

Inducing the Lyndon Array

Felipe A. Louza¹, Sabrina Mantaci², Giovanni Manzini^{3,4},
Marinella Sciortino², and Guilherme P. Telles⁵

¹ Department of Computing and Mathematics, University of São Paulo, Brazil
louza@usp.br

² Dipartimento di Matematica e Informatica, University of Palermo, Italy
{sabrina.mantaci,marinella.sciortino}@unipa.it

³ University of Eastern Piedmont, Alessandria, Italy

⁴ IIT CNR, Pisa, Italy

giovanni.manzini@uniupo.it

⁵ Institute of Computing, University of Campinas, Brazil
gpt@ic.unicamp.br

Abstract. In this paper we propose a variant of the induced suffix sorting algorithm by Nong (TOIS, 2013) that computes simultaneously the Lyndon array and the suffix array of a text in $O(n)$ time using $\sigma + O(1)$ words of working space, where n is the length of the text and σ is the alphabet size. Our result improves the previous best space requirement for linear time computation of the Lyndon array. In fact, all the known linear algorithms for Lyndon array computation use suffix sorting as a preprocessing step and use $O(n)$ words of working space in addition to the Lyndon array and suffix array. Experimental results with real and synthetic datasets show that our algorithm is not only space-efficient but also fast in practice.

Keywords: Lyndon array, Suffix array, induced suffix sorting, lightweight algorithms

1 Introduction

The suffix array is a central data structure for string processing. Induced suffix sorting is a remarkably powerful technique for the construction of the suffix array. Induced sorting was introduced by Itoh and Tanaka [10] and later refined by Ko and Aluru [11] and by Nong *et al.* [19,18]. In 2013, Nong [17] proposed a space efficient linear time algorithm based on induced sorting, called SACA-K, which uses only $\sigma + O(1)$ words of working space, where σ is the alphabet size and the working space is the space used in addition to the input and the output. Since a small working space is a very desirable feature, there have been many algorithms adapting induced suffix sorting to the computation of data structures related to the suffix array, such as the Burrows-Wheeler transform [21], the Φ -array [8], the LCP array [4,14], and the document array [13].

The Lyndon array of a string is a powerful tool that generalizes the idea of Lyndon factorization. In the Lyndon array (LA) of string $T = T[1] \dots T[n]$ over

the alphabet Σ , each entry $\text{LA}[i]$, with $1 \leq i \leq n$, stores the length of the longest Lyndon factor of T starting at that position i . Bannai *et al.* [2] used Lyndon arrays to prove the conjecture by Kolpakov and Kucherov [12] that the number of runs (maximal periodicities) in a string of length n is smaller than n . In [3] the authors have shown that the computation of the Lyndon array of T is strictly related to the construction of the Lyndon tree [9] of the string $\$T$ (where the symbol $\$$ is smaller than any symbol of the alphabet Σ).

In this paper we address the problem of designing a space economical linear time algorithm for the computation of the Lyndon array. As described in [5,15], there are several algorithms to compute the Lyndon array. It is noteworthy that the ones that run in linear time (cf. [1,3,5,6,15]) use the sorting of the suffixes (or a partial sorting of suffixes) of the input string as a preprocessing step. Among the linear time algorithms, the most space economical is the one in [5] which, in addition to the $n \log \sigma$ bits for the input string plus $2n$ words for the Lyndon array and suffix array, uses a stack whose size depends on the structure of the input. Such stack is relatively small for non pathological texts, but in the worst case its size can be up to n words. Therefore, the overall space in the worst case can be up to $n \log \sigma$ bits plus $3n$ words.

In this paper we propose a variant of the algorithm SACA-K that computes in linear time the Lyndon array as a by-product of suffix array construction. Our algorithm uses overall $n \log \sigma$ bits plus $2n + \sigma + O(1)$ words of space. This bound makes our algorithm the one with the best worst case space bound among the linear time algorithms. Note that the $\sigma + O(1)$ words of working space of our algorithm is optimal for strings from alphabets of constant size. Our experiments show that our algorithm is competitive in practice compared to the other linear time solutions to compute the Lyndon array.

2 Background

Let $T = T[1] \dots T[n]$ be a string of length n over a fixed ordered alphabet Σ of size σ , where $T[i]$ denotes the i -th symbol of T . We denote $T[i, j]$ as the factor of T starting from the i -th symbol and ending at the j -th symbol. A suffix of T is a factor of the form $T[i, n]$ and is also denoted as T_i . In the following we assume that any integer array of length n with values in the range $[1, n]$ takes n words ($n \log n$ bits) of space.

Given $T = T[1] \dots T[n]$, the i -th rotation of T begins with $T[i + 1]$, corresponding to the string $T' = T[i + 1] \dots T[n]T[1] \dots T[i]$. Note that, a string of length n has n possible rotations. A string T is a *repetition* if there exists a string S and an integer $k > 1$ such that $T = S^k$, otherwise it is called *primitive*. If a string is primitive, all of its rotations are different.

A primitive string T is called a *Lyndon word* if it is the lexicographical least among its rotations. For instance, the string $T = abanba$ is not a Lyndon word, while it is its rotation $aabanb$ is. A *Lyndon factor* of a string T is a factor of T that is a Lyndon word. For instance, anb is a Lyndon factor of $T = abanba$.

Definition 1. Given a string $T = T[1] \dots T[n]$, the Lyndon array (LA) of T is an array of integers in the range $[1, n]$ that, at each position $i = 1, \dots, n$, stores the length of the longest Lyndon factor of T starting at i :

$$\text{LA}[i] = \max\{\ell \mid T[i, i + \ell - 1] \text{ is a Lyndon word}\}.$$

The suffix array (SA) [16] of a string $T = T[1] \dots T[n]$ is an array of integers in the range $[1, n]$ that gives the lexicographic order of all suffixes of T , that is $T_{\text{SA}[1]} < T_{\text{SA}[2]} < \dots < T_{\text{SA}[n]}$. The inverse suffix array (ISA) stores the inverse permutation of SA, such that $\text{ISA}[\text{SA}[i]] = i$. The suffix array can be computed in $O(n)$ time using $\sigma + O(1)$ words of working space [17].

Usually when dealing with suffix arrays it is convenient to append to the string T a special end-marker symbol $\$$ (called sentinel) that does not occur elsewhere in T and $\$$ is smaller than any other symbol in Σ . Here we assume that $T[n] = \$$. Note that the values $\text{LA}[i]$, for $1 \leq i \leq n - 1$ do not change when the symbol $\$$ is appended at the position n . Also, string $T = T[1] \dots T[n - 1]\$$ is always primitive.

Given an array of integers A of size n , the next smaller value (NSV) array of A , denoted NSV_A , is an array of size n such that $\text{NSV}_A[i]$ contains the smallest position $j > i$ such that $A[j] < A[i]$, or $n + 1$ if such a position j does not exist. Formally:

$$\text{NSV}_A[i] = \min\{\{n + 1\} \cup \{i < j \leq n \mid A[j] < A[i]\}\}.$$

As an example, in Figure 1 we consider the string $T = \text{banaananaanana}\$$, and its Suffix Array (SA), Inverse Suffix Array (ISA), Next Smaller Value array of the ISA (NSV_{ISA}), and Lyndon Array (LA). We also show all the Lyndon factors starting at each position of T .

If the SA of T is known, the Lyndon array LA can be computed in linear time thanks to the following lemma that rephrases a result in [9]:

Lemma 1. *The factor $T[i, i + \ell - 1]$ is the longest Lyndon factor of T starting at i iff $T_i < T_{i+k}$, for $1 \leq k < \ell$, and $T_i > T_{i+\ell}$. Therefore, $\text{LA}[i] = \ell$. \square*

Lemma 1 can be reformulated in terms of the inverse suffix array [5], such that $\text{LA}[i] = \ell$ iff $\text{ISA}[i] < \text{ISA}[i + k]$, for $1 \leq k < \ell$, and $\text{ISA}[i] > \text{ISA}[i + \ell]$. In other words, $i + \ell = \text{NSV}_{\text{ISA}}[i]$. Since given ISA we can compute NSV_{ISA} in linear time using an auxiliary stack [7,20] of size $O(n)$ words, we can then derive LA, in the same space of NSV_{ISA} , in linear time using the formula:

$$\text{LA}[i] = \text{NSV}_{\text{ISA}}[i] - i, \text{ for } 1 \leq i \leq n. \quad (1)$$

Overall, this approach uses $n \log \sigma$ bits for T plus $2n$ words for LA and ISA, and the space for the auxiliary stack.

Alternatively, LA can be computed in linear time from the Cartesian tree [22] built for ISA [3]. Recently, Franek *et al.* [6] observed that LA can be computed in linear time during the suffix array construction algorithm by Baier [1] using overall $n \log \sigma$ bits plus $2n$ words for LA and SA plus $2n$ words for auxiliary

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
T =	b	a	n	a	a	n	a	n	a	a	n	a	n	a	\$
SA =	15	14	9	4	12	7	2	10	5	1	13	8	3	11	6
ISA =	10	7	13	4	9	15	6	12	3	8	14	5	11	2	1
NSV _{ISA} =	2	4	4	9	7	7	9	9	14	12	12	14	14	15	16
LA =	1	2	1	5	2	1	2	1	5	2	1	2	1	1	1
	b	a	n	a	a	n	a	n	a	a	n	a	n	a	\$
Lyndon factors		n		a n		a n		a n		a n		a n			
		n		n		n		n		n		n			

Fig. 1. SA, ISA, NSV_{ISA}, LA and all Lyndon factors for $T = \text{banaananaanana}\$$

integer arrays. Finally, Louza *et al.* [15] introduced an algorithm that computes LA in linear time during the Burrows-Wheeler inversion, using $n \log \sigma$ bits for T plus $2n$ words for LA and an auxiliary integer array, plus a stack with twice the size as the one used to compute NSV_{ISA} (see Section 4).

Summing up, the most economical linear time solution for computing the Lyndon array is the one based on (1) that requires, in addition to T and LA, n words of working space plus an auxiliary stack. The stack size is small for non pathological inputs but can use n words in the worst case (see also Section 4). Therefore, considering only LA as output, the working space is $2n$ words in the worst case.

2.1 Induced Suffix Sorting

The algorithm SACA-K [17] uses a technique called induced suffix sorting to compute SA in linear time using only $\sigma + O(1)$ words of working space. In this technique each suffix T_i of $T[1, n]$ is classified according to its lexicographical rank relative to T_{i+1} .

Definition 2. A suffix T_i is *S-type* if $T_i < T_{i+1}$, otherwise T_i is *L-type*. We define T_n as *S-type*. A suffix T_i is *LMS-type* (leftmost S-type) if T_i is *S-type* and T_{i-1} is *L-type*.

The type of each suffix can be computed with a right-to-left scanning of T [18], or otherwise it can be computed on-the-fly in constant time during Nong's algorithm [17, Section 3]. By extension, the type of each symbol in T can be classified according to the type of the suffix starting with such symbol. In particular $T[i]$ is LMS-type if and only if T_i is LMS-type.




Fig. 2. Induced suffix sorting steps (SACA-K) for $T = banaananaana\$$

Definition 3. An LMS-factor of T is a factor that begins with a LMS-type symbol and ends with the following LMS-type symbol.

We remark that LMS-factors do not establish a factorization of T since each of them overlaps with the following one by one symbol. By convention, $T[n, n]$ is always an LMS-factor. The LMS-factors of $T = banaananaana\$$ are shown in Figure 2, where the type of each symbol is also reported. The LMS types are the grey entries. Notice that in SA all suffixes starting with the same symbol $c \in \Sigma$ can be partitioned into a c -bucket. We will keep an integer array $C[1, \sigma]$ where $C[c]$ gives either the first (head) or last (tail) available position of the c -bucket. Then, whenever we insert a value into the head (or tail) of a c -bucket, we increase (or decrease) $C[c]$ by one. An important remark is that within each c -bucket S-type suffixes are larger than L-type suffixes. Figure 2 shows a running example of algorithm SACA-K for $T = banaananaana\$$.

Given all LMS-type suffixes of $T[1, n]$, the suffix array can be computed as follows:

Steps:

1. Sort all LMS-type suffixes recursively into SA^1 , stored in $\text{SA}[1, n/2]$.
2. Scan SA^1 from right-to-left, and insert the LMS-suffixes into the tail of their corresponding c -buckets in SA .
3. Induce L-type suffixes by scanning SA left-to-right: for each suffix $\text{SA}[i]$, if $T_{\text{SA}[i]-1}$ is L-type, insert $\text{SA}[i] - 1$ into the head of its bucket.
4. Induce S-type suffixes by scanning SA right-to-left: for each suffix $\text{SA}[i]$, if $T_{\text{SA}[i]-1}$ is S-type, insert $\text{SA}[i] - 1$ into the tail of its bucket.

Step 1 considers the string T^1 obtained by concatenating the lexicographic names of all the consecutive LMS-factors (each different string is associated with a symbol that represents its lexicographic rank). Note that T^1 is defined over an alphabet of size $O(n)$ and that its length is at most $n/2$. The SACA-K algorithm is applied recursively to sort the suffixes of T^1 into SA^1 , which is stored in the first half of SA . Nong *et al.* [18] showed that sorting the suffixes of T^1 is equivalent to sorting the LMS-type suffixes of T . We will omit details of this step, since our algorithm will not modify it.

Step 2 obtains the sorted order of all LMS-type suffixes from SA^1 scanning it from right-to-left and bucket sorting then into the tail of their corresponding c -buckets in SA . Step 3 induces the order of all L-type suffixes by scanning SA from left-to-right. Whenever suffix $T_{\text{SA}[i]-1}$ is L-type, $\text{SA}[i] - 1$ is inserted in its final (corrected) position in SA .

Finally, Step 4 induces the order of all S-type suffixes by scanning SA from right-to-left. Whenever suffix $T_{\text{SA}[i]-1}$ is S-type, $\text{SA}[i] - 1$ is inserted in its final (correct) position in SA .

Theoretical costs. Overall, algorithm SACA-K runs in linear time using only an additional array of size $\sigma + O(1)$ words to store the bucket array [17].

3 Inducing the Lyndon array

In this section we show how to compute the Lyndon array (LA) during Step 4 of algorithm SACA-K described in Section 2.1. Initially, we set all positions $\text{LA}[i] = 0$, for $1 \leq i \leq n$. In Step 4, when SA is scanned from right-to-left, each value $\text{SA}[i]$, corresponding to $T_{\text{SA}[i]}$, is read in its final (correct) position i in SA . In other words, we read the suffixes in decreasing order from $\text{SA}[n], \text{SA}[n-1], \dots, \text{SA}[1]$. We now show how to compute, during iteration i , the value of $\text{LA}[\text{SA}[i]]$.

By Lemma 1, we know that the length of the longest Lyndon factor starting at position $\text{SA}[i]$ in T , that is $\text{LA}[\text{SA}[i]]$, is equal to ℓ , where $T_{\text{SA}[i]+\ell}$ is the next suffix (in text order) that is smaller than $T_{\text{SA}[i]}$. In this case, $T_{\text{SA}[i]+\ell}$ will be the first suffix in $T_{\text{SA}[i]+1}, T_{\text{SA}[i]+2}, \dots, T_n$ that has not yet been read in SA , which means that $T_{\text{SA}[i]+\ell} < T_{\text{SA}[i]}$. Therefore, during Step 4, whenever we read $\text{SA}[i]$,

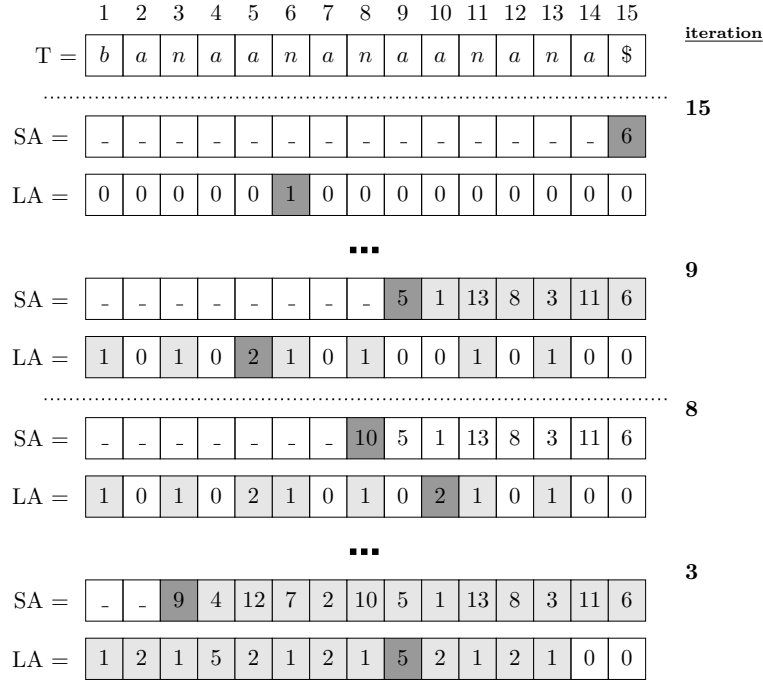


Fig. 3. Running example for $T = \text{banaananaanana}\$$.

we compute $LA[SA[i]]$ by scanning $LA[SA[i] + 1, n]$ to the right up to the first position $LA[SA[i] + \ell] = 0$, and we set $LA[SA[i]] = \ell$.

The correctness of this procedure follows from the fact that every position in $LA[1, n]$ is initialized with zero, and if $LA[SA[i] + 1], LA[SA[i] + 2], \dots, LA[SA[i] + \ell - 1]$ are no longer equal to zero, their corresponding suffixes has already been read in positions larger than i in $SA[i, n]$, and such suffixes are larger (lexicographically) than $T_{SA[i]}$. Then, the first position we find $LA[SA[i] + \ell] = 0$ corresponds to a suffix $T_{SA[i] + \ell}$ that is smaller than $T_{SA[i]}$, which was still not read in SA . Also, $T_{SA[i] + \ell}$ is the next smaller suffix (in text order) because we read $LA[SA[i] + 1, n]$ from left-to-right.

Figure 3 illustrates iterations $i = 15, 9$, and 3 of our algorithm for $T = \text{banaananaanana}\$$. For example, at iteration $i = 9$, the suffix T_5 is read at position $SA[9]$, and the corresponding value $LA[5]$ is computed by scanning $LA[6], LA[7], \dots, LA[15]$ up to finding the first empty position, which occurs at $LA[7 = 5 + 2]$. Therefore, $LA[5] = 2$.

At each iteration $i = n, n - 1, \dots, 1$, the value of $LA[SA[i]]$ is computed in additional $LA[SA[i]]$ steps, that is our algorithm adds $O(LA[i])$ time for each iteration of SACA-K.

Therefore, our algorithm runs in $O(n \cdot \text{avelyn})$ time, where $\text{avelyn} = \sum_{i=1}^n \text{LA}[i]/n$. Note that computing LA does not need extra memory on top of the space for LA[1, n]. Thus, the working space is the same as SACA-K, which is $\sigma + O(1)$ words.

Lemma 2. *The Lyndon array and the suffix array of a string $T[1, n]$ over an alphabet of size σ can be computed simultaneously in $O(n \cdot \text{avelyn})$ time using $\sigma + O(1)$ words of working space, where avelyn is equal to the average value in LA[1, n]. \square*

In the next sections we show how to modify the above algorithm to reduce both its running time and its working space.

3.1 Reducing the running time to $O(n)$

We now show how to modify the above algorithm to compute each LA entry in constant time. To this end, we store for each position LA[i] the next smaller position ℓ such that LA[ℓ] = 0. We define two additional pointer arrays NEXT[1, n] and PREV[1, n]:

Definition 4. *For $i = 1, \dots, n - 1$, $\text{NEXT}[i] = \min\{\ell \mid i < \ell \leq n \text{ and } \text{LA}[\ell] = 0\}$. In addition, we define $\text{NEXT}[n] = n + 1$.*

Definition 5. *For $i = 2, \dots, n$, $\text{PREV}[i] = \ell$, such that $\text{NEXT}[\ell] = i$ and $\text{LA}[\ell] = 0$. In addition, we define $\text{PREV}[1] = 0$.*

The above definitions depend on LA and therefore NEXT and PREV are updated as we compute additional LA entries. Initially, we set $\text{NEXT}[i] = i + 1$ and $\text{PREV}[i] = i - 1$, for $1 \leq i \leq n$. Then, at each iteration $i = n, n - 1, \dots, 1$, when we compute LA[j] with $j = \text{SA}[i]$ setting:

$$\text{LA}[j] = \text{NEXT}[j] - j \tag{2}$$

we update the pointers arrays as follows:

$$\text{NEXT}[\text{PREV}[j]] = \text{NEXT}[j], \quad \text{if } \text{PREV}[j] > 0 \tag{3}$$

$$\text{PREV}[\text{NEXT}[j]] = \text{PREV}[j], \quad \text{if } \text{NEXT}[j] < n + 1 \tag{4}$$

The cost of computing each LA entry is now constant, since only two additional computations (Equations 3 and 4) are needed. Because of the use of the arrays PREV and NEXT the working space of our algorithm is now $2n + \sigma + O(1)$ words.

Theorem 1. *The Lyndon array and the suffix array of a string $T[1, n]$ over an alphabet of size σ can be computed simultaneously in $O(n)$ time using $2n + \sigma + O(1)$ words of working space. \square*

3.2 Getting rid of a pointer array

We now show how to reduce the working space of Section 3.1 by storing only one array, say $A[1, n]$, keeping NEXT/PREV information together. In a glance, we store NEXT initially into the space of $A[1, n]$, then we reuse $A[1, n]$ to store the (useful) entries of PREV.

Note that, whenever we write $LA[j] = \ell$, the value in $A[j]$, that is $NEXT[j]$ is no more used by the algorithm. Then, we can reuse $A[j]$ to store $PREV[j+1]$. Also, we know that if $LA[j] = 0$ then $PREV[j+1] = j$. Therefore, we can redefine PREV in terms of A:

$$PREV[j] = \begin{cases} j-1 & \text{if } LA[j-1] = 0 \\ A[j-1] & \text{otherwise.} \end{cases} \quad (5)$$

The running time of our algorithm remains the same since we have added only one extra verification to obtain $PREV[j]$ (Equation 5). Observe that whenever $NEXT[j]$ is overwritten the algorithm does not need it anymore. The working space is therefore reduced to $n + \sigma + O(1)$ words.

Theorem 2. *The Lyndon array and the suffix array of a string $T[1, n]$ over an alphabet of size σ can be computed simultaneously in $O(n)$ time using $n + \sigma + O(1)$ words of working space. \square*

3.3 Getting rid of both pointer arrays

Finally, we show how to use the space of $LA[1, n]$ to store both the auxiliary array $A[1, n]$ and the final values of LA. First we observe that it is easy to compute $LA[i]$ when T_i is an L-type suffix.

Lemma 3. $LA[j] = 1$ iff T_j is an L-type suffix, or $i = n$.

Proof. If T_j is an L-type suffix, then $T_j > T_{j+1}$ and $LA[j] = 1$. By definition $LA[n] = 1$. \square

Notice that at Step 4 during iteration $i = n, n-1, \dots, 1$, whenever we read an S-type suffix T_j , with $j = SA[i]$, its succeeding suffix (in text order) T_{j+1} has already been read in some position in the interval $SA[i+1, n]$ (T_{j+1} have induced the order of T_j). Therefore, the LA-entries corresponding to S-type suffixes are always inserted on the left of a block (possibly of size one) of non-zero entries in $LA[1, n]$.

Moreover, whenever we are computing $LA[j]$ and we have $NEXT[j] = j+k$ (stored in $A[j]$), we know the following entries $LA[j+1], LA[j+2], \dots, LA[j+k-1]$ are no longer zero, and we have to update $A[j+k-1]$, corresponding to $PREV[j+k]$ (Equation 5). In other words, we update PREV information only for right-most entry of each block of non empty entries, which corresponds to a position of an L-type suffix because S-type are always inserted on the left of a block.

Then, at the end of the modified Step 4, if $A[i] < i$ then T_i is an L-type suffix, and we know that $LA[i] = 1$. On the other hand, the values with $A[i] > i$ remain

equal to $\text{NEXT}[i]$ at the end of the algorithm. And we can use them to compute $\text{LA}[i] = \text{A}[i] - i$ (Equation 2).

Thus, after the completion of Step 4, we sequentially scan $\text{A}[1, n]$ overwriting its values with LA as follows:

$$\text{LA}[j] = \begin{cases} 1 & \text{if } \text{A}[j] < j \\ \text{A}[j] - j & \text{otherwise.} \end{cases} \quad (6)$$

The running time of our algorithm is still linear, since we added only a linear scan over $\text{A}[1, n]$ at the end of Step 4. On the other hand, the working space is reduced to $\sigma + O(1)$ words, since we need to store only the bucket array $\text{C}[1, \sigma]$.

Theorem 3. *The Lyndon array and the suffix array of a string of length n over an alphabet of size σ can be computed simultaneously in $O(n)$ time using $\sigma + O(1)$ words of working space. \square*

Note that the bounds on the working space given in the above theorems assume that the output consists of SA and LA. If one is interested in LA only, then the working space of the algorithm is $n + \sigma + O(1)$ words which is still smaller than the working space of the other linear time algorithms that we discussed in Section 2.

4 Experiments

We compared the performance of our algorithm, called SACA-K+LA, with algorithms to compute LA in linear time by Franek *et al.* [5,9] (NSV-LYNDON), Baier [1,6] (BAIER-LA), and Louza *et al.* [15] (BWT-LYNDON). We also compared a version of Baier’s algorithm that computes LA and SA together (BAIER-LA+SA). We considered the three linear time alternatives of our algorithm described in Sections 3.1–3.3. We tested all three versions since one could be interested in the fastest algorithm regardless of the space usage. We used four bytes for each computer word so the total space usage of our algorithms was respectively $17n$, $13n$ and $9n$ bytes. We included also the performance of SACA-K [17] to evaluate the overhead added by the computation of LA in addition to the SA.

The experiments were conducted on a machine with an Intel Xeon Processor E5-2630 v3 20M Cache 2.40-GHz, 384 GB of internal memory and a 13 TB SATA storage, under a 64 bits Debian GNU/Linux 8 (kernel 3.16.0-4) OS. We implemented our algorithms in ANSI C. The time was measured with `clock()` function of C standard libraries and the memory was measured using `malloc_count` library⁶. The source-code is publicly available at <https://github.com/felipelouza/lyndon-array/>.

We used string collections from the Pizza & Chili dataset⁷. In particular, the datasets `einstein-de`, `kernel`, `fib41` and `cere` are highly repetitive texts⁸, and

⁶ https://github.com/bingmann/malloc_count

⁷ <http://pizzachili.dcc.uchile.cl/texts.html>

⁸ <http://pizzachili.dcc.uchile.cl/repcorpus.html>

dataset	σ	$n/2^{20}$	LA			LA and SA				SA
			NSV-LYNDON [9]	BAIER-LA [1,6]	BWT-LYNDON [15]	BAIER-LA+SA [1,6]	SACA-K+LA-17n	SACA-K+LA-13n	SACA-K+LA-9n	SACA-K [17]
<code>pitches</code>	133	53	0.15	0.20	0.20	0.26	0.26	0.22	0.18	0.13
<code>sources</code>	230	201	0.26	0.28	0.32	0.37	0.46	0.41	0.34	0.24
<code>xml</code>	97	282	0.29	0.31	0.35	0.42	0.52	0.47	0.38	0.27
<code>dna</code>	16	385	0.39	0.28	0.49	0.43	0.69	0.60	0.52	0.36
<code>english.1GB</code>	239	1,047	0.46	0.39	0.56	0.57	0.84	0.74	0.60	0.42
<code>proteins</code>	27	1,129	0.44	0.40	0.53	0.66	0.89	0.69	0.58	0.40
<code>einstein-de</code>	117	88	0.34	0.28	0.38	0.39	0.57	0.54	0.44	0.31
<code>kernel</code>	160	246	0.29	0.29	0.39	0.38	0.53	0.47	0.38	0.26
<code>fib41</code>	2	256	0.34	0.07	0.45	0.18	0.66	0.57	0.46	0.32
<code>cere</code>	5	440	0.27	0.09	0.33	0.17	0.43	0.41	0.35	0.25
<code>bbba</code>	2	100	0.04	0.02	0.05	0.03	0.05	0.04	0.03	0.03

Table 1. Running time ($\mu\text{s}/\text{input byte}$).

the `english.1G` is the first 1GB of the original `english` dataset. We also created an artificial repetitive dataset, called `bbba`, consisting of a string T with 100×2^{20} copies of b followed by one occurrence of a , that is, $T = b^{n-2}a$. This dataset represents a worst-case input for the algorithms that use a stack (NSV-LYNDON and BWT-LYNDON).

Table 1 shows the running time of each algorithm in $\mu\text{s}/\text{input byte}$. The results show that our algorithm is competitive in practice. In particular, the version SACA-K+LA-9n was only about 1.35 times slower than the fastest algorithm (BAIER-LA) for non-repetitive datasets, and 2.92 times slower for repetitive datasets. Also, the performance of SACA-K+LA-9n and BAIER-LA+SA were very similar. Finally, the overhead of computing LA in addition to SA was small: SACA-K+LA-9n was 1.42 times slower than SACA-K, whereas BAIER-LA+SA was 1.55 times slower than BAIER-LA, on average. Note that SACA-K+LA-9n was consistently faster than SACA-K+LA-13n and SACA-K+LA-17n, so using more space does not yield any advantage.

Table 2 shows the peak space consumed by each algorithm given in bytes per input symbol. The smallest values were obtained by NSV-LYNDON, BWT-LYNDON and SACA-K+LA-9n. In details, the space used by NSV-LYNDON and BWT-LYNDON was $9n$ bytes plus the space used by the stack. The stack space was negligible (about 10KB) for almost all datasets, except for `bbba` where the stack used $4n$ bytes for NSV-LYNDON and $8n$ bytes for BWT-LYNDON (the

dataset	σ	$n/2^{20}$	LA			LA and SA				SA
			NSV-LYNDON [9]	BAIER-LA [1,6]	BWT-LYNDON [15]	BAIER-LA+SA [1,6]	SACA-K+LA-17n	SACA-K+LA-13n	SACA-K+LA-9n	SACA-K [17]
pitches	133	53	9	17	9	17	17	13	9	5
sources	230	201	9	17	9	17	17	13	9	5
xml	97	282	9	17	9	17	17	13	9	5
dna	16	385	9	17	9	17	17	13	9	5
english.1GB	239	1,047	9	17	9	17	17	13	9	5
proteins	27	1,129	9	17	9	17	17	13	9	5
einstein-de	117	88	9	17	9	17	17	13	9	5
kernel	160	246	9	17	9	17	17	13	9	5
fib41	2	256	9	17	9	17	17	13	9	5
cere	5	440	9	17	9	17	17	13	9	5
bbba	2	100	13	17	17	17	17	13	9	5

Table 2. Peak space (bytes/input size).

number of stack entries is the same, but each stack entry consists of a pair of integers). On the other hand, our algorithm, SACA-K+LA-9n, used exactly $9n + 1024$ bytes for all datasets.

5 Conclusions

We have introduced an algorithm for computing simultaneously the suffix array and Lyndon array (LA) of a text using induced suffix sorting. The most space-economical variant of our algorithm uses only $n + \sigma + O(1)$ words of working space making it the most space economical LA algorithm among the ones running in linear time; this includes both the algorithm computing the SA and LA and the ones computing only the LA. The experiments have shown our algorithm is only slightly slower than the available alternatives, and that computing the SA is usually the most expensive step of all linear time LA construction algorithms. A natural open problem is to devise a linear time algorithm to construct only the LA using $o(n)$ words of working space.

Acknowledgments

The authors thank Uwe Baier for kindly providing the source codes of algorithms BAIER-LA and BAIER-LA+SA, and Prof. Nalvo Almeida for granting access to the machine used for the experiments.

Funding: F.A.L. was supported by the grant #2017/09105-0 from the São Paulo Research Foundation (FAPESP). G.M. was partially supported by PRIN grant 2017WR7SHH, by INdAM-GNCS Project 2019 *Innovative methods for the solution of medical and biological big data* and by the LSBC_19-21 Project from the University of Eastern Piedmont. S.M. and M.S. are partially supported by MIUR-SIR project CMACBioSeq *Combinatorial methods for analysis and compression of biological sequences* grant n. RBSI146R5L. G.P.T. acknowledges the support of Brazilian agencies Conselho Nacional de Desenvolvimento Científico e Tecnológico (CNPq) and Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES).

References

1. Baier, U.: Linear-time suffix sorting — a new approach for suffix array construction. In: Proc. Annual Symposium on Combinatorial Pattern Matching (CPM). pp. 23:1–23:12 (2016)
2. Bannai, H., I, T., Inenaga, S., Nakashima, Y., Takeda, M., Tsuruta, K.: The "runs" theorem. *SIAM J. Comput.* **46**(5), 1501–1514 (2017)
3. Crochemore, M., Russo, L.M.: Cartesian and Lyndon trees. *Theoretical Computer Science* (2018). <https://doi.org/10.1016/j.tcs.2018.08.011>
4. Fischer, J.: Inducing the LCP-Array. In: Proc. Workshop on Algorithms and Data Structures (WADS). pp. 374–385 (2011)
5. Franek, F., Islam, A.S.M.S., Rahman, M.S., Smyth, W.F.: Algorithms to compute the Lyndon array. In: Proc. PSC. pp. 172–184 (2016)
6. Franek, F., Paracha, A., Smyth, W.F.: The linear equivalence of the suffix array and the partially sorted Lyndon array. In: Proc. PSC. pp. 77–84 (2017)
7. Goto, K., Bannai, H.: Simpler and faster Lempel Ziv factorization. In: 2013 Data Compression Conference, DCC 2013, Snowbird, UT, USA, March 20-22, 2013. pp. 133–142 (2013)
8. Goto, K., Bannai, H.: Space efficient linear time Lempel-Ziv factorization for small alphabets. In: Proc. IEEE Data Compression Conference (DCC). pp. 163–172 (2014)
9. Hohlweg, C., Reutenauer, C.: Lyndon words, permutations and trees. *Theor. Comput. Sci.* **307**(1), 173–178 (2003)
10. Itoh, H., Tanaka, H.: An efficient method for in memory construction of suffix arrays. In: Proceedings of the sixth Symposium on String Processing and Information Retrieval (SPIRE '99). pp. 81–88. IEEE Computer Society Press (1999)
11. Ko, P., Aluru, S.: Space efficient linear time construction of suffix arrays. In: Proc. 14th Symposium on Combinatorial Pattern Matching (CPM '03). pp. 200–210. Springer-Verlag LNCS n. 2676 (2003)
12. Kolpakov, R.M., Kucherov, G.: Finding maximal repetitions in a word in linear time. In: Proc. FOCS. pp. 596–604 (1999)

13. Louza, F.A., Gog, S., Telles, G.P.: Inducing enhanced suffix arrays for string collections. *Theor. Comput. Sci.* **678**, 22–39 (2017)
14. Louza, F.A., Gog, S., Telles, G.P.: Optimal suffix sorting and LCP array construction for constant alphabets. *Inf. Process. Lett.* **118**, 30–34 (2017)
15. Louza, F.A., Smyth, W.F., Manzini, G., Telles, G.P.: Lyndon array construction during Burrows-Wheeler inversion. *J. Discrete Algorithms* **50**, 2–9 (2018)
16. Manber, U., Myers, G.: Suffix arrays: a new method for on-line string searches. *SIAM Journal on Computing* **22**(5), 935–948 (1993)
17. Nong, G.: Practical linear-time $O(1)$ -workspace suffix sorting for constant alphabets. *ACM Trans. Inf. Syst.* **31**(3), 15 (2013)
18. Nong, G., Zhang, S., Chan, W.H.: Two efficient algorithms for linear time suffix array construction. *IEEE Trans. Comput.* **60**(10), 1471–1484 (2011)
19. Nong, G., Zhang, S., Chan, W.H.: Linear suffix array construction by almost pure induced-sorting. In: *Proc. IEEE Data Compression Conference (DCC)*. pp. 193–202 (2009)
20. Ohlebusch, E.: *Bioinformatics Algorithms: Sequence Analysis, Genome Rearrangements, and Phylogenetic Reconstruction*. Oldenbusch Verlag (2013)
21. Okanohara, D., Sadakane, K.: A linear-time Burrows-Wheeler transform using induced sorting. In: *Proc. International Symposium on String Processing and Information Retrieval (SPIRE)*. pp. 90–101 (2009)
22. Vuillemin, J.: A unifying look at data structures. *Commun. ACM* **23**(4), 229–239 (1980)