suffixes in Section 4.4. The main difference is that we use binary search to find the bucket head. Specifically, for every $SA[i]$, let $j = SA[i] - 1$. If $T[j]$ is L-type, then place $suf(j)$ into the LF-entry of its bucket, and increase the head counter by one. To specify how to place the L-suffix into the LF-entry of its bucket in SA, We only specify the case where $N_L > 3$ for the bucket. The other cases with $N_L \leq 3$ are similar and simpler. Let L_i denote the index of i -th L-suffix which needs to be placed into the LF-entry of this bucket. We distinguish the following four cases:

Index	l	$l + 1$	$l + 2$	$l + 3$...	$l + N_L - 1(r_L)$	$l + N_L$...	r
Type	L	L	L	L	...	L	S	...	S
Case (1) :									
SA	$SA[l]$	B_H	E	$SA[l + 3]$...	B_T	$SA[l + N_L]$...	$SA[r]$
SA	L_1	B_H	1	$SA[l + 3]$...	B_T	$SA[l + N_L]$...	$SA[r]$
Case (2) :									
SA	L_1	B_H	1	$SA[l + 3]$...	B_T	$SA[l + N_L]$...	$SA[r]$
SA	L_1	B_H	2	L_2	...	B_T	$SA[l + N_L]$...	$SA[r]$
⋮									
SA	L_1	B_H	$N_L - 3$	L_2	...	B_T	$SA[l + N_L]$...	$SA[r]$
SA	L_1	R_2	L_2	L_3	...	B_T	$SA[l + N_L]$...	$SA[r]$
Case (3) :									
SA	L_1	R_2	L_2	L_3	...	B_T	$SA[l + N_L]$...	$SA[r]$
SA	L_1	L_2	L_3	L_4	...	R_1	$SA[l + N_L]$...	$SA[r]$
Case (4) :									
SA	L_1	L_2	L_3	L_4	...	R_1	$SA[l + N_L]$...	$SA[r]$
SA	L_1	L_2	L_3	L_4	...	L_{N_L}	$SA[l + N_L]$...	$SA[r]$

- (1) If $SA[l + 1] = B_H$ and $SA[l + 2] = E$: The first L-suffix (i.e. L_1) needs to be placed into this bucket. We let $SA[l] = j$ and $SA[l + 2] = 1$ (use $SA[l + 2]$ as the counter to denote the number of L-suffixes have been placed so far). Recall that $suf(j)$ is the current L-suffix we want to place.
- (2) If $SA[l + 1] = B_H$ and $SA[l + 2] \neq E$: The L-suffixes (related to the current bucket) except the first L-suffix (L_1) and the last two L-suffixes (L_{N_L-1} and L_{N_L}) need to be placed. Let $c = SA[l + 2]$ (counter). If $SA[l + c + 2] \neq B_T$, we let $SA[l + c + 2] = j$ and $SA[l + 2] = c + 1$. Otherwise (this is the third to last L-suffix, i.e. L_{N_L-2}), we shift these $c - 1$ L-suffixes to the left by one position (i.e., move $SA[l + 3 \dots r_L - 1]$ to $SA[l + 2 \dots r_L - 2]$) and let $SA[r_L - 1] = j$ and $SA[l + 1] = R_2$.
- (3) If $SA[l + 1] = R_2$: The penultimate L-suffix (i.e. L_{N_L-1}) needs to be placed. We scan this bucket to find r_L such that $SA[r_L] = B_T$. Then, we move $SA[l + 2 \dots r_L - 1]$ to $SA[l + 1 \dots r_L - 2]$. After, we let $SA[r_L - 1] = j$ and $SA[r_L] = R_1$.
- (4) Otherwise, the last L-suffix (i.e. L_{N_L}) needs to be placed. We scan this bucket to find r_L such that $SA[r_L] = R_1$. Then let $SA[r_L] = j$.

We have the following lemmas and theorem.

⁸ Note that we are able to do the binary search in SA, though there are special symbols (i.e. B_H, B_T, E) in SA. Since the longest continuous special symbol entries in SA is 2, i.e., any three of continuous entries in SA must have at least one suffix entry (this entry represents the index of a suffix of T).

Lemma 20. *When we scan $SA[i]$ in the induced sorting step, we can identify whether $T[SA[i] - 1]$ is L-type or S-type in $O(1)$ time. The only exception is when $\text{suf}(SA[i] - 1)$ is the last L-suffix which needs to be inserted into the bucket $T[SA[i] - 1]$. This special case requires $O(N_L)$ time to be identified, where N_L denotes the number of L-suffixes in its bucket (i.e. $N_L = r_L - l + 1$).*

Proof. Let $j = SA[i] - 1$. First if $T[j] \neq T[j + 1]$, by definition it is trivial. Otherwise, $T[j] = T[j + 1]$. Similar to Lemma 14, we only need to know whether all L-suffixes in the bucket $T[j]$ (i.e., bucket $T[SA[i]]$) have already been sorted or not. In our algorithm, we use the extended interior counter trick which maintains the counters of the buckets. So we can identify whether all L-suffixes in the bucket $T[j]$ have already been sorted or not immediately except when $\text{suf}(j)$ is the last L-suffixes which needs to be placed into the bucket $T[j]$ (corresponding to case (4)). However, we can scan this bucket from left to right to identify whether the special symbol R_1 exists or not. If exists, which means there is one L-suffix remained, this must be the $\text{suf}(j)$. Otherwise, all L-suffixes in the bucket $T[j]$ have already been sorted. This scanning operation takes $O(N_L)$ time. \square

Lemma 21. *All the suffixes can be sorted correctly from the sorted S-suffixes in $O(n \log n)$ time.*

Proof. All S-suffixes have been sorted correctly according to Lemma 19. All L-suffixes can be sorted correctly from the sorted S-suffixes using induced sorting steps according to Lemma 1. For the time complexity, in the scanning process, we use $O(\log n)$ time to do binary search to find the head of its bucket for each L-suffix (identified by Lemma 20), and use the extended interior counter trick to find the LF-entry in constant time except for the last two L-suffixes of the bucket. We need to scan all L-suffixes in this bucket in order to find the final position (corresponding to case (4) in Step 2) and shift these L-suffixes by one position (corresponding to case (3) in Step 2). This only takes $O(n)$ time overall since each bucket is scanned at most twice. To sum up, this sorting step takes $O(n \log n)$ time. \square

Theorem 6. *Our Algorithm takes $O(n \log n)$ time and $O(1)$ workspace to compute the suffix array for string T over general alphabets, where T is read-only and only comparisons between characters are allowed.*

Proof. All steps in our algorithm take $O(n \log n)$ time. Besides, the size of recursive problem T_1 is not larger than half of $|T|$. We have $T(n) = T(n/2) + O(n \log n)$, thus $T(n) = O(n \log n + \frac{n}{2} \log \frac{n}{2} + \dots) = O(n \log n)$. For workspace, every step uses $O(1)$ workspace, and in the recursive subproblem we can also reuse the $O(1)$ workspace. Moreover, at the same recursive level, the different steps can reuse the $O(1)$ workspace too. \square

6. Conclusion

In this paper, we present optimal in-place algorithms for suffix sorting over (read-only) integer alphabets and read-only general alphabets. All of them are optimal both in time and space. Concretely, we provide the first optimal linear time in-place suffix sorting algorithm for (read-only) integer alphabets. Our algorithms solve the open problem posed by Franceschini and Muthukrishnan in ICALP 2007 [17]. For the read-only general alphabets, we provide an optimal in-place $O(n \log n)$ time suffix sorting algorithm, which recovers the result obtained by Franceschini and Muthukrishnan [17]. Our algorithm is arguably simpler.

There is a surge of interests in developing external memory algorithms for suffix sorting in recent years (see e.g., [41, 42]). Many such algorithms are extensions of existing lightweight internal memory algorithms. It would be interesting to investigate the external memory setting and see whether our tricks and data structures are applicable in this setting. We also plan to consider other string processing problems that are tightly connected with suffix array, such as compressed suffix arrays, longest common prefixes, Burrows-Wheeler transform and Lempel-Ziv factorization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The authors would like to thank Ge Nong, Gonzalo Navarro, Travis Gagie and anonymous reviewers for many helpful and constructive suggestions, which help us to improve the presentation of the paper significantly. Jian Li and Zhize Li are supported in part by the National Natural Science Foundation of China Grant 61822203, 61772297, 61632016 and the Zhongguancun Haihua Institute for Frontier Information Technology, Turing AI Institute of Nanjing and Xi'an Institute for Interdisciplinary Information Core Technology. Hongwei Huo is supported in part by National Natural Science Foundation of China Grant 61373044.

Appendix A. Running examples

A.1. Induced sorting all L-suffixes from the sorted S-suffixes

In this appendix, we give a running example for the standard induced sorting step which needs the bucket array and type array explicitly. Assume that all indices of the sorted S-suffixes are already in their correct positions in SA (i.e., in the tail of their corresponding buckets in SA). We scan SA from left to right (i.e., from SA[0] to SA[n − 1]). When we scan SA[i], let $j = SA[i] - 1$. If $T[j]$ is L-type, i.e., $\text{suf}(j)$ is an L-suffix (indicated by the type array), we place the index of $\text{suf}(j)$ (i.e. j) into the LF-entry of bucket $T[j]$, and then let the LF-pointer of this bucket $T[j]$ point to the next entry. The LF-pointers are maintained by the bucket array. If $\text{suf}(j)$ is an S-suffix, we do nothing (since all S-suffixes are already sorted in the correct positions).

The idea of induced sorting is that the lexicographical order between $\text{suf}(i)$ and $\text{suf}(j)$ are decided by the order of $\text{suf}(i + 1)$ and $\text{suf}(j + 1)$ if $\text{suf}(i)$ and $\text{suf}(j)$ are in the same bucket (i.e., $T[i] = T[j]$). Considering two L-suffixes $\text{suf}(i)$ and $\text{suf}(j)$ in the same bucket, we have $\text{suf}(i + 1) < \text{suf}(j + 1)$ and $\text{suf}(j + 1) < \text{suf}(i + 1)$ by the definition of L-suffix. Since we scan SA from left to right, $\text{suf}(i + 1)$ and $\text{suf}(j + 1)$ must appear earlier than $\text{suf}(i)$ and $\text{suf}(j)$. Hence the correctness of induced sorting is not hard to prove by induction. Note that the order of $\text{suf}(i)$ and $\text{suf}(j)$ with $T[i] \neq T[j]$ is already correct, since we always place the L-suffixes in their corresponding buckets.

Example: We use the following running example to demonstrate the induced sorting step. Consider a string $T[0 \dots 12] = \text{"2113311331210"}$ (the integer alphabets).

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S

$T[2]$ is S-type since $T[2] < T[3]$. The S-substrings are {11, 1331, 11, 1331, 1210, 0}. The S-suffixes are { $\text{suf}(1)$, $\text{suf}(2)$, $\text{suf}(5)$, $\text{suf}(6)$, $\text{suf}(9)$, $\text{suf}(12)$ }.

Now, we show the induced sorting step in our running example. Suppose all indices of the sorted S-suffixes (i.e., 1, 2, 5, 6, 9, 12) are already stored in the tail of their corresponding buckets in SA: (E denotes an Empty entry.)

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(12)	(E	1	5	9	2	6)	(E	E)	(E	E	E	E)
Bucket	(0)	(1	1	1	1	1	1)	(2	2)	(3	3	3	3)

(The entries between a pair of parentheses denote a bucket in SA which are these suffixes that start with the same character. The heads of bucket 0, 1, 2, 3 are 0, 1, 7, 9, respectively.)

The scanning process is as follows. An arrow on top of a number indicates that it is the current entry we are scanning.

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(12)	(11	1	5	9	2	6)	(E	E)	(E	E	E	E)
SA	(12)	(11	1	5	9	2	6)	(10	E)	(E	E	E	E)
SA	(12)	(11	1	5	9	2	6)	(10	0)	(E	E	E	E)
SA	(12)	(11	1	5	9	2	6)	(10	0)	(4	E	E	E)
SA	(12)	(11	1	5	9	2	6)	(10	0)	(4	8	E	E)
SA	(12)	(11	1	5	9	2	6)	(10	0)	(4	8	3	E)
SA	(12)	(11	1	5	9	2	6)	(10	0)	(4	8	3	7)

We first scan $SA[0] = 12$. We place 11 (since $T[11]$ is L-type) to the LF-entry of bucket 1 (i.e., $SA[1]$), note that the LF-pointer of bucket 1 initially points to $SA[1]$ (head of bucket 1). Next, we scan $SA[1] = 11$, and we place 10 ($T[10]$ is also L-type) to the LF-entry of its bucket (i.e., bucket 2), and so on. □

A.2. Induced sorting all L-suffixes from the sorted LMS-suffixes

In this appendix, we show the induce sorting step which sorting the L-suffixes from LMS-suffixes in our example. Note that the empty entries can be ignored, and all L-suffixes can still be sorted correctly.

Example: Suppose all LMS-suffixes (i.e., 1, 5, 9, 12) are already sorted in the tail of their corresponding bucket in SA: (E denotes an Empty entry.)

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
T	2	1	1	3	3	1	1	3	3	1	2	1	0
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(12)	(E	E	E	1	5	9)	(E	E)	(E	E	E	E)
Bucket	(0)	(1	1	1	1	1	1)	(2	2)	(3	3	3	3)

The scanning process is as follows. An arrow on top of a number indicates that it is the current entry we are scanning. When we are scanning an empty entry in SA, we ignore this entry (i.e., do nothing).

Index	0	1	2	3	4	5	6	7	8	9	10	11	12
Type	L	S	S	L	L	S	S	L	L	S	L	L	S
SA	(<u>12</u>)	(<u>11</u>)	E	E	1	5	9)	(E	E)	(E	E	E	E)
SA	(12)	(<u>11</u>)	E	E	1	5	9)	(<u>10</u>)	E)	(E	E	E	E)
SA	(12)	(11	E	E	<u>1</u>)	5	9)	(10	<u>0</u>)	(E	E	E	E)
SA	(12)	(11	E	E	1	<u>5</u>)	9)	(10	0)	(<u>4</u>)	E	E	E)
SA	(12)	(11	E	E	1	5	<u>9</u>)	(10	0)	(<u>4</u>)	<u>8</u>)	E	E)
SA	(12)	(11	E	E	1	5	9)	(10	0)	(<u>4</u>)	<u>8</u>)	<u>3</u>)	E)
SA	(12)	(11	E	E	1	5	9)	(10	0)	(4	<u>8</u>)	3	<u>7</u>)

Appendix B. Handling special symbols

We use at most five special symbols (e.g. U (Unique), E (Empty)) in this paper. The special symbols are only used to simplify the argument, and we do not need any additional assumption. Now we describe how to handle these special symbols. Note that we introduce these special symbols and use them in our interior counter trick (see Section 3.3, 3.7 and 4.4). Recall that, if one uses the bucket array (which needs extra $\max\{|\Sigma|, n/2\}$ words workspace) without using the interior counter trick, then the suffix array SA would only contain the indices of T (i.e., $\{0, 1, \dots, n - 1\}$) (see the preliminary section). Now the key point is to distinguish whether an entry in SA is one of these five integers (chosen to replace the special symbols) or just an index as before.

First, consider the simpler case where $n < 2^{\lceil \log n \rceil} - 5$. We can simply use integers $\{n, n + 1, \dots, n + 4\}$ as these five special symbols since they are different from all indices (identifiable) and each of them can be stored in one entry of SA same as an index of T .

Otherwise, we use any five integers (belong to $[0, n - 1]$) as these five special symbols. Without loss of generality, we assume that these five integers are $\{n - 5, n - 4, \dots, n - 1\}$. Then we use ten extra variables (as the previous bucket array) to indicate the head/tail of the five buckets (which the five suffixes $\{\text{suf}(n - 5), \dots, \text{suf}(n - 1)\}$ belong to) in SA and their LF/RF-entries. Thus we do not need the interior counter trick for these five buckets. Note that we can obtain the head/tail of these five buckets by scanning T once to count how many characters are smaller/larger than these five characters $\{T[n - 5], \dots, T[n - 1]\}$, respectively. To identify an entry of SA, if the entry belongs to one of these five buckets, it is just an index. Otherwise, we check its value: if it equals to one of $\{n - 5, n - 4, \dots, n - 1\}$, then it is a special symbol since none of these five integers belongs to this bucket.

There is one more subtlety. In our interior counter trick, there is a counter for each bucket to count how many suffixes have been placed into this bucket. Thus there are three types of entries in SA: 1) indices of T ; 2) these five integers (as special symbols); 3) counters for the buckets. We want to point out that the position of each counters is always fixed to be adjacent to the special symbol. Besides, the value of any counter is less than its bucket size minus 2 since the counter is deleted when we place the last two indices into its bucket (see e.g. the figure in Section 4.4). Moreover, it is not hard to verify that the counter can only conflict with the special symbol E (which happens when we first insert an index to a bucket). Thus we choose to use $n - 1$ to denote the special symbol E . There is no conflict since the counter is always less than $n - 2$.

Appendix C. Restoring the original string T

In this appendix, we show that we can restore the string T in our first algorithm (Section 3) for the integer alphabet $\{1, 2, \dots, |\Sigma|\}$ with $|\Sigma| \leq n$. Note that for our second algorithm (Section 4), we do not need to restore T since we consider the harder read-only case, i.e., T is read-only. To restore exactly the same string T , we additionally assume that the alphabet Σ is consecutive, i.e., every letter in Σ should appear in T .

First, we can see that upon the termination of our algorithm, suffix array SA contains the indices of all suffixes of T which are in lexicographical order. Note that if we do not modify T , we will have the following observation.

Observation 5. For each suffix $\text{suf}(\text{SA}[i])$ in SA, let b_i denote its bucket character (i.e., the first common character). We have that $T[\text{SA}[i]] = b_i$.

The key point to recover T is that we need to maintain the equality relationship of the characters of T . So if we modify T to T' under this condition such that $T'[i] = T'[j]$ (or $T'[i] \neq T'[j]$, resp.) if and only if $T[i] = T[j]$ (or $T[i] \neq T[j]$, resp.). Then, we can recover T from SA and T' using linear time (scan SA once) and $O(1)$ workspace from above Observation 5. Now, we need to modify the first renaming step in our algorithm to rename each character $T[i]$ to be its bucket tail (note that this modification maintains the equality relationship). This change only leads to some minor change in the details in the later induced sorting step. In the induced sorting step, since we let all $T[i]$ points to its bucket tail, the induced sort LMS-suffixes or S-suffixes are the same as before. The only thing we need to explain in the induced sorting step is that we induced sort L-suffixes from the sorted LMS-suffixes since there are no pointers which point to the bucket head (see Step 1 of Section 3.7 which sort all L-suffixes from the sorted LMS-suffixes using induced sorting). However, we can fix this step using our interior counter trick we developed. Now, we describe the details. We distinguish the buckets in SA into two types: one type does not contain LMS-suffixes, the other contains LMS-suffixes. These two types are easy to identify since the LMS-suffixes have already been sorted in the tail of their corresponding buckets in SA.

Type 1. The buckets do not contain LMS-suffixes: In this type, we initialize the bucket in the following steps. Scan T from right to left. For every $T[i]$ which is L-type and its bucket is this type, do the following:

- (1) If $SA[T[i]] = \text{Empty}$, let $SA[T[i]] = \text{Unique1}$ (unique L-type character in this bucket).
- (2) If $SA[T[i]] = \text{Unique1}$, let $SA[T[i]] = \text{Multi1}$ and $SA[T[i] - 1] = 2$ (number of L-type characters in this bucket is 2).
- (3) If $SA[T[i]] = \text{Multi1}$, increase $SA[T[i] - 1]$ by one. ($SA[T[i] - 1]$ denote the number of L-type characters in this bucket).

After this initialization, the head of this type bucket can be indicated by $SA[t]$ and $SA[t - 1]$, where t is its bucket tail.

Type 2. The buckets contain LMS-suffixes: In this type, we initialize the bucket in the following steps. Scan T from right to left. For every $T[i]$ which is L-type and its bucket is this type, do the following:

- (1) If $SA[T[i]]$ is an index, shift these LMS-suffixes (which are sorted in this bucket tail) to left by one position and let $SA[T[i]] = \text{Unique2}$ (unique L-type character in this bucket).
- (2) If $SA[T[i]] = \text{Unique2}$, shift these LMS-suffixes (which have been shifted by one position) to left by one position again and let $SA[T[i] - 1] = 2$ (number of L-type characters in this bucket is 2).
- (3) If $SA[T[i]] = \text{Multi2}$, increase $SA[T[i] - 1]$ by one. ($SA[T[i] - 1]$ denote the number of L-type characters in this bucket).

After this initialization, the head of this type bucket can be indicated by $SA[t]$ and $SA[t - 1]$ too, where t is its bucket tail.

Now, all L-suffixes can be sorted using induced sort like before, but their indices are not in their final positions in SA. We need scan T once more from right to left to compute the number of suffixes in each bucket, then shift these sorted L-suffixes to their bucket head (it is not hard to see that this shift step can be done in linear time). Now, all L-suffixes are placed in their final positions in SA. Then using induced sort as before we can sort all S-suffixes as well. In conclusion, we have the following lemma.

Lemma 22. *The original string T can be restored using linear time and $O(1)$ workspace.*

Finally, we note that if Σ is not consecutive, according to the above algorithm, we may get a string equivalent to T . For example, if $T = "13343226"$ where '5' does not appear, then we can recover an equivalent text to T , which is "13343225", i.e., '6' is replaced by '5'.

References

- [1] Z. Li, J. Li, H. Huo, Optimal in-place suffix sorting, in: Proceedings of the 25th International Symposium on String Processing and Information Retrieval, SPIRE, in: LNCS, vol. 11147, Springer, 2018, pp. 268–284.
- [2] Z. Li, J. Li, H. Huo, Optimal in-place suffix sorting, in: Proceedings of the 28th IEEE Data Compression Conference, DCC, 2018, p. 422.
- [3] U. Manber, G. Myers, Suffix arrays: a new method for on-line string searches, in: Proceedings of the First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, Society for Industrial and Applied Mathematics, 1990, pp. 319–327.
- [4] E.M. McCreight, A space-economical suffix tree construction algorithm, J. ACM 23 (1976) 262–272.
- [5] M. Farach, Optimal suffix tree construction with large alphabets, in: Proceedings of the 38th Annual Symposium on Foundations of Computer Science, FOCS, IEEE, 1997, pp. 137–143.
- [6] P. Ferragina, G. Manzini, Opportunistic data structures with applications, in: Proceedings of the 41st Annual Symposium on Foundations of Computer Science, FOCS, IEEE, 2000, pp. 390–398.
- [7] M.I. Abouelhoda, S. Kurtz, E. Ohlebusch, The enhanced suffix array and its applications to genome analysis, in: Algorithms in Bioinformatics, Springer, 2002, pp. 449–463.
- [8] R. Grossi, J.S. Vitter, Compressed suffix arrays and suffix trees with applications to text indexing and string matching, SIAM J. Comput. 35 (2005) 378–407.
- [9] H. Huo, L. Chen, J.S. Vitter, Y. Nekrich, A practical implementation of compressed suffix arrays with applications to self-indexing, in: Proceedings of IEEE Data Compression Conference, DCC, IEEE, 2014, pp. 292–301.
- [10] H. Huo, Z. Sun, S. Li, J.S. Vitter, X. Wang, Q. Yu, J. Huan, Cs2a: a compressed suffix array-based method for short read alignment, in: Proceedings of IEEE Data Compression Conference, DCC, IEEE, 2016, pp. 271–278.
- [11] M. Burrows, D.J. Wheeler, A block-sorting lossless data compression algorithm, Technical Report 124, 1994.

- [12] J. Ziv, A. Lempel, Compression of individual sequences via variable-rate coding, *IEEE Trans. Inf. Theory* 24 (1978) 530–536.
- [13] M.I. Abouelhoda, S. Kurtz, E. Ohlebusch, Replacing suffix trees with enhanced suffix arrays, *J. Discret. Algorithms* 2 (2004) 53–86.
- [14] J. Kärkkäinen, P. Sanders, Simple linear work suffix array construction, in: *Proceedings of the 30th International Colloquium on Automata, Languages and Programming, ICALP, Springer, 2003*, pp. 943–955.
- [15] P. Ko, S. Aluru, Space efficient linear time construction of suffix arrays, in: *Combinatorial Pattern Matching, CPM, Springer, 2003*, pp. 200–210.
- [16] J. Kärkkäinen, P. Sanders, S. Burkhardt, Linear work suffix array construction, *J. ACM* 53 (2006) 918–936.
- [17] G. Franceschini, S. Muthukrishnan, In-place suffix sorting, in: *Proceedings of the 34th International Colloquium on Automata, Languages and Programming, ICALP, Springer-Verlag, 2007*, pp. 533–545.
- [18] G. Nong, S. Zhang, W.H. Chan, Linear suffix array construction by almost pure induced-sorting, in: *Proceedings of IEEE Data Compression Conference, DCC, IEEE, 2009*, pp. 193–202.
- [19] G. Nong, S. Zhang, W.H. Chan, Two efficient algorithms for linear time suffix array construction, *IEEE Trans. Comput.* 60 (2011) 1471–1484.
- [20] G. Nong, Practical linear-time $O(1)$ -workspace suffix sorting for constant alphabets, *ACM Trans. Inf. Syst.* 31 (2013) 15.
- [21] S.J. Puglisi, W.F. Smyth, A.H. Turpin, A taxonomy of suffix array construction algorithms, *ACM Comput. Surv.* 39 (2007) 4.
- [22] J. Dhaliwal, S.J. Puglisi, A. Turpin, Trends in suffix sorting: a survey of low memory algorithms, in: *Proceedings of the Thirty-Fifth Australasian Computer Science Conference*, vol. 122, Australian Computer Society, Inc., 2012, pp. 91–98.
- [23] D.K. Kim, J.S. Sim, H. Park, K. Park, Linear-time construction of suffix arrays, in: *Combinatorial Pattern Matching, CPM, Springer, 2003*, pp. 186–199.
- [24] G. Manzini, P. Ferragina, Engineering a lightweight suffix array construction algorithm, in: *European Symposium on Algorithms, ESA, Springer, 2002*, pp. 698–710.
- [25] S. Burkhardt, J. Kärkkäinen, Fast lightweight suffix array construction and checking, in: *Combinatorial Pattern Matching, CPM, Springer, 2003*, pp. 55–69.
- [26] W.-K. Hon, K. Sadakane, W.-K. Sung, Breaking a time-and-space barrier in constructing full-text indices, in: *Proceedings of the 44th Annual Symposium on Foundations of Computer Science, FOCS, IEEE, 2003*, pp. 251–260.
- [27] M.A. Maniscalco, S.J. Puglisi, Faster lightweight suffix array construction, in: *Proc. of International Workshop on Combinatorial Algorithms, IWOCAL, Citeseer, 2006*, pp. 16–29.
- [28] G. Nong, S. Zhang, Optimal lightweight construction of suffix arrays for constant alphabets, in: *Algorithms and Data Structures, Springer, 2007*, pp. 613–624.
- [29] T.M. Chan, J.I. Munro, V. Raman, Selection and sorting in the restore model, in: *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA, Society for Industrial and Applied Mathematics, 2014*, pp. 995–1004.
- [30] G. Nong, S. Zhang, W.H. Chan, Linear time suffix array construction using d -critical substrings, in: *Combinatorial Pattern Matching, CPM, Springer, 2009*, pp. 54–67.
- [31] K.-B. Schürmann, J. Stoye, An incomplex algorithm for fast suffix array construction, *Softw. Pract. Exp.* 37 (2007) 309–329.
- [32] K. Sadakane, A fast algorithm for making suffix arrays and for Burrows-Wheeler transformation, in: *Proceedings of IEEE Data Compression Conference, DCC, IEEE, 1998*, pp. 129–138.
- [33] D. Baron, Y. Bresler, Antisequential suffix sorting for BWT-based data compression, *IEEE Trans. Comput.* 54 (2005) 385–397.
- [34] N.J. Larsson, K. Sadakane, Faster suffix sorting, *Theor. Comput. Sci.* 387 (2007) 258–272.
- [35] D.K. Kim, J. Jo, H. Park, A fast algorithm for constructing suffix arrays for fixed-size alphabets, in: *Experimental and Efficient Algorithms, Springer, 2004*, pp. 301–314.
- [36] H. Itoh, H. Tanaka, An efficient method for in memory construction of suffix arrays, in: *International Symposium on String Processing and Information Retrieval, SPIRE, IEEE, 1999*, pp. 81–88.
- [37] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, *Introduction to Algorithms*, MIT Press, Cambridge, 2001.
- [38] G. Jacobson, Space-efficient static trees and graphs, in: *Proceedings of the 30th Annual Symposium on Foundations of Computer Science, FOCS, IEEE, 1989*, pp. 549–554.
- [39] D. Clark, Compact pat trees, Ph.D. thesis, University of Waterloo, 1996.
- [40] J. Salowe, W. Steiger, Simplified stable merging tasks, *J. Algorithms* 8 (1987) 557–571.
- [41] P. Ferragina, T. Gagie, G. Manzini, Lightweight data indexing and compression in external memory, *Algorithmica* 63 (2012) 707–730.
- [42] G. Nong, W.H. Chan, S.Q. Hu, Y. Wu, Induced sorting suffixes in external memory, *ACM Trans. Inf. Syst.* 33 (2015) 12.