






Parallel Suffix Sorting for Large String Analytics

Zhihui Du¹, Sen Zhang², and David A. Bader¹

¹ New Jersey Institute of Technology, Newark, NJ, USA
{zhihui.du,bader}@njit.edu

² State University of New York, College at Oneonta, Oneonta, NY, USA
zhangs@oneonta.edu

Abstract. The suffix array is a fundamental data structure to support string analysis efficiently. It took about 26 years for the sequential suffix array construction algorithm to achieve $\mathcal{O}(n)$ time complexity and in-place sorting. In this paper, we develop the D -Limited Parallel Induce ($DLPI$) algorithm, the first $\mathcal{O}(\frac{n}{p})$ time parallel suffix array construction algorithm. The basic idea of $DLPI$ includes two aspects: dividing the $\mathcal{O}(n)$ size problem into p reduced sub-problems with size $\mathcal{O}(\frac{n}{p})$ so we can handle them on p processors in parallel; developing an efficient parallel induce sorting method to achieve correct order for all the reduced sub-problems. The complete algorithm description is given to show the implementation method of the proposed idea. The time and space complexity analysis and proof are also given to show the correctness and efficiency of the proposed algorithm. The proposed $DLPI$ algorithm can handle large strings with scalable performance.

Keywords: Suffix Array · String Algorithm · Parallel Sorting · String Analysis · Optimal Algorithm

1 Introduction

Suffix arrays were initially introduced by Manber and Myers [18] as a space efficient alternative to suffix trees [21]. Suffix arrays can be widely used in string processing, data compression, text indexing, information retrieval and computational biology. Since the volume of string data is increasing constantly, high performance suffix array construction algorithms (SACAs) have been a challenging problem. Thirteen years after the suffix array was proposed, the first linear time algorithm for suffix sorting over integer alphabets was achieved by three research groups, Ko and Aluru [12], Kärkkäinen and Sanders [9] and Kim *et al.* [11] at almost the same time. They reduced the time complexity of suffix array construction algorithms from original $\mathcal{O}(n \log(n))$ to $\mathcal{O}(n)$. These sequential algorithms are optimal in terms of asymptotic time complexity. Furthermore, many lightweight algorithms [1, 8, 19, 20] with small working space were developed. Especially, Nong *et al.* [22] can achieve $\mathcal{O}(1)$ space complexity for constant

alphabets and Li *et al.* [17] can achieve $\mathcal{O}(1)$ in-place sorting for read-only integer alphabets. This also took about thirteen years to reduce the working space complexity from $\mathcal{O}(n)$ to $\mathcal{O}(1)$.

Many parallel SACAs have also been developed. For examples, Futamura *et al.* [4] gave a very early effort to implement a parallel SACA based on the sequential prefix-doubling method. Shun’s problem-based benchmark suite (PBBS) [26] leveraged the task-parallel Cilk Plus programming model in its parallel multicore skew algorithm implementation. Osipov [23] and Deo and Keely [2] implemented the parallel Difference Cover 3 [10] or skewed algorithm on GPU. Homann *et al.* [6] introduced the mkESA tool on multithreaded CPUs that could parallelize the sequential induce copy method. Lao *et al.* [14, 15] implemented their parallel recursive algorithm on multicore computers. All the parallel methods can significantly improve the practical performance compared with the corresponding sequential methods. Yet, none of them can handle very large string on many (p) processors in $\mathcal{O}(\frac{n}{p})$ time. To achieve scalable performance, we need a parallel SACA with $\mathcal{O}(\frac{n}{p})$ time complexity. The major contributions of this paper are as follows.

- A high level parallel suffix sorting framework is proposed. This framework aims to divide a large string’s suffix sorting problem ($T(n, p)$) into many even size reduced sub-problems ($T(\frac{n}{p}, 1)$) and the large problem can be solved by handling the many reduced sub-problems on p processors in parallel. In other words, $T(n, p) = T(\frac{n}{p}, 1)$.
- The first parallel suffix array construction algorithm *DLPI* with $\mathcal{O}(\frac{n}{p})$ time is presented. *DLPI* is optimal in terms of asymptotic time complexity.

2 Problem Description

We first give some basic definitions and notations to present the problem clearly.

Definition 1. *Suffix Array:* Given a string $S = S[0..n - 1]$ with n characters, the string’s suffix array (SA) is an array of integers providing the indices of suffixes of S in lexicographical order. This means that $\forall i < j$, we have $\text{suf}(i') < \text{suf}(j')$, where $i' = SA[i]$, $j' = SA[j]$ and $\text{suf}(k)$ is the suffix $S[k..n - 1]$.

Definition 2. *Read-only integer alphabets:* The alphabets Σ is a set of characters ($\Sigma \subseteq \mathbb{Z}$) that can be used to build a string. Given a string $S = S[0..n - 1]$ with n characters, $\forall S[i], 0 \leq i < n$, we have $S[i] \in \Sigma$. At the same time, the given string S cannot be changed during the procedure of building its suffix array. Since different characters can be encoded as different integers, we assume $\forall S[i]$, we have $S[i] \in \{x | 1 \leq x \leq |\Sigma|\}$.

In this paper, our problem is based on read-only integer alphabets instead of constant alphabets, which have only constant characters, or integer alphabets, whose input strings can be updated during the sorting procedure. The constant or integer alphabets is a special case of our problem.

The proposed problem is as follows. Given a very large string S built from a read-only integer alphabets Σ with length n and a parallel random access machine (PRAM) with p processors, can we have a parallel algorithm to build the suffix array of S in $\mathcal{O}(\frac{n}{p})$ time?

3 Algorithm Design

Unlike the existing parallel SACAs, we do not try to explore the parallelism in the framework of sequential SACAs. Instead, we first build a parallel framework that aims to divide the whole problem into many reduced sub-problems; and then develop a parallel induce method to solve all the reduced sub-problems.

Definition 3. *Order of Suffix Sets:* Given two non-empty suffix sets Set_1 and Set_2 of a string S , if $\forall x \in Set_1, \forall y \in Set_2$, their lexicographical order meets $x < y$ (or $x > y$), then we define $Set_1 < Set_2$ (or $Set_1 > Set_2$).

In this section, we propose an idea to sort the suffixes of a long string in two steps. First, we construct many (p) suffix subsets to cover all the suffixes. The suffix subsets are ordered, but suffixes in each suffix subset are not sorted. Then, we sort each suffix subset in parallel into its own sub-suffix array and achieve the complete suffix array by combining the different sub-suffix arrays corresponding to different suffix subsets together.

Algorithm 1: DLPI Algorithm

```

1  Function DLPI(String,  $p$ )
2      Step (1) Build parallel reduced subproblems
3      1.1 Divide all suffixes of  $S$  into  $p$  suffix subsets
            $SubSet_1, \dots, SubSet_p, \forall 1 \leq i \leq p, |SubSet_i| = \mathcal{O}(\frac{n}{p})$ 
4      1.2 Call Parallel Suffix SubSets Sorting function  $SA = PSSS(SubSet_1, \dots, SubSet_p)$ 
5      1.3 Evenly select  $(p - 1)$  splitters from each processor  $p_i$ 's returned suffix array  $SA[i]$ 
6      1.4 Add the  $(p - 1) \times p$  splitters into each subset to get  $SpSubSet_i, 1 \leq i \leq p$ 
7      1.5 Call Parallel Suffix SubSets Sorting function
            $SA = PSSS(SpSubSet_1, \dots, SpSubSet_p)$ 
8      1.6 According to the returned  $SA$ , divide all suffixes into  $p$  ordered subsets that meet
            $OSubSet_1 < \dots < OSubSet_p$ 
9      Step (2) Sort reduced subproblems in parallel
10     2.1 Call Parallel Suffix SubSets Sorting function  $SA = PSSS(OSubSet_1, \dots, OSubSet_p)$ 
11     2.2 return SA
12 end
    
```

3.1 Algorithm Framework

In Algorithm 1, we present the framework of our parallel suffix array construction algorithm D -Limited Parallel Induce ($DLPI$). This framework transforms a large $T(n, p)$ problem, where n represents the problem size and p represents the number processors of the PRAM, into p parallel $T(\frac{n}{p}, 1)$ problems, which means that each single problem is of size $\mathcal{O}(\frac{n}{p})$ and can be handled in one processor.

In line 3 of Algorithm 1, for all the n suffixes of a given string S , we assign them into p subsets evenly. The $PSSS$ function will generate different sub-suffix arrays corresponding to different subsets and we can select $(p - 1)$ different splitters [5] that divide each subset evenly (line 5). Then we augment each subset by adding $(p - 1) \times (p - 1)$ splitters from the rest of the subsets and call $PSSS$ again on the augmented subsets. Afterwards, we utilize the ordered $p \times (p - 1)$ splitters as guides to assign all the suffixes (in a round-robin order) into p ordered subsets, $OSubSet_1, \dots, OSubSet_p$ (lines 6–8). Here, all suffixes are assigned into p ordered subsets with the size of $\mathcal{O}(\frac{n}{p})$.

The second step is straightforward, where we just call *PSSS* again to generate the order of suffixes in different subsets and then combine them together to obtain the complete suffix array *SA* (lines 10–11).

3.2 Parallel Induce Method

In Algorithm 2, we describe the essential function *PSSS* that can support parallel induce on all reduced sub-problems. The basic idea of this function is that we first construct p much smaller strings to express the different sub-problems. The suffixes with different short prefixes can be sorted easily and the novel parallel induce method is used to derive the order of suffixes with long and the same prefixes.

Algorithm 2: Parallel Suffix SubSets Sorting Algorithm

```

1  Function PSSS(SubSet1, ..., SubSetp)
2  Step (1) Sort suffixes of each subsets and distinguish Fixed and Changeable suffixes
3  Build D-limited shrunk strings DS-S1, ..., DS-Sp according to different subsets
4  forall (i in 1..p) do
5      ESA[i][] = SeqOptSA(DS-Si)
6      Remove all indices ESA[i][j] that are not in Seti and get SA[i][] corresponding to
          SubSeti
7      var mg = -1
8      for (j in 0..|SubSeti|-1) do
9          if (Suf(SA[i][j]) and its closest suffix in SA have the same D prefix) then
10             Flag[i][j] = Changeable
11             if (Suf(SA[i][j]) is the first Changeable suffix of a new group) then
12                 mg++
13                 ChgGrp[i][mg].head = j; ChgGrp[i][mg].num = 1
14             end
15             ChgGrp[i][mg].num++
16         end
17         else
18             Flag[i][j] = Fixed
19         end
20     end
21 end
22 Step (2) Induce the order of Changeable suffixes in each Changeable suffix group
23 2.1 Build aligned subsets AliSubSet1, ..., AliSubSetp for Changeable suffix groups
24 2.2 Generate the new suffix array AliSA of the suffixes just like the previous step (1)
25 2.3 Generate the distinguishable tail suffix array DTA for suffixes in the Changeable
    suffix groups
26 2.4 Induce the correct order of all Changeable suffixes in SA based on DTA and AliSA
27 return SA
28 end

```

We introduce the first step of *PSSS* function at first.

Definition 4. *D*-limited substring and *D*-limited shrunk string: Given a constant *D*, a string *S* with length *n* and one of its suffix subset *SubSet*, if two suffixes *suf*(*i*) ∈ *SubSet* and *suf*(*j*) ∈ *SubSet*, where *i* < *j* and no other suffix sits between *i* and *j* in *SubSet* (we will let *j* = *n* if no such *suf*(*j*) in *SubSet*), then the *D*-limited substring of *suf*(*i*) is the substring *S*[*i*..*j* - 1] if *j* - *i* ≤ *D* or *S*[*i*..*i* + *D* - 1] if *j* - *i* > *D*. The *D*-limited shrunk string *DS* of *S* is the string by concatenating all *D*-limited substrings from *SubSet* together according to their original order in *S*.

Definition 5. *D-prefix substring:* Given a constant D , a string S with length n , and an index i , the D -prefix substring of $\text{suf}(i)$ is the substring $S[i..i+D-1]$ if $i < n - D$ or $S[i..n]$ otherwise.

Building Reduced Strings. The first step of this function is building p much smaller D -limited shrunk strings DS_1, \dots, DS_p so each processor can handle one smaller string in parallel (line 3). We use D -prefix substrings to replace the original suffixes.

We will call the existing optimal sequential SACA *SeqOptSA* [17] to generate the extended suffix array for the given shrunk string. Since we do not need to compare the suffixes not included in the given subset, we remove the indices of such suffixes in the extended suffix array and get the exact suffix array SA (lines 5–6). We use a two-dimension array to express the partitioned data in different processors. The cardinality of the first dimension stands for the number of processors and that of the second dimension stands for the maximum number of suffixes assigned to different processors.

For the suffix whose order can be decided based on its D -prefix substring, its rank in the suffix array is correct. If there are two or more suffixes whose D -prefix substrings are exactly the same, their ranks in SA should be induced based on their complete suffixes. We use a two-dimension array *Flag* to mark the correct rank as Fixed and the rank to be induced as Changeable. At the same time, we use a two-dimension array *ChgGrp* to manage the clustered Changeable suffixes by their D -prefix substrings. $\text{ChgGrp}[i][mg]$ keeps the current group of Changeable suffixes on processor i . $\text{ChgGrp}[i][mg].\text{head}$ is the rank of the first suffix in the corresponding suffix array and $\text{ChgGrp}[i][mg].\text{num}$ is the total number of suffixes in the current group (lines from 7 to 20).

Based on the *ChgGrp* data structure, the induce sorting method works as follows. When we know the smallest suffix within the group mg , we just need to switch the rank of the smallest suffix with that of the head suffix, advance $\text{ChgGrp}[i][mg].\text{head}$ by one, and reduce $\text{ChgGrp}[i][mg].\text{num}$ by 1. If a suffix can split the suffixes into two ordered subsets, we will put the suffix at the correct position in its SA and split its Changeable suffix group into two smaller groups. This way, we can induce one suffix at its correct position. When $\text{ChgGrp}[i][mg].\text{num}$ is one, all suffixes in the Changeable group mg have been correctly ranked. The suffixes in different groups can be induced in parallel.

The second step is to induce the correct ranks of Changeable suffixes (line 22). The basic idea is building induce chain for all the Changeable suffixes; then identifying the tail suffix that can distinguish the Changeable suffix from other suffixes; inducing the order for each Changeable suffix based on the tail suffixes. It includes four substeps and we will present the detailed descriptions as follows.

Definition 6. *Aligned suffix set:* Given a Changeable suffix group CG and a non negative integer k , the set $\{\text{suf}(x) \mid \forall e \in CG, e = \text{suf}(y) \wedge x = y + D \times k \wedge x < n\}$ is the k aligned suffix set of CG .

Building Aligned Suffix Sets. In the first substep (line 23) we build p completely new suffix subsets $AliSubSet_1, \dots, AliSubSet_p$ that are used to induce the correct order of all the Changeable suffixes. Suffixes in an aligned suffix set will be assigned to the same processor so we can get their order based on each processor's suffix array.

For all the Changeable suffix groups, we can generate all of their k aligned suffix set. We will merge some overlapping aligned suffix sets and assign these sets into p processors and form p suffix subsets $AliSubSet_1, \dots, AliSubSet_p$.

Generating SA for Aligned Suffix Sets. In the second substep, we may employ the similar method as before (lines from 2 to 21) to generate the suffix array of each aligned suffix subset. Here we use $AliSA$ to express the new suffix array corresponding to the aligned suffixes. $AliFlag$ has the similar meaning as before to mark the Fixed and Changeable suffixes.

Now we have obtained two sets of Fixed and Changeables, we will use the later set of Fixed suffixes to induce the previous set of Changeable suffixes by using a data structure called DTA (to be further defined next).

Building the Distinguishable Tail Suffix Array. In the third substep, we will build an array DTA to store the suffixes that can be used to distinguish one Changeable suffix from other suffixes in the same Changeable group.

Definition 7. *Distinguishable Tail Suffix:* For any Changeable suffix $suf(x)$ in a Changeable suffix group $ChgGrp$, its distinguishable tail suffix $suf(DTA(x))$ is the suffix that can distinguish the order of $suf(x)$ from the other Changeable suffixes according to $suf(DTA(x))$'s D -prefix substring.

We will transfer the index t of suffix $suf(t)$ whose flag is Fixed to its left suffix $suf(t - D)$ and let $DTA[t - D] = t$ if $suf(t - D)$ exists and it is a changeable suffix. This procedure will continue to the head of the string along the induce chain of $suf(t)$. The challenge here is that we should do it in parallel. The basic idea is as follows.

We first assign all suffixes to different processors based on the indices of different suffixes evenly and each processor only checks about $\frac{n}{p}$ suffixes. For suffixes assigned to the current processor i , we will cluster them $\frac{n}{p}$ into D classes based on their indices' modulo D values. Each processor will scan every class from its end suffix to its start suffix, in parallel. The index of the Fixed suffix $suf(f)$ will be passed to its left suffix $suf(f - D)$ one by one until the new Fixed suffix is met. Then the index of the new Fixed suffix will replace the old one and be passed to the left suffix. In order to pass the distinguishable indices across processors, we use a 2-D temporary array $tmp[D][p]$ consisting of $(D \times p)$ elements. If an ending suffix $suf(e)$ (there are D of them) is changeable, we let $tmp[e\%D][i] = -1$ and $DTA[e] = -(i)$, where i refers to the id of the processor (ranged from 0 to $(p - 1)$) that means that $suf(e)$'s DTA value $DTA(e)$ is unknown and it will get its value from $tmp[e\%D][i]$. Each Changable $suf(j)$ on

the processor i that cannot get its distinguishable tail suffix from upto its last D -prefix suffix on the same processor will point to the same element $tmp[j\%D][i]$.

After this, we will scan the temporary array from end to start for different modulo values. For current temporary element $tmp[d][i]$ that corresponds to the (i) processor and d^{th} class, if $sf(r)$ is the first suffix of its right processor, $r\%D = d$, and $DTA[r] > 0$, we will let $tmp[d][i] = DTA[r]$. If not, we will let $tmp[d][i] = tmp[d][i + 1]$. This means that the temporary array update will start from $tmp[d][p - 2]$ and end with $tmp[d][0]$ (for all d in $[0..D - 1]$), sequentially. Finally, each processor will check its suffixes whose DTA values are still negative and update them with their respective corresponding temporary values. This way, we can propagate the distinguishable tail suffixes from end to start in parallel.

Inducing the Order of Changeable Suffixes. In the fourth substep, we know the distinguishable tail suffix of each Changeable suffix. We can use such information to induce the order of suffixes in each Changeable suffix group. For each Changeable suffix group, first we will use its closest distinguishable tail suffixes to distinguish the corresponding Changeable suffixes from others. Then, we will induce the correct order of all Changeable suffixes based on their distinguishable tail suffixes' indices from small to large. The order of different Changeable suffix groups can be induced in parallel.

4 Complexity Analysis

In this section, we adopt the widely used Parallel Random Access Machine (PRAM) model [7] to analyze our parallel algorithm. The time complexity of the proposed algorithm is $\mathcal{O}(\frac{n}{p})$ and the space complexity is $\mathcal{O}(n)$. We will prove that every step of our algorithm can be done in $\mathcal{O}(\frac{n}{p})$ time and at most $\mathcal{O}(n)$ working space (the space except the input string S and the returned suffix array) is needed to generate the complete suffix array.

The *DLPI* function gives the framework of our algorithm. For substep 1.1 of step 1, we can assign the suffixes of the given string S with length n into p parts and each has $\mathcal{O}(\frac{n}{p})$ elements in $\mathcal{O}(\frac{n}{p})$ time. The p D -limited shrunk strings will need $\mathcal{O}(p \times D \times \frac{n}{p}) = \mathcal{O}(D \times n) = \mathcal{O}(n)$ space. For substep 1.2, we will give the time and space complexity of the parallel induce function *PSSS* later. Selecting $(p - 1)$ splitters for each processor based on its returned suffix array and adding them into different subsets are straightforward and can also be done in $\mathcal{O}(\frac{n}{p})$ time. Here we assume $p^3 < n$, when we add $(p - 1)^2$ new elements to each subset, each subset will have $\mathcal{O}(\frac{n}{p}) + \mathcal{O}((p - 1)^2) \leq \mathcal{O}(\frac{n}{p}) + \mathcal{O}(\frac{n}{p}) = \mathcal{O}(\frac{n}{p})$ elements. So, the total working space will also be $\mathcal{O}(p \times \frac{n}{p}) = \mathcal{O}(n)$.

For substep 1.5, just like before, we will discuss the time complexity of *PSSS* later. From substeps 1.3 to 1.6, we know that the total number of elements between two closest splitters cannot be larger than $\mathcal{O}(\frac{n}{p^2})$. So, when we combine p parts of elements divided by the same splitters together into one subset, its

size cannot be larger than $\mathcal{O}(\frac{n}{p})$. At the same time, the elements of each subset will be no larger than $\mathcal{O}(\frac{n}{p})$. Based on this conclusion, it is feasible for us to build p ordered subsets according to the $p \times (p - 1)$ splitters.

Hence, we can claim that *DLPI* function can generate the complete suffix array of a given string S with length n in $\mathcal{O}(\frac{n}{p})$ time on p processors using $\mathcal{O}(n)$ space if the parallel induce function *PSSS* can return the suffix array for each suffix subset in $\mathcal{O}(\frac{n}{p})$ time on p processors using $\mathcal{O}(n)$ working space.

Theorem 1. *For a string S with length n , if its suffixes are assigned to p given subsets with size $\mathcal{O}(\frac{n}{p})$, then the corresponding D -prefix substrings of each subset can be sorted in $\mathcal{O}(\frac{n}{p})$ time on p processors with $\mathcal{O}(n)$ space.*

Proof. We can build D -limited shrunk strings based on given p suffix subsets of the string S in parallel. The shrunk strings can be done by directly concatenating all the D -limited substrings corresponding to the suffixes in each subset directly. This work will take $\mathcal{O}(\frac{n}{p})$ time with $\mathcal{O}(n)$ space. Then, we may employ the existing in-place sequential linear suffix array algorithm *SeqOptSA* to directly return their corresponding extended suffix arrays in $\mathcal{O}(\frac{n}{p})$ time. The extended suffix arrays will contain more indices than each subset's elements. So, we need to remove the additional indices. This can also be done in at most $\mathcal{O}(\frac{n}{p})$ time. Totally, $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space will be needed to sort the D -prefix substrings of suffixes in all the given subsets.

Lemma 1. *All suffixes can be marked as Fixed or Changeable suffixes and clustered into groups in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.*

Proof. To mark all suffixes as Fixed or Changeable, a *Flag*[1.. p] array with $\mathcal{O}(n)$ space will be needed. To store the Changeable group information, at most $\mathcal{O}(\frac{n}{2})$ space for a Changeable suffix array *ChgGrp*[1.. p] will be needed because the suffixes can be divided into at most $(\frac{n}{2})$ groups. Based on the returned suffix array, each processor can compare any suffix's D -prefix substring with its neighbor to check if they are the same. The different D -prefix substrings mean that the corresponding suffixes can be marked as Fixed; otherwise, they will be marked as Changeable. The entire character comparison operations for any processor should be $\mathcal{O}(D \times \frac{n}{p}) = \mathcal{O}(\frac{n}{p})$. Clustering the Changeable suffixes based on their D -prefix substrings and storing the group information into *ChgGrp* are similar. So, the marking and clustering operations can be done in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.

Lemma 2. *The aligned subsets $AliSubSet_1, \dots, AliSubSet_p$ that each is no more than $\mathcal{O}(\frac{n}{p})$ elements can be built in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.*

Proof. The total number of aligned suffixes cannot be larger than $\mathcal{O}(n)$. Since we combine some overlapping aligned suffix sets and assign them to different processors evenly, the total number of suffixes assigned to one processor cannot be larger than $\mathcal{O}(\frac{n}{p})$. The total number of suffixes in the Changeable suffix groups cannot be larger than $\mathcal{O}(n)$, and the total number of suffixes in all aligned suffix

sets cannot be larger than $\mathcal{O}(n)$ either. So, for the first substep, totally at most $\mathcal{O}(n)$ space will be needed to store all the suffixes. Generating at most $\mathcal{O}(\frac{n}{p})$ suffixes $AliSubSet_1, \dots, AliSubSet_p$ for each processor from the Changeable suffix groups is straightforward and can be done at most in $\mathcal{O}(\frac{n}{p})$ time.

Corollary 1. *The D -prefix substrings of subsets $AliSubSet_1, \dots, AliSubSet_p$ can be sorted in $\mathcal{O}(\frac{n}{p})$ time on p processors with $\mathcal{O}(n)$ space.*

Proof. $AliSubSet_1, \dots, AliSubSet_p$ are p suffix subsets and each of them have at most $\mathcal{O}(\frac{n}{p})$ suffixes. Based on theorem 1, we can get the corollary and the third substep can be done in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.

Subsequently, again all suffixes of AliSubSets can be marked as Fixed or Changeable suffixes and clustered into groups in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space, a direct application of Lemma 1 on the new subsets i.e. AliSubsets.

Lemma 3. *The distinguishable tail suffix array DTA can be generated in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.*

Proof. We can allocate the DTA array with size n to cover all suffixes. So, $\mathcal{O}(n)$ space is sufficient. The basic idea of distinguishable tail suffix generation is passing the closest Fixed suffix to the current Changeable suffix and storing the Fixed suffix's index in DTA . The short passing path will be easy to implement. In order to reduce the passing time for a very long passing path, our implementation method divides the long passing path into multiple parallel subpaths. The suffix passing can be done on different subpaths in parallel. We allocate at most p temporary memory space to transfer the index across different processors. Since all the suffixes assigned to one processor cannot be larger than $\mathcal{O}(\frac{n}{p})$, the first scan procedure can be done in $\mathcal{O}(\frac{n}{p})$ time for all the processors. Then we let one processor pass the value in the temporary memory one by one from end to start. So, at most $\mathcal{O}(p)$ time is needed. Finally, during the last scan, every processor will assign the suffixes with the value of the temporary memory space if they point to this memory space. The third substep will need at most $\mathcal{O}(\frac{n}{p})$ time. So, totally, $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space are needed to generate DTA .

Lemma 4. *Inducing the order of all Changeable suffixes based on DTA and $AliSA$ can be done in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.*

Proof. Generating the relative order of suffixes in each Changeable group based on its DTA can be done in $\mathcal{O}(\frac{n}{p})$ time on each processor because the length of each suffix to be sorted will be no more than D (D or $<D$ at the end of the string), and we have at most $\mathcal{O}(\frac{n}{p})$ such suffixes for each processor. It will need to scan all the corresponding distinguishable tail suffixes to induce the order of Changeable suffixes. The total number of distinguishable tail suffixes is the same as the total number of Changeable suffixes that is no more than $\mathcal{O}(\frac{n}{p})$ on each processor. So, the induce procedure also can be done in $\mathcal{O}(\frac{n}{p})$ time. The total space to keep the $AliSA$ and the temporary string is no more than $\mathcal{O}(n)$. So the fourth substep can also be done in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space.

Theorem 2. For a string S with length n , its suffix array can be generated in $\mathcal{O}(\frac{n}{p})$ time on p processors with $\mathcal{O}(n)$ space in parallel.

Proof. Based on the above theorem and the lemmas, every algorithm step can be done in $\mathcal{O}(\frac{n}{p})$ time and $\mathcal{O}(n)$ space. So, after adding them together, we will get the conclusion.

5 Related Work

There have been many works on the suffix array construction algorithm since suffix array was invented in 1990 by Manber and Myers [18]. “Induce” is an essential technique in suffix sorting. Although prefix-doubling [24] adopts the induce technique, it cannot reduce the problem size step by step. This is why it cannot achieve $\mathcal{O}(n)$ time complexity. The following works [9, 11, 12] recursively solve the problem by constructing a reduced problem and employing the induce technique to sort the suffixes.

All existing parallel suffix array construction algorithms were trying to parallelize one or combined sequential algorithms. Futamura *et al.* [4] gave the early effort to parallel the prefix-doubling method. Larsson *et al.* [16] implemented optimized methods based on the previous prefix-doubling technology and improved its performance in parallel. Osipov *et al.* [23] implemented prefix-doubling algorithm on GPUs. Flick and Aluru [3]’s parallel MPI-based implementation of the prefix-doubling method can achieve very high practical performance on human genome datasets. Kulla *et al.* [13] parallelized the sequential DC3 method, which regularly samples the string to build a smaller $\frac{2}{3}n$ problem. Deo *et al.* [2] further implement the DC3 method on GPUs. Shun [25]’s parallel skew (DC3) algorithm could achieve good performance on shared-memory multi-core computers. Wang *et al.* [27] implemented a hybrid prefix-doubling and DC3 method on GPUs to improve the existing GPU methods significantly. Lao *et al.* [14, 15] employed pipeline technology to parallelize their previous sequential linear algorithms on multicore computers.

The existing sequential algorithm framework is the barrier to the existing parallel methods of achieving scalable performance. We develop a parallel framework and propose a parallel induce method to achieve $\mathcal{O}(\frac{n}{p})$ time complexity.

6 Conclusion

The novel idea provided in this paper is the concept of D -limited shrunk substrings that divides the complete problem with size n into p reduced sub-problems with size $\mathcal{O}(\frac{n}{p})$. An optimal parallel suffix array construction algorithm to handle the problem in $\mathcal{O}(\frac{n}{p})$ time complexity (p is the number of parallel processors and we assume $p^3 < n$) is critical for us to handle large strings (built on an integer alphabet) with scalable performance. The critical technology is parallel induce. The suffixes with long repeat prefixes can induce their order based on their distinguishable tail suffixes in parallel. We take advantage of the existing

optimal sequential suffix array construction algorithm as an independent execution unit to generate the order of all D -prefix substrings that can be used to separate suffixes with long repeat prefixes from those with short unique prefixes.

The simplicity and the $\mathcal{O}(\frac{n}{p})$ time complexity make the proposed D -Limited Parallel Induce (*DLPI*) algorithm very promising to handle huge strings with scalable performance. *DLPI* is the first parallel suffix array construction algorithm with $\mathcal{O}(\frac{n}{p})$ time complexity. We will focus on further reducing the total working space $\mathcal{O}(n)$ in the future work.

Acknowledgement. This research was funded in part by NSF grant number CCF-2109988.

References

1. Burkhardt, S., Kärkkäinen, J.: Fast lightweight suffix array construction and checking. In: Baeza-Yates, R., Chávez, E., Crochemore, M. (eds.) CPM 2003. LNCS, vol. 2676, pp. 55–69. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44888-8_5
2. Deo, M., Keely, S.: Parallel suffix array and least common prefix for the GPU. In: Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, pp. 197–206 (2013). <https://doi.org/10.1145/2442516.2442536>
3. Flick, P., Aluru, S.: Parallel distributed memory construction of suffix and longest common prefix arrays. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 1–10 (2015)
4. Futamura, N., Aluru, S., Kurtz, S.: Parallel suffix sorting. Electrical Engineering and Computer Science - All Scholarship 64 (2001). <https://surface.syr.edu/eecs/64>
5. Helman, D.R., JáJá, J., Bader, D.A.: A new deterministic parallel sorting algorithm with an experimental evaluation. J. Exp. Algorithmics (JEA) **3**, 4-es (1998). <https://doi.org/10.1145/297096.297128>
6. Homann, R., Flier, D., Giegerich, R., Rehmsmeier, M.: mkESA: enhanced suffix array construction tool. Bioinformatics **25**(8), 1084–1085 (2009). <https://doi.org/10.1093/bioinformatics/btp112>
7. JáJá, J.: An Introduction to Parallel Algorithms, vol. 10, p. 133889. Addison-Wesley, Reading (1992)
8. Kärkkäinen, J.: Fast BWT in small space by blockwise suffix sorting. Theoret. Comput. Sci. **387**(3), 249–257 (2007). <https://doi.org/10.1016/j.tcs.2007.07.018>
9. Kärkkäinen, J., Sanders, P.: Simple linear work suffix array construction. In: Baeten, J.C.M., Lenstra, J.K., Parrow, J., Woeginger, G.J. (eds.) ICALP 2003. LNCS, vol. 2719, pp. 943–955. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-45061-0_73
10. Kärkkäinen, J., Sanders, P., Burkhardt, S.: Linear work suffix array construction. J. ACM (JACM) **53**(6), 918–936 (2006). <https://doi.org/10.1145/1217856.1217858>
11. Kim, D.K., Sim, J.S., Park, H., Park, K.: Linear-time construction of suffix arrays. In: Baeza-Yates, R., Chávez, E., Crochemore, M. (eds.) CPM 2003. LNCS, vol. 2676, pp. 186–199. Springer, Heidelberg (2003). https://doi.org/10.1007/3-540-44888-8_14

12. Ko, P., Aluru, S.: Space efficient linear time construction of suffix arrays. *J. Discret. Algorithms* **3**(2–4), 143–156 (2005). <https://doi.org/10.1016/j.jda.2004.08.002>
13. Kulla, F., Sanders, P.: Scalable parallel suffix array construction. *Parallel Comput.* **33**(9), 605–612 (2007)
14. Lao, B., Nong, G., Chan, W.H., Pan, Y.: Fast induced sorting suffixes on a multi-core machine. *J. Supercomput.* **74**(7), 3468–3485 (2018). <https://doi.org/10.1007/s11227-018-2395-5>
15. Lao, B., Nong, G., Chan, W.H., Xie, J.Y.: Fast in-place suffix sorting on a multicore computer. *IEEE Trans. Comput.* **67**(12), 1737–1749 (2018). <https://doi.org/10.1109/TC.2018.2842050>
16. Larsson, N.J., Sadakane, K.: Faster suffix sorting. *Theoret. Comput. Sci.* **387**(3), 258–272 (2007). <https://doi.org/10.1016/j.tcs.2007.07.017>
17. Li, Z., Li, J., Huo, H.: Optimal in-place suffix sorting. *Inf. Comput.* **285**, 104818 (2022)
18. Manber, U., Myers, G.: Suffix arrays: a new method for on-line string searches. *SIAM J. Comput.* **22**(5), 935–948 (1993). <https://doi.org/10.1137/0222058>
19. Maniscalco, M.A., Puglisi, S.J.: Faster lightweight suffix array construction. In: *Proceedings of International Workshop on Combinatorial Algorithms (IWOCA)*, pp. 16–29 (2006)
20. Manzini, G., Ferragina, P.: Engineering a lightweight suffix array construction algorithm. *Algorithmica* **40**(1), 33–50 (2004). <https://doi.org/10.1007/s00453-004-1094-1>
21. McCreight, E.M.: A space-economical suffix tree construction algorithm. *J. ACM (JACM)* **23**(2), 262–272 (1976). <https://doi.org/10.1145/321941.321946>
22. Nong, G.: Practical linear-time $O(1)$ -workspace suffix sorting for constant alphabets. *ACM Trans. Inf. Syst. (TOIS)* **31**(3), 1–15 (2013). <https://doi.org/10.1145/2493175.2493180>
23. Osipov, V.: Parallel suffix array construction for shared memory architectures. In: Calderón-Benavides, L., González-Caro, C., Chávez, E., Ziviani, N. (eds.) *SPIRE 2012*. LNCS, vol. 7608, pp. 379–384. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34109-0_40
24. Puglisi, S.J., Smyth, W.F., Turpin, A.H.: A taxonomy of suffix array construction algorithms. *ACM Comput. Surv. (CSUR)* **39**(2), 4-es (2007). <https://doi.org/10.1145/1242471.1242472>
25. Shun, J.: Fast parallel computation of longest common prefixes. In: *SC 2014: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 387–398. IEEE (2014). <https://doi.org/10.1109/SC.2014.37>
26. Shun, J., et al.: Brief announcement: the problem based benchmark suite. In: *Proceedings of the Twenty-Fourth Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 68–70 (2012). <https://doi.org/10.1145/2312005.2312018>
27. Wang, L., Baxter, S., Owens, J.D.: Fast parallel skew and prefix-doubling suffix array construction on the GPU. *Concurr. Comput. Pract. Exp.* **28**(12), 3466–3484 (2016). <https://doi.org/10.1002/cpe.3867>