

# Scalable Construction of Text Indexes

Timo Bingmann<sup>1</sup>, Simon Gog<sup>1</sup>, and Florian Kurpicz<sup>2</sup>

<sup>1</sup> Institute of Theoretical Informatics,  
Karlsruhe Institute of Technology, Germany

<sup>2</sup> Department of Computer Science,  
Technische Universität Dortmund, Germany

**Abstract.** The suffix array is the key to efficient solutions for myriads of string processing problems in different applications domains, like data compression, data mining, or Bioinformatics. With the rapid growth of available data, suffix array construction algorithms had to be adapted to advanced computational models such as external memory and distributed computing. In this article, we present five suffix array construction algorithms utilizing the new algorithmic big data batch processing framework Thrill, which allows us to process input sizes in orders of magnitude that have not been considered before.

## 1 Introduction

Suffix arrays [15,7] are the basis for many text indexes and string algorithms. Suffix array construction is theoretically linear work, but practical suffix sorting is computationally intensive and often limits the applicability of advanced text data structures on large datasets. While fast sequential algorithms exist in the RAM model [17,18], these are limited by the CPU power and RAM size of a single machine. External memory algorithms on a single machine are limited by disk [4,3,9], and often have long running times due to mostly sequential computation or limited I/O bandwidth.

Most suffix array construction algorithms focus only on sequential computation models. However, while the volume of data is increasing, the speed of individual CPU cores is not. This leaves us no choice but to consider shared-memory parallelism and distributed cluster computation to gain considerable speedups in the future.

Most suffix array construction algorithms (SACAs) employ a subset of three basic suffix sorting principles: *prefix doubling*, *recursion* and *inducing* [20]. The last type, *inducing*, is the basis for the fastest *sequential* suffix array construction algorithms [17,18], but yields only well to parallelization for small alphabets [13], and does not appear to be a promising approach for distributed environments. Recently, a fast distributed prefix doubling implementation using MPI has been presented [6]. While they report high speeds for very small inputs, we could not successfully run their implementation on large inputs. Furthermore, using hundreds of high performance machines for small inputs is not dollar-cost-efficient.

We propose to use the big data framework *Thrill*, which supports distributed external memory algorithms for suffix sorting of large inputs.

After giving a short introduction to Thrill in Section 1.2, we provide a detailed description of our SACA implementations in Section 2. Section 2.1 considers multiple variants of prefix doubling algorithms, and Section 2.2 discusses the recursive difference cover algorithms DC3 and DC7.

## 1.1 Related Work

There exists numerous work on sequential SACAs, see [20,5] for two overview articles. Research in this area is still active, as just this year another theoretically optimal SACA has been presented that combines ideas used in prefix doubling and inducing [1]. Kärkkäinen et al. [11,10] presented a linear time SACA, the so called DC3 algorithm, that works well in multiple advanced models of computation such as external memory and also parallel and distributed environments. Kulla and Sanders showed the scalability of the DC3 algorithm in a distributed environment [12]. More recently, Flick and Aluru presented an implementation of a prefix doubling algorithm in MPI that can also compute the longest common prefix array [6]. SACAs have also been considered in external memory, where in theory the DC3 algorithm [11] is optimal. Dementiev et al. [4] compared multiple implementations of prefix doubling and DC3 for external memory in practice. Lately, Kärkkäinen et al. [9,8] presented two different external memory SACAs.

Related to SACAs are construction algorithms for the suffix tree and the Burrows-Wheeler transform (BWT), which are often used in Bioinformatics. In this domain one can, however, make special assumptions such as that the input text is fairly random (like DNA), or that one wishes to compress multiple very similar texts (like multiple genome sequences of the same species). These practical assumptions yield suffix sorting implementations tailored to their applications, like straight-forward parallel radix sort [16,22] or merging of multiple BWTs generated in parallel [21]. On general text these implementations, however, have super-linear theoretical running time.

## 1.2 A Short Introduction into Thrill

We implemented five suffix array construction algorithms using the distributed big data batch computation framework *Thrill* [2]. Thrill works with *distributed immutable arrays* (DIAs) storing tuples. Items in DIAs cannot be accessed directly, instead Thrill provides a rich set of DIA operations which can be used to transform DIAs (we use and describe only a subset of the operations Thrill provides). Each DIA operation can be instantiated with appropriate user-defined functions for constructing complex algorithms.

**Filter( $f$ )** takes a  $\text{DIA}\langle A \rangle X$  and a function  $f: A \rightarrow \text{bool}$ , and returns the  $\text{DIA}\langle A \rangle$  containing  $[x \in X \mid f(x)]$  within which the order of items is maintained.

**Map( $f$ )** applies the function  $f: A \rightarrow B$  to each item in the input  $\text{DIA}\langle A \rangle X$ , and returns a  $\text{DIA}\langle B \rangle Y$  with  $Y[i] = f(X[i])$  for all  $i = 0, \dots, |X| - 1$ .

**Window<sub>k</sub>(w)** and **FlatWindow<sub>k</sub>(w')** takes an input  $\text{DIA}\langle A \rangle X$  and a window function  $w: \mathbb{N}_0 \times A^k \rightarrow B$ . The operation scans over  $X$  with a window of size  $k$  and applies  $w$  once to each set of  $k$  consecutive items from  $X$  and their index in  $X$ . The final  $k - 1$  indexes with less than  $k$  consecutive items are delivered to  $w$  as partial windows padded with sentinel values. The result of all invocations of  $w$  is returned as a  $\text{DIA}\langle B \rangle$  containing  $|X|$  items in the order.

**FlatWindow** is a variant of **Window** which takes a input  $\text{DIA}\langle A \rangle X$  and a window function  $w': \mathbb{N}_0 \times A^k \rightarrow \text{list}(B)$ . The only difference compared to **Window** is, that  $w'$  can *emit* zero or more items that are concatenated in the resulting  $\text{DIA}\langle B \rangle$  in the order they are emitted.<sup>3</sup>

**PrefixSum(s)** Given an input  $\text{DIA}\langle A \rangle X$  and an associative operation  $s: A \times A \rightarrow A$  (by default  $s = +$ ), **PrefixSum** returns a  $\text{DIA}\langle A \rangle Y$  such that  $Y[0] = X[0]$  and  $Y[i] = s(Y[i - 1], X[i])$  for all  $i = 1, \dots, |X| - 1$ .

**Sort(c)** sorts an input  $\text{DIA}\langle A \rangle X$  with respect to a less-comparison function  $c: A \times A \rightarrow \text{bool}$ . If **Sort** is called without a comparison function, we assume the tuples are compared component-wise with the first component being most significant, the second component the second most significant, and so on.

**Merge(X<sub>1</sub>, ..., X<sub>n</sub>, c)** Given a set of sorted  $\text{DIA}\langle A \rangle$ s  $X_1, \dots, X_n$  and a less-comparison function  $c: A \times A \rightarrow \text{bool}$ , **Merge** returns  $\text{DIA}\langle A \rangle Y$  that contains all tuples of  $X_1, \dots, X_n$  and is sorted with respect to  $c$ . If **Merge** is called without a comparison function we compare the tuples component-wise (see **Sort**).

**Union(X<sub>1</sub>, ..., X<sub>n</sub>)** Given a set of  $\text{DIA}\langle A \rangle$ s  $X_1, \dots, X_n$ , **Union** returns  $\text{DIA}\langle A \rangle Y = \bigcup_{i=1}^n X_i$  containing all items of the input in an arbitrary order.

**Zip(X<sub>1</sub>, ..., X<sub>n</sub>, f)** Given a set of  $\text{DIA}$ s  $X_1, \dots, X_n$  of type  $A_1, \dots, A_n$  of equal size ( $|X_1| = \dots = |X_n|$ ) and a function  $f: A_1 \times \dots \times A_n \rightarrow B$ , **Zip** returns  $\text{DIA}\langle B \rangle Y$  with  $Y[i] = f(X_1[i], \dots, X_n[i])$  for all  $i = 0, \dots, |X_1| - 1$ .

**ZipWithIndex(f)** Given an input  $\text{DIA}\langle A \rangle X$  and a function  $f: (\mathbb{N}_0, A) \rightarrow B$ , **ZipWithIndex** returns  $\text{DIA}\langle B \rangle Y$  with  $Y[i] = f(i, X[i])$  for all  $i = 0, \dots, |X| - 1$

**Max(c)** Given an input  $\text{DIA}\langle A \rangle X$ , **Max** returns the maximum item  $m = \max_c X$  with respect to a less-comparison function  $c: A \times A \rightarrow \text{bool}$ . By default (if **Max** is called without a comparison function) the tuples are compared component-wise (see **Sort**).

**Size()** Given an input  $\text{DIA}\langle A \rangle X$ , **Size** returns the number of items in  $X$ , i.e.,  $|X|$ .

**Thrill** applies chains of functions (*method chaining*) to a  $\text{DIA}$ , e.g., if we have a  $\text{DIA}\langle \mathbb{N}_0 \rangle N = \{0, 1, 2, \dots, 9\}$  and want to compute the prefix sum of all odd elements, then we write  $N.\text{Filter}(a \mapsto (a \bmod 2 = 1)).\text{PrefixSum}()$ . Using chaining, the operations form a data-flow style graph of  $\text{DIA}$  operations. Drawings of this graph help to give a visual impression of the dependencies between the operations. As the data-flow drawings in this paper are generated from our actual **Thrill** implementation, they contain some additional nodes. These are only needed for performance (**Cache**) and due to the way **Thrill** code is chained (**Collapse**).

<sup>3</sup> We say the items are *emitted*, as in other  $\text{DIA}$  operations more than one item can be created per call of the function  $w$  to the output  $\text{DIA}\langle B \rangle$ , while *return* exits the function.

## 2 Scalable Suffix Array Construction Algorithms

In this section, we describe the suffix array construction algorithms that we have implemented in Thrill. First we describe algorithms based on prefix doubling, and then we present two implementations based on recursion. SACAs based on inducing do not appear to be a promising approach in a distributed environment.

Given is a text  $T$  of length  $|T| = n$  over an alphabet  $\Sigma$ . We call the substring  $T[i, n]$  the  $i$ -th *suffix* of  $T$ . The *suffix array* (SA) for  $T$  is a permutation of  $[0, n)$  such that  $T[\text{SA}[i], n] \leq_{\text{lex}} T[\text{SA}[j], n]$  for all  $0 \leq i \leq j < n$ . The inverse permutation of SA is called the *inverse suffix array* (ISA) and the *lexicographic rank* of the  $i$ -th suffix is  $\text{ISA}[i]$ . While the ranks of all suffixes are distinct, we will often use the notion of a *lexicographic name*. Lexicographic names are representatives of suffixes which need not be distinct but do respect the lexicographic ordering, i.e.,  $n_i$  and  $n_j$  are lexicographic names of two suffixes  $T[i, n] < T[j, n]$  iff  $n_i \leq n_j$ .

### 2.1 Prefix Doubling Algorithms

The goal of a *prefix doubling* algorithm is to give each suffix of  $T$  a lexicographic name such that the name corresponds to the rank of the suffix in the (partial) SA. The names are computed using prefixes of length  $2^k$  of the suffixes for  $k = 1, \dots, \lceil \log_2 |T| \rceil$ . During each step, we double the length of these prefixes (hence the name of this type of algorithm). We can compute the name for the prefix  $T[i, i + 2^k]$  using the already computed names of the prefixes  $T[i, i + 2^{k-1}]$  and  $T[i + 2^{k-1}, i + 2^k]$ .

Algorithm 1 describes the basic structure of the prefix doubling algorithms presented in this section. The corresponding data-flow graph is shown in Figure 1. The whole algorithm requires one DIA  $N$  storing tuples and one DIA  $S$  storing triples. Initially,  $S$  contains the triples  $(i, T[i], T[i + 1])$  for all  $i = 0, \dots, n - 1$  where we assume that  $T[n] = \$$ , see Line 2. These triples contain a text position and the *name pair* for that position, i.e, the two names that are required to compute the new name for the suffix starting at the text position. Next in Line 4, we sort  $S$  with respect to the

---

#### Algorithm 1: Generic Prefix Doubling algorithm.

---

```

1 function PrefixDoubling( $T \in \text{DIA}(\Sigma)$ )
2    $S := T.\text{Window}_2((i, [t_0, t_1]) \mapsto (i, t_0, t_1))$  // Create initial triples  $(i, T[i], T[i + 1])$ .
3   for  $k := 1$  to  $\lceil \log_2 |T| \rceil - 1$  do
4      $S := S.\text{Sort}((i, r_0, r_1) \text{ by } (r_0, r_1))$  // Sort triples by name pair.
5      $N := S.\text{FlatWindow}_2((i, [a, b]) \mapsto \text{CmpName}(i, a, b))$  // Map to names 0 or  $i$ .
6     if  $N.\text{Filter}((i, r) \mapsto (r = 0)).\text{Size}() = 1$  then // If all names distinct, then
7        $\lfloor$  return  $N.\text{Map}((i, r) \mapsto i)$  // return names as suffix array,
8      $N := N.\text{PrefixSum}((i, r), (i', r') \mapsto (i', \max(r, r')))$  // else calculate new names
9      $S := \text{Generate new name pairs using } N$  // and run next refinement iteration.

```

---

---

**Algorithm 2:** Identifications of suffix array intervals.

---

```

1 function CmpName( $j \in \mathbb{N}_0, (i, r_0, r_1), (i', r'_0, r'_1) \in N$ )
2   if  $j = 0$  then
3     emit  $(i, 0)$  // First DIA item has no offset.
4     emit  $\begin{cases} (i', j) & \text{if } (r_0, r_1) \neq (r'_0, r'_1), \text{ // Add sentinel if rank pairs alter.} \\ (i', 0) & \text{otherwise. // } T[i, n) \text{ and } T[i', n) \text{ get the same new name.} \end{cases}$ 

```

---

name pair as we know that the names correspond to the ranks of the suffixes. Now we prepare the computation of the new names using the functions `CmpName()` that takes the current position  $i$  in  $S$  and the items  $S[i]$  and  $S[i + 1]$  as input and emits a tuple consisting of a text position and a new name, see Algorithm 2. We know that the suffixes are sorted with respect to their name pairs. Therefore, we can scan  $S$  and mark every position where the name pair differs from its predecessor. `CmpName()` marks these non-unique names pairs by giving them the name 0. All unique names pairs get a name equal to their current position in  $S$ . If there is only one suffix with name 0 we know that all names differ and that we have finished the computation, see Line 6. Otherwise, we can use a the DIA operation `PrefixSum()` to set the name of the tuple to the largest preceding name, i.e., the the new name which is unique if the name was not 0 and the preceding name is not 0, see Line 8. Now each suffix has a new, more refined name. The next step (see Line 9) is to identify the ranks of the suffixes required for the next doubling step. During the  $k$ -th doubling step, we fill  $S$  with one triple for each index  $i = 0, \dots, |T| - 1$  that contains the current name of the suffix at position  $i$  and the current name of the suffix at position  $i + 2^{k-1}$ . This is also the step, where the prefix doubling algorithms presented here differ. Next, we show two different approaches to compute the name pairs for the next refining iteration.

**Prefix Doubling using Sorting.** In the seminal suffix array paper by Manber and Myers [15], the presented SACA was a prefix doubling algorithm using sorting. This idea was refined by Dementiev et al. [4] who presented an external memory SACA that we adapted to Thrill. The idea is to compute the new name pairs by sorting the old names with respect to the starting position of the suffix, see Algorithm 3. We make use of the fact that during each iteration we know for each suffix the suffix whose current name is required to compute the new, refined name. Hence, we can sort the tuples containing the starting positions of the suffixes and their current name in such a way that if there is another name required for a name pair, then it is the name of the succeeding tuple, see Line 2. To do so, we use the following less-comparator  $\langle_{\text{op}}^k : (\mathbb{N}_0, \mathbb{N}_0) \times (\mathbb{N}_0, \mathbb{N}_0) \rightarrow \text{bool}$  (see Equation 1) in Algorithm 3:

$$(i, r) \langle_{\text{op}}^k (i', r') = \begin{cases} i \text{ div } 2^k < i' \text{ div } 2^k & \text{if } i \equiv i' \pmod{2^k}, \\ i \bmod 2^k < i' \bmod 2^k & \text{otherwise.} \end{cases} \quad (1)$$

---

**Algorithm 3:** Prefix Doubling using sorting.

---

```
1 function PrefixDoublingSorting( $k \in \mathbb{N}_0$ )
2    $N := N.\text{Sort}(<_{\text{op}}^k)$  // Sort such that names required for renaming are consecutive.
3    $S := N.\text{Window}_2 \left( (j, [(i, r_0, r_1), (i', r'_0, r'_1)]) \mapsto \begin{cases} (i, r_0, r'_0) & \text{if } i + 2^k = i', \\ (i, r_0, 0) & \text{otherwise.} \end{cases} \right)$ 
```

---

---

**Example 4:** Example of prefix doubling using sorting in Thrill.

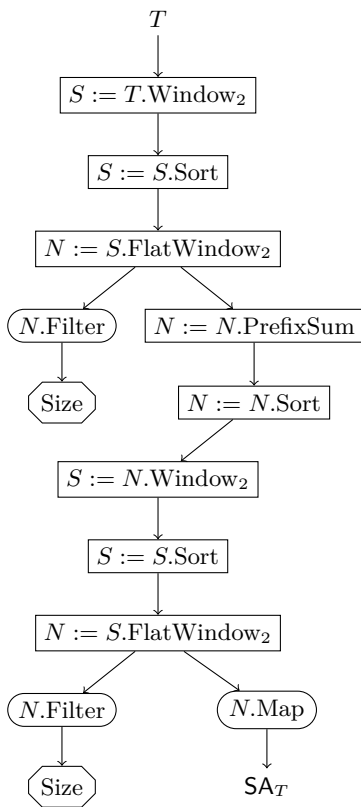
---

```
1  $T = [\text{b, d, a, c, b, d, a, c, b}]$ 
2  $S = [(0, \text{d, b}), (1, \text{b, a}), (2, \text{a, c}), (3, \text{c, b}), (4, \text{b, d}), (5, \text{d, a}), (6, \text{a, c}), (7, \text{c, b}), (8, \text{b, \$})]$  // 1.2
3  $\mathbf{k = 1}$  // 1.3
4  $S = [(2, \text{a, c}), (6, \text{a, c}), (8, \text{b, \$}), (0, \text{b, d}), (4, \text{b, d}), (3, \text{c, b}), (7, \text{c, b}), (1, \text{d, a}), (5, \text{d, a})]$  // 1.4
5  $N = [(2, 0), (6, 0), (8, 2), (0, 3), (4, 0), (3, 5), (7, 0), (1, 7), (5, 0)]$  // 1.5
6 4 items with rank 0 // 1.6
7  $N = [(2, 0), (6, 0), (8, 2), (0, 3), (4, 3), (3, 5), (7, 5), (1, 7), (5, 7)]$  // 1.8
8  $N = [(0, 3), (2, 0), (4, 3), (6, 0), (8, 2), (1, 7), (3, 5), (5, 7), (7, 5)]$  // 3.2
9  $S = [(0, 3, 0), (2, 0, 3), (4, 3, 0), (6, 0, 2), (8, 2, 0), (1, 7, 5), (3, 5, 7), (5, 7, 5), (7, 5, 0)]$  // 3.3
10  $\mathbf{k = 2}$  // 1.3
11  $S = [(6, 0, 2), (2, 0, 3), (8, 2, 0), (0, 3, 0), (4, 3, 0), (7, 5, 0), (3, 5, 7), (1, 7, 5), (5, 7, 5)]$  // 1.4
12  $N = [(6, 0), (2, 1), (8, 2), (0, 3), (4, 0), (7, 5), (3, 6), (1, 7), (5, 0)]$  // 1.5
13 2 items with rank 0 // 1.6
14  $N = [(6, 0), (2, 1), (8, 2), (0, 3), (4, 3), (7, 5), (3, 6), (1, 7), (5, 7)]$  // 1.8
15  $N = [(0, 3), (4, 3), (8, 2), (1, 7), (5, 7), (2, 1), (6, 0), (3, 6), (7, 5)]$  // 3.2
16  $S = [(0, 3, 3), (4, 3, 2), (8, 2, 0), (1, 7, 7), (5, 7, 0), (2, 1, 0), (6, 0, 0), (3, 6, 5), (7, 5, 0)]$  // 3.3
17  $\mathbf{k = 3}$  // 1.3
18  $S = [(6, 0, 0), (2, 1, 0), (8, 2, 0), (4, 3, 2), (0, 3, 3), (7, 5, 0), (3, 6, 5), (5, 7, 0), (1, 7, 7)]$  // 1.4
19  $N = [(6, 0), (2, 1), (8, 2), (4, 3), (0, 4), (7, 5), (3, 6), (5, 7), (1, 8)]$  // 1.5
20 1 item with rank 0 // 1.6
21 Result:  $[6, 2, 8, 4, 0, 7, 3, 5, 1]$  // 1.7
```

---

After sorting using the  $<_{\text{op}}^k$ -comparator, we need to ensure that two consecutive names are the ones required to compute the new name, since the required name may not exist due to the length of the text. This occurs during the  $k$ -th iteration for each suffix beginning at a text position greater than  $n - 2^k$ . In this case we use the sentinel name 0 which compares smaller than any valid name, see Line 3. In both cases, we return one triple for each position, consisting of a text position, the current name of the suffix beginning at that position and the name of the suffix  $2^k$  positions to the right (if it exists and 0 otherwise).

Now we give an example of prefix doubling using sorting in Thrill, see Example 4. We compute the suffix array of the text  $T = \text{bdacbdacb}$ . The comment at the end of each line refers to the line of code responsible for the change from the previous line where  $x.y$  denotes line  $y$  in Algorithm  $x$ .



**Fig. 1.** DIA data-flow graph of two iterations of prefix doubling with sorting.

**Prefix Doubling using the Inverse Suffix Array.** During the  $k$ -th doubling step, we compute a name for each suffix and hence for each position of the text. Algorithm 5 describes how we obtain the rank of the required suffixes using the inverse suffix array. This approach has been considered in a distributed environment [6] and is based on the work of Larsson and Sadakane [14] who proposed to use the inverse suffix for prefix doubling. If we sort the names based on their position in the text, we get the *partial* inverse suffix array (partial, as the inverse suffix array does not necessarily contain the final position of all suffixes in the SA). Using this partial inverse suffix array, we can get the current rank of each suffix by its text position. For each position  $i$ , we need the rank of the  $(i + 2^k)$ -th suffix. To get this rank, we scan over the DIA with a window of width  $2^k$ , i.e., the same as shifting the partial inverse suffix array by  $2^k$  positions and appending 0s until its length is  $|T|$  again.

---

**Algorithm 5:** Prefix Doubling using the inverse suffix array.

---

```
1 function PrefixDoublingISA( $k \in \mathbb{N}_0$ )
2    $N := N.\text{Sort}(\{(i, r) \text{ by } i\})$  // Compute partial ISA.
3    $S := N.\text{Window}_{2^{k+1}}\left(j, [(i, r), \dots, (i', r')]\right) \mapsto \begin{cases} (i, r, r') & \text{if } j + 2^k < |T|, \\ (i, r, 0) & \text{otherwise.} \end{cases}$ 
```

---

---

**Example 6:** Example of prefix doubling using the inverse suffix array in Thrill.

---

```
1  $T = [\text{b, d, a, c, b, d, a, c, b}]$ 
2  $S = [(0, \text{b, d}), (1, \text{d, a}), (2, \text{a, c}), (3, \text{c, b}), (4, \text{b, d}), (5, \text{d, a}), (6, \text{a, c}), (7, \text{c, b}), (8, \text{b, } \$)]$  // 1.2
3  $k = 1$  // 1.3
4  $S = [(2, \text{a, c}), (6, \text{a, c}), (8, \text{b, } \$), (0, \text{b, d}), (4, \text{b, d}), (3, \text{c, b}), (7, \text{c, b}), (1, \text{d, a}), (5, \text{d, a})]$  // 1.4
5  $N = [(2, 0), (6, 0), (8, 2), (0, 3), (4, 0), (3, 5), (7, 0), (1, 7), (5, 0)]$  // 1.5
6 4 items with rank 0 // 1.6
7  $N = [(2, 0), (6, 0), (8, 2), (0, 3), (4, 3), (3, 5), (7, 5), (1, 7), (5, 7)]$  // 1.8
8  $N = [(0, 3), (1, 7), (2, 0), (3, 5), (4, 3), (5, 7), (6, 0), (7, 5), (8, 2)]$  // 5.2
9  $S = [(0, 3, 0), (1, 7, 5), (2, 0, 3), (3, 5, 7), (4, 3, 0), (5, 7, 5), (6, 0, 2), (7, 5, 0), (8, 2, 0)]$  // 5.3
10  $k = 2$  // 1.3
11  $S = [(6, 0, 2), (2, 0, 3), (8, 2, 0), (0, 3, 0), (4, 3, 0), (7, 5, 0), (3, 5, 7), (1, 7, 5), (5, 7, 5)]$  // 1.4
12  $N = [(6, 0), (2, 1), (8, 2), (0, 3), (4, 0), (7, 5), (3, 6), (1, 7), (5, 0)]$  // 1.5
13 2 items with rank 0 // 1.6
14  $N = [(6, 0), (2, 1), (8, 2), (0, 3), (4, 3), (7, 5), (3, 6), (1, 7), (5, 7)]$  // 1.8
15  $N = [(0, 3), (1, 7), (2, 1), (3, 6), (4, 3), (5, 7), (6, 0), (7, 5), (8, 2)]$  // 5.2
16  $S = [(0, 3, 3), (1, 7, 7), (2, 1, 0), (3, 6, 5), (4, 3, 2), (5, 7, 0), (6, 0, 0), (7, 5, 0), (8, 2, 0)]$  // 5.3
17  $k = 3$  // 1.3
18  $S = [(6, 0, 0), (2, 1, 0), (8, 2, 0), (4, 3, 2), (0, 3, 3), (7, 5, 0), (3, 6, 5), (5, 7, 0), (1, 7, 7)]$  // 1.4
19  $N = [(6, 0), (2, 1), (8, 2), (4, 3), (0, 4), (7, 5), (3, 6), (5, 7), (1, 8)]$  // 1.5
20 1 item with rank 0 // 1.6
21 Result:  $[6, 2, 8, 4, 0, 7, 3, 5, 1]$  // 1.7
```

---

Again, we give an example of the algorithm for the same text as before, see Example 6. Also, the comments refer to the algorithm and line responsible for the change as in the previous example.

**Prefix Doubling with Discarding.** In the algorithms described above, we always sort and consider all suffixes (name pairs) for the following renaming. Even though some of them are already at their correct position, i.e., have a unique name. Now, we present an algorithm which extends the prefix doubling algorithm using sorting such that only the following suffixes are sorted: suffixes that do not yet have a unique name and suffixes that have an unique name but are required to compute a name pair (for a suffix that does not yet have a unique name). All other suffixes are *discarded* and are not considered for the computation anymore. This technique has also been considered for external memory suffix array construction [4].



---

**Algorithm 7: Prefix Doubling with Discarding.**


---

```

1 function PrefixDoublingDiscarding( $T \in \text{DIA}(\Sigma)$ )
2    $S := T.\text{Window}_2((i, [t_0, t_1]) \mapsto (i, t_0, t_1))$  // Create initial triples  $(i, T[i], T[i + 1])$ .
3    $S := S.\text{Sort}((i, r_0, r_1) \text{ by } (r_0, r_1))$  // Sort triples by name pairs.
4    $N := S.\text{FlatWindow}_2((i, [a, b]) \mapsto \text{CmpName}(i, a, b))$  // Map names to 0 or  $i$ .
5    $N := N.\text{PrefixSum}(((i, r), (i', r')) \mapsto (i', \max(r, r')))$  // Calculate initial names.
6   for  $k := 1$  to  $\lceil \log_2 |T| \rceil$  do
7      $P := N.\text{FlatWindow}_3((i, [a, b, c]) \mapsto \text{Unique}(a, b, c, i))$  // Compute states of items.
8      $P := \text{Union}(P, U).\text{Sort}(<_{\text{op}}^k)$  // Concatenate undiscarded items and sort them.
9      $P := P.\text{FlatWindow}_3((i, [a, b, c]) \mapsto \text{NPairs}(i, a, b, c, k))$  // Compute new name
10     $D' := P.\text{Filter}((i, r_0, r_1, s) \mapsto (s = \mathbf{d}))$  // pairs and update state. Then find and
11     $D := \text{Union}(D, D').\text{Map}((i, r_0, r_1, s) \mapsto (i, a, r_0))$  // store newly discarded items.
12     $U' := P.\text{Filter}((i, r_0, r_1, s) \mapsto (s = \mathbf{u}))$  // Separate the already unique items and the
13     $U := U'.\text{Map}((i, r_0, r_1, s) \mapsto (i, r_0, s))$  // items that still need to be sorted. Former
14     $I' := P.\text{Filter}((i, r_0, r_1, s) \mapsto (s = \mathbf{n}))$  // are only needed to compute the name pairs
15     $I := I'.\text{Map}((i, r_0, r_1, s) \mapsto (i, r_0, r_1))$  // and stored in  $U$ . Latter are stored in  $I$ .
16    if  $I.\text{Size}() = 0$  then
17      return  $D.\text{Sort}((i, r) \text{ by } r).\text{Map}((i, r) \mapsto i)$  // If all items are unique return SA.
18     $M := I.\text{FlatWindow}_2((i, [a, b]) \mapsto \text{NameDiscarding}(i, a, b))$  // Form names that
19     $M := M.\text{PrefixSum}(((i, r_0, r_1, r_2), (i', r'_0, r'_1, r'_2)) \mapsto (i', \max(r'_0, r'_0), \max(r'_1, r'_1), r'_2))$ 
20     $N := M.\text{Map}((i, r_0, r_1, r_2) \mapsto (i, r_2 + (r_1 - r_0)))$  // comply with the old names.

```

---

Initially, Algorithm 7 behaves like the generic prefix doubling algorithm (see Figure 2 for the data-flow graph). We compute name pairs for consecutive text positions (line 2) and compute the names for all suffixes the same way we do in the generic algorithm. Next, we add a *state* to the triples  $(i, r_1, r_2)$ , i.e., creating 4-tuples  $(i, r_1, r_2, s)$ , indicating whether a name pair is *unique* (**u**) or *not unique* (**n**) (see function Unique, Algorithm 8). All 4-tuples that are unique do not need a new name but they may still be required to compute the new name of another suffix. Hence we add a third state, a 4-tuple that is unique gets the state *discarded* (**d**) if it is not required for the computation of a different name. Those tuples can easily be identified by looking at three consecutive tuples after they have been sorted using the less-comparator described in Equation 1. Let  $a = (i, r_1, r_2, s)$ ,  $b = (i', r'_1, r'_2, s')$  and  $c = (i'', r''_1, r''_2, s'')$  be three continuous tuples with  $s''$  being unique. If either  $s$  or  $s'$  is unique, then  $c$  can be discarded because both  $a$  and  $b$  will get a unique name pair during this iteration. Otherwise (if  $s$  and  $s'$  are not unique) then  $c$  cannot be discarded as  $a$  will not get a unique name pair during this iteration and we require the name of  $c$  during the next iteration to compute the name pair (see Function NPairs, Algorithm 8, Lines 7–18). While computing the final state we also create the new name pairs required for the new name if the state is not unique, as otherwise the name is final.

Since we do not consider all tuples during the course of Algorithm 7 we need to change the renaming based on the name pairs. Up to now, we were able to give

---

**Algorithm 8:** Prefix Doubling with Discarding (Additional Functions)

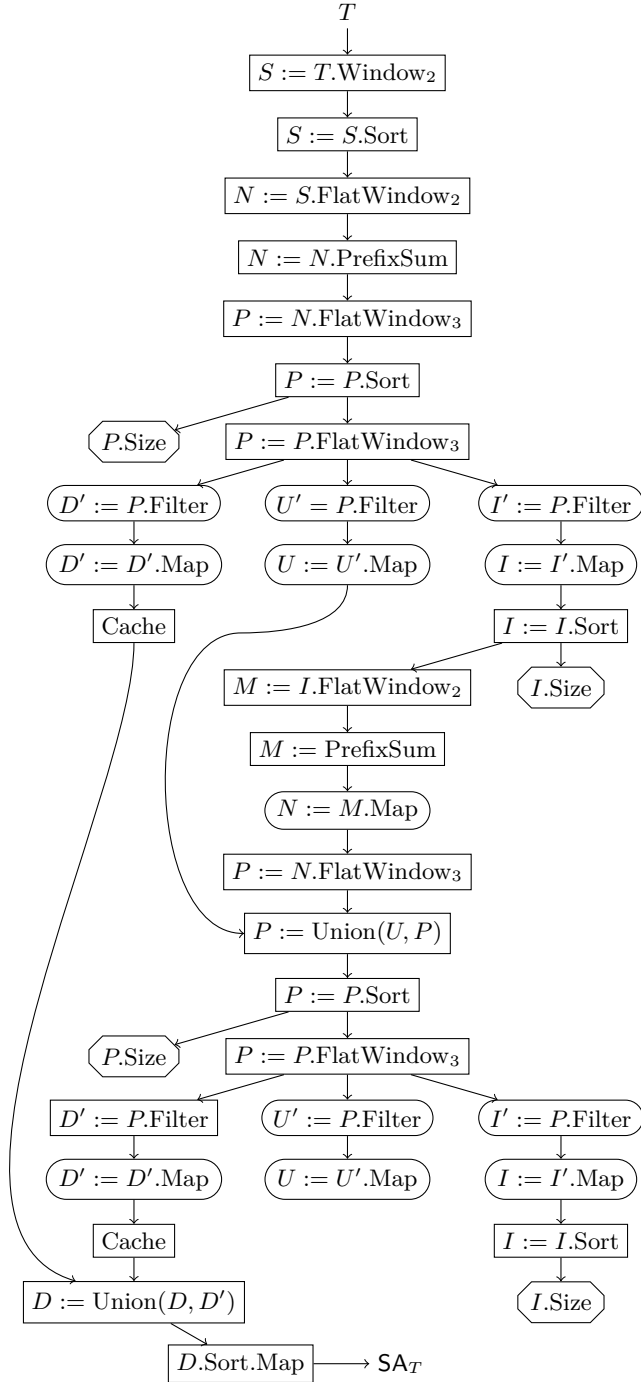
---

```

1 function Unique( $j \in \mathbb{N}_0, (i, r), (i', r'), (i'', r'') \in N$ )
2   if  $j = 0$  then
3     emit  $\begin{cases} (i, r, \mathbf{u}) & \text{if } r \neq r', \text{ // First item is unique} \\ (i, r, \mathbf{n}) & \text{otherwise. // if its ranks differ from its successor.} \end{cases}$ 
4   else if  $j + 2 = l$  then
5     emit  $\begin{cases} (i'', r'', \mathbf{u}) & \text{if } r' \neq r'', \text{ // Final item is unique} \\ (i'', r'', \mathbf{n}) & \text{otherwise. // if its ranks differs from its precursor.} \end{cases}$ 
6   emit  $\begin{cases} (i', r', \mathbf{u}) & \text{if } r \neq r' \text{ and } r' \neq r'', \text{ // An item is} \\ (i', r', \mathbf{n}) & \text{otherwise. // unique if its ranks are unique.} \end{cases}$ 
7 function NPairs( $j \in \mathbb{N}_0, (i, r, s), (i', r', s'), (i'', r'', s'') \in P, k \in \mathbb{N}_0$ )
8   if  $j = 0$  then
9     emit  $\begin{cases} (i, r, 0, \mathbf{d}) & \text{if } s = \mathbf{u}, \text{ // The first two items can be discarded} \\ (i', r', 0, \mathbf{d}) & \text{if } s' = \mathbf{u}. \text{ // if they are unique. Emit } \leq 2 \text{ items.} \end{cases}$ 
10  else if  $j + 2 = l$  then
11    if  $s' = \mathbf{n}$  then
12      emit  $\begin{cases} (i', r', r'', \mathbf{n}) & \text{if } i' + 2^k = i'', \text{ // If the last two items of the} \\ (i', r', 0, \mathbf{n}) & \text{otherwise. // DIA are undecided, then we need} \end{cases}$ 
13    if  $s'' = \mathbf{n}$  then
14      emit  $(i'', r'', 0, \mathbf{n})$  // to fuse the ranks required for renaming.
15    if  $s = \mathbf{n}$  then
16      emit  $\begin{cases} (i, r, r', \mathbf{n}) & \text{if } i + 2^k = i', \text{ // The ranks for renaming are} \\ (i, r, 0, \mathbf{n}) & \text{otherwise. // consecutive and fused accordingly.} \end{cases}$ 
17    if  $s'' = \mathbf{u}$  then
18      emit  $\begin{cases} (i'', r'', 0, \mathbf{d}) & \text{if } s = \mathbf{u} \text{ or } s' = \mathbf{u}, \text{ // Unique items are dis-} \\ (i'', r'', 0, \mathbf{u}) & \text{otherwise. // carded if uncalled-for in future renaming.} \end{cases}$ 
19 function NameDiscarding( $j \in \mathbb{N}_0, l \in \mathbb{N}_0, (i, r_0, r_1), (i', r'_0, r'_1) \in I$ )
20   if  $j = 0$  then
21     emit  $(i, 1, 1, r_0)$  // The new names must comply with the old ones.
22   emit  $\begin{cases} (i', j + 2, j + 2, r'_0) & \text{if } r_0 \neq r'_0 \text{ and } r_1 \neq r'_1, \text{ // The first rank de-} \\ (i', 1, j + 2, r'_0) & \text{else if } r_0 = r'_0, \text{ // terminates the group and new} \\ (i', 1, 1, r'_0) & \text{otherwise. // names are consistent within groups.} \end{cases}$ 

```

---



**Fig. 2.** DIA data-flow graph of two iterations of prefix doubling with discarding.

names starting at 0 and continue based on the (preliminary) position in SA. If we discard tuples this approach is not feasible any more as we need to consider the names of already discarded tuples. During the  $k$ -th iteration, all suffixes that do not have a unique name form consecutive intervals in SA. Within these intervals all suffixes that cannot be distinguished by their first  $2^k$  characters share the same name. These names are extended, i.e., increased such that the new name is always at least as great as the previous name and greater than the rank of the first preceding suffix that can be distinguished using the first  $2^k$  characters of the suffixes (lines 18–20). At the beginning of the next iteration we add all unique names to the new names and check if they can be discarded. As soon as all names are unique (line 16) we know that all names have been discarded and can compute SA by sorting the discarded tuples by their names (line 17).

**Prefix Quadrupling.** In the prefix doubling algorithms described above, during the  $k$ -th doubling step, we consider substrings of length  $2^k$ . This can be generalized to substrings of length  $a^k$  for any  $a \in \mathbb{N}_0$  with  $a > 1$ . The prefix doubling algorithms using sorting are I/O optimal for 5-tuples in external memory and in practice using 4-tuples, i.e., prefix quadrupling has the advantage that less memory is required for storing the tuples and that the I/O-volume is just 1.5% worse compared to prefix quintupling [4]. The change within the algorithms can be kept to a minimum as we just require rank quadruples instead of rank pairs. Also, the comparison and the computation of the new names have to be adapted accordingly.

## 2.2 Difference Cover Algorithms – DC3 and DC7, aka skew3 and skew7

In 2003, the DC3 aka skew3 suffix sorting algorithm and its generalization, DCX and skewX, was proposed by Kärkkäinen, Sanders, and Burkhardt [10,11]. They employ recursion on a subset of the suffixes to reach linear running time in the sequential RAM model, which translates to sorting complexity in the external memory and PRAM models. While the reference implementation by the authors is in the sequential RAM model, the algorithms were later implemented for external memory [4,23], and DC3 was implemented for distributed memory using MPI [12].

The DCX algorithms are based on scanning, sorting, and merging, and hence are asymptotically optimal in many models provided optimal theoretical base algorithms. As Thrill supplies all of these base algorithms as scalable distributed algorithmic primitives, implementing DCX is a natural choice.

The key notion of DCX is to recursively calculate the ranks of suffixes in only a *difference cover* of the original text. A set  $D \subseteq \mathbb{N}_0$  is a difference cover for  $n \in \mathbb{N}_0$ , if  $\{(i - j) \bmod n \mid i, j \in D\} = \{0, \dots, n - 1\}$ . Examples of difference covers are  $D_3 = \{1, 2\}$  for  $n = 3$ ,  $D_7 = \{0, 1, 3\}$  for  $n = 7$ , and  $D_{13} = \{0, 1, 3, 9\}$  for  $n = 13$ . In general, a difference cover of size  $\mathcal{O}(\sqrt{n})$  can be calculated for any  $n$  in  $\mathcal{O}(\sqrt{n})$  time [11].

The broad steps of the DC3 algorithm are the following:

1. Calculate ranks for all suffixes starting at positions  $i$  in the difference cover  $D_3 = \{1, 2\}$ . This is done by sorting the triples  $(T[i], T[i + 1], T[i + 2])$  for  $i \in D_3$ , calculating lexicographic names, and recursively calling a suffix sorting algorithm on a reduced string of size  $\frac{2}{3}n$  if necessary. The result of step 1 are two arrays,  $R_1$  and  $R_2$ , containing the ranks of suffixes  $i = 1 \bmod 3$  and  $i = 2 \bmod 3$ .
2. Scan text  $T$ ,  $R_1$ , and  $R_2$  to generate three arrays:  $S_0$ ,  $S_1$ , and  $S_2$ , where array  $S_j$  contains one tuple for each suffix  $i$  with  $i = j \bmod 3$ . The arrays store tuples containing the two next ranks from  $R_1$  and  $R_2$  and all characters from  $T$  up to the next ranks. This is exactly the information required such that the following merge step is able to deduce the suffix array.
3. Sort  $S_0$ ,  $S_1$ , and  $S_2$  and merge them using a custom comparison function which compares the suffixes represented in the tuples using characters and ranks. Only a constant number of characters and ranks need to be accessed in each comparison. Output the suffix array using the indices stored in tuples.

The first two steps of the difference cover suffix sorting algorithms can be seen as preparation for the final merge in Step 3. Step 1 delivers ranks for all suffixes  $i \in D_3$  in  $R_1$  and  $R_2$ . In Step 2 tuples are created in  $S_0$ ,  $S_1$ , and  $S_2$  which are constructed from the recursively calculated ranks and characters from the text. The tuples are designed such that the comparison function can fully determine the final suffix array. The complete DC3 implementation in Thrill algorithm code is shown as Algorithm 9, and Example 10 shows the transcript of a run with the text  $T = \text{dbacbacbd}$ . Figure 4 shows the dataflow graph of DC7 instead of DC3, which is slightly more complex but shows the algorithmic structure better. In the algorithm code we omitted some details on padding and sentinels needed for inputs that are not a multiple of the difference cover size.

Goal of Lines 2–24 is to calculate  $R_1$  and  $R_2$  (step 1). This is done by performing the following steps:

1. Scan the text  $T$  using a FlatWindow operation and create triples  $(i, c_0, c_1, c_2)$  for all indices  $i$  in the difference cover  $D_3 = \{1, 2\}$  (lines 2–4).
2. Sort the triples as  $S$ , scan  $S$  and use a prefix sum to calculate lexicographic names  $N$  (lines 5–11). The lexicographic names are constructed in the prefix sum from 0 and 1 indicators. The value 0 is used if two lexicographic consecutive triples are equal, which means they are assigned the same lexicographic name; the value 1 increments the name in the prefix sum and assigns the unequal triple a new name.
3. Check if all lexicographic names are different by comparing the highest lexicographic name against the maximum possible (lines 12–13)
4. If all lexicographic names are different, then  $I_S$ , which contains the indexes of  $S$ , is already the suffix array of the suffixes in  $D_3$  (lines 21–24). Hence,  $R_1$  and  $R_2$  can be created directly: the suffix array  $I_S$  only needs to be inverted and split by mod3. Because Thrill’s Filter operation is performed locally, we interleave the future  $R_1$  and  $R_2$  parts using the Sort operation such that the two arrays are balanced on the distributed system after the Filter.

---

**Algorithm 9: DC3 Algorithm in Thrill.**


---

```

1 function DC3( $T \in \text{DIA}(\Sigma)$ )
2    $T_3 := T.\text{FlatWindow}_3((i, [c_0, c_1, c_2]) \mapsto \text{MakeTriples}(i, c_0, c_1, c_2))$ 
3   with function MakeTriples ( $i \in \mathbb{N}_0, c_0, c_1, c_2 \in \Sigma$ )
4     if  $i \neq 0 \pmod 3$  then emit ( $i, c_0, c_1, c_2$ )           // Make triples  $i \in D_3$ .
5    $S := T_3.\text{Sort}((i, c_0, c_1, c_2)$  by  $(c_0, c_1, c_2)$ )           // Sort triples lexicographically.
6    $I_S := S.\text{Map}((i, c_0, c_1, c_2) \mapsto i)$                      // Extract sorted indices.
7    $N' := S.\text{FlatWindow}_2((i, [p_0, p_1]) \mapsto \text{CmpTriple}(i, p_0, p_1))$  // Compare triples.
8   with function CmpTriple( $i \in \mathbb{N}_0, p_0 = (c_0, c_1, c_2), p_1 = (c'_0, c'_1, c'_2)$ ) // Emit one
9     if  $i = 0$  then emit 0 // sentinel for index 0, and 0 or 1
10    emit (if  $(c_0, c_1, c_2) = (c'_0, c'_1, c'_2)$  then 0 else 1) // depending on previous tuple.
11   $N := N'.\text{PrefixSum}()$  // Use prefix sum to calculate names.
12   $n_{\text{sub}} = \lceil 2|T|/3 \rceil, n_{\text{mod}1} = \lceil |T|/3 \rceil$  // Size of recursive problem and mod 1 part of  $T_R$ 
13  if  $N.\text{Max}() + 1 = n_{\text{sub}}$  then // If duplicate names exist, sort names back to
14     $T'_R := \text{Zip}([I_S, N], (i, n) \mapsto (i, n)).\text{Sort}((i, n)$  by  $(i \pmod 3, i \text{ div } 3)$ ) // string order
15     $\text{SA}_R := \text{DC3}(T'_R.\text{Map}((i, n) \mapsto n))$  // as  $T_1 \oplus T_2$  and call suffix sorter.
16     $I'_R := \text{SA}_R.\text{ZipWithIndex}((i, r) \mapsto (i, r))$  // Invert resulting suffix array, but
17     $I_R := I'_R.\text{Sort}((i, r)$  by  $(i \pmod n_{\text{mod}1}, i)$ ) // interleave ISA for better locality
18     $R_1 := I_R.\text{Filter}((i, r) \mapsto i < n_{\text{mod}1}).\text{Map}((i, r) \mapsto r + 1)$  // after separating ISA
19     $R_2 := I_R.\text{Filter}((i, r) \mapsto i \geq n_{\text{mod}1}).\text{Map}((i, r) \mapsto r + 1)$  // into  $R_1$  and  $R_2$ .
20  else // Else, if all names/triples are unique, then  $I_S$  is already the suffix array.
21     $R := I_S.\text{ZipWithIndex}((i, r) \mapsto (i, r))$  // Invert it to get ISA, but
22     $I_R := R.\text{Sort}((i, r)$  by  $(i \text{ div } 3, i)$ ) // interleave ISA for better locality
23     $R_1 := I_R.\text{Filter}((i, r) \mapsto i = 1 \pmod 3).\text{Map}((i, r) \mapsto r + 1)$  // after separating it
24     $R_2 := I_R.\text{Filter}((i, r) \mapsto i = 2 \pmod 3).\text{Map}((i, r) \mapsto r + 1)$  // into  $R_1$  and  $R_2$ .
25   $\hat{T}_3 := T.\text{FlatWindow}_3((i, [c_0, c_1, c_2]) \mapsto \text{MakeTriples}(i, c_0, c_1, c_2))$ 
26  with function MakeTriples ( $i \in \mathbb{N}_0, c_0, c_1, c_2 \in \Sigma$ ) // Prepare Zip with all
27    if  $i = 0 \pmod 3$  then emit ( $c_0, c_1, c_2$ ) // triples  $i \notin D_3$ .
28   $Z' := \text{Zip}([\hat{T}_3, R_1, R_2], ((c_0, c_1, c_2), (r_1, r_2)) \mapsto (c_0, c_1, c_2, r_1, r_2))$  // Pull chars and ranks
29   $Z := Z'.\text{Window}_2((i, [(z_1, z_2)]) \mapsto (i, z_1, z_2))$  // using Zip from three arrays
30   $S'_0 := Z.\text{Map}((i, (c_0, c_1, c_2, r_1, r_2), (\bar{c}_0, \bar{c}_1, \bar{c}_2, \bar{r}_1, \bar{r}_2)) \mapsto (3i + 0, c_0, c_1, r_1, r_2))$  // to make
31   $S'_1 := Z.\text{Map}((i, (c_0, c_1, c_2, r_1, r_2), (\bar{c}_0, \bar{c}_1, \bar{c}_2, \bar{r}_1, \bar{r}_2)) \mapsto (3i + 1, c_1, r_1, r_2))$  // arrays of
32   $S'_2 := Z.\text{Map}((i, (c_0, c_1, c_2, r_1, r_2), (\bar{c}_0, \bar{c}_1, \bar{c}_2, \bar{r}_1, \bar{r}_2)) \mapsto (3i + 2, c_2, r_2, \bar{c}_0, \bar{r}_1))$ 
33   $S_0 := S'_0.\text{Sort}((i, c_0, c_1, r_1, r_2)$  by  $(c_0, r_1)$ ) // representatives for each
34   $S_1 := S'_1.\text{Sort}((i, c_1, r_1, r_2)$  by  $(r_1)$ ) // suffix class.
35   $S_2 := S'_2.\text{Sort}((i, c_2, r_2, \bar{c}_0, \bar{r}_1)$  by  $(r_2)$ )
36  return Merge( $[S_0, S_1, S_2]$ , CompareDC3).Map( $(i, \dots) \mapsto i$ ) // Merge sorted
37  with function CompareDC3( $z_1, z_2$ ) // representatives to deliver final suffix array.
38     $(c_0, r_1) < (c'_0, r'_1)$  if  $z_1 = (i, c_0, c_1, r_1, r_2) \in S_0, z_2 = (i', c'_0, c'_1, r'_1, r'_2) \in S_1,$ 
39     $(c_0, r_1) < (c'_1, r'_2)$  if  $z_1 = (i, c_0, c_1, r_1, r_2) \in S_0, z_2 = (i', c'_1, r'_1, r'_2) \in S_1,$ 
40     $(c_0, c_1, r_2) < (c'_2, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 = (i, c_0, c_1, r_1, r_2) \in S_0, z_2 = (i', c'_2, r'_2, \bar{c}'_0, \bar{r}'_1) \in S_2,$ 
41     $(r_1) < (r'_1)$  if  $z_1 = (i, c_1, r_1, r_2) \in S_1, z_2 = (i', c'_1, r'_1, r'_2) \in S_1,$ 
42     $(r_1) < (r'_2)$  if  $z_1 = (i, c_1, r_1, r_2) \in S_1, z_2 = (i', c'_2, r'_2, \bar{c}'_0, \bar{r}'_1) \in S_2,$ 
43     $(r_2) < (r'_2)$  if  $z_1 = (i, c_2, r_2, \bar{c}_0, \bar{r}_1) \in S_2, z_2 = (i', c'_2, r'_2, \bar{c}'_0, \bar{r}'_1) \in S_2,$ 
44    and symmetrically if  $z_1 \in S_i, z_2 \in S_j$  with  $i > j$  .

```

---

---

**Example 10:** Example of DC3 Algorithm in Thrill.

---

```

1  $T = [d, b, a, c, b, a, c, b, d]$  // Example text  $T$ .
2  $T_3 = [(1, b, a, c), (2, a, c, b), (4, b, a, c), (5, a, c, b), (7, b, d, \$), (8, d, \$, \$)]$  // Triples  $i \in D_3$ .
3  $S = [(2, a, c, b), (5, a, c, b), (1, b, a, c), (4, b, a, c), (7, b, d, \$), (8, d, \$, \$)]$  // Sorted triples.
4  $I_S = [2, 5, 1, 4, 7, 8]$  // Indexes extracted from sorted triples.
5  $N' = [0, 0, 1, 0, 1]$  // 0/1 indicators depending if triples are unequal or equal.
6  $N = [0, 0, 1, 1, 2, 3]$  // Prefix sum of 0/1 indicators delivers lexicographic names.
7  $n_{\text{sub}} = 6, n_{\text{mod}1} = 3$  // Calculate result size directly.
8 Condition  $(N.\text{Max}() + 1 = 4) \neq (6 = n_{\text{sub}})$ , so follow recursion branch.
9  $T''_R = [(2, 0), (5, 0), (1, 1), (4, 1), (7, 2), (8, 3)]$  // Zip lexicographic names and their string
10  $T'_R = [(1, 1), (4, 1), (7, 2), (2, 0), (5, 0), (8, 3)]$  // index, and sort them to string order
11  $T_R = [1, 1, 2, 0, 0, 3]$  // to construct the recursive subproblem.
12  $\text{SA}_R = [3, 4, 0, 1, 2, 5]$  // Recursively calculate suffix array of  $T_R$ .
13  $I'_R = [(0, 3), (1, 4), (2, 0), (3, 1), (4, 2), (5, 5)]$  // Add index positions to suffix array
14  $I_R = [(0, 2), (3, 0), (1, 3), (4, 1), (2, 4), (5, 5)]$  // and sort into interleaved  $R_1$  and  $R_2$  ranks.
15  $R_1 = [3, 4, 5], R_2 = [1, 2, 6]$  // Filter  $R_1$  and  $R_2$  from  $I_R$ .
16  $\hat{T}_3 = [(d, b, a), (c, b, a), (c, b, d)]$  // Prepare triples  $i \notin D_3$ .
17  $Z' = [(d, b, a, 3, 1), (c, b, a, 4, 2), (c, b, d, 5, 6)]$  // Zip  $\hat{T}_3, R_1, R_2$  to make arrays  $S_i$ .
18  $S'_0 = [(0, d, b, 3, 1), (3, c, b, 4, 2), (6, c, b, 5, 6)]$  // Construct  $(i, c_0, c_1, r_1, r_2) \in S_0$ ,
19  $S'_1 = [(1, b, 3, 1), (4, b, 4, 2), (7, b, 5, 6)]$  //  $(i, c_1, r_1, r_2) \in S_1$ , and
20  $S'_2 = [(2, a, 1, c, 4), (5, a, 2, c, 5), (8, d, 6, \$, 0)]$  //  $(i, c_2, r_2, \bar{c}_0, \bar{r}_1) \in S_2$ 
21  $S_0 = [(3, c, b, 4, 2), (6, c, b, 5, 6), (0, d, b, 3, 1)]$  // as representatives of suffixes,
22  $S_1 = [(1, 3, b, 1), (4, 4, b, 2), (7, 5, b, 6)]$  // sort them among themselves
23  $S_2 = [(2, 1, a, c, 4), (5, 2, a, c, 5), (8, 6, d, \$, 0)]$  // such that merging delivers
24 Result:  $[2, 5, 1, 4, 7, 3, 6, 8, 0]$  // the final suffix array.

```

---

5. Otherwise, prepare a recursive subproblem  $T_R$  to calculate the ranks.

- (a) Sort the lexicographic names back into string order such that  $T_R = T_1 \oplus T_2$  where  $\oplus$  is string concatenation (line 14).  $T_1$  represents the complete text  $T$  using the lexicographic names of all triples  $i = 1 \pmod 3$ , and  $T_2$  is another complete copy of  $T$  with triples  $i = 2 \pmod 3$ . By replacing the triples with lexicographic names, the original text is reduced by  $\frac{2}{3}$ .
- (b) Recursively call any suffix sorting algorithm (e.g. DC3) on  $T_R$  (line 15).
- (c) Invert the permutation  $\text{SA}_R$  to gain ranks  $R_1$  and  $R_2$  of triples of  $T$  in  $D_3$ , again interleave  $\text{ISA}_R$  such that  $R_1$  and  $R_2$  are balanced on the workers after the Filter.

With  $R_1$  and  $R_2$  from Step 1 (lines 2–24), the objective of Step 2 is to create  $S_0$ ,  $S_1$ , and  $S_2$  in Lines 25–35. Each suffix  $i$  has exactly one representative in the array  $S_j$  where  $j = i \pmod 3$ . Its representative contains the recursively calculated ranks of the two following suffixes in the difference cover from  $R_1$  and  $R_2$ , and the characters  $T[i], T[i+1], T[i+2], \dots$  up to (but excluding) the next known rank.

For DC3 these are  $R_1[\frac{i}{3}]$ ,  $R_2[\frac{i}{3}]$ ,  $T[i]$ , and  $T[i+1]$  for a suffix  $i = 0 \pmod 3$  in  $S_0$ .  $R_1[\frac{i}{3}]$  is the rank of the suffix  $T[i+1, n)$  and  $R_2[\frac{i}{3}]$  is the rank of suffix  $T[i+2, n)$ ,

which are both in the difference cover. We write the tuple as  $(i, c_0, c_1, r_1, r_2)$  where the indexes are interpreted relative to  $i \bmod 3$ . Each suffix  $i = 1 \bmod 3$  in  $S_1$  stores  $R_1[\frac{i-1}{3}]$ ,  $R_2[\frac{i-1}{3}]$ , and  $T[i]$ , and we write the tuples as  $(i, c_1, r_1, r_2)$  where the indexes again are relative to  $i \bmod 3$ . And lastly, each suffix  $i = 2 \bmod 3$  in  $S_2$  stores  $R_1[\frac{i-2}{3} + 1]$ ,  $R_2[\frac{i-2}{3}]$ ,  $T[i]$ , and  $T[i + 1]$ , because  $R_1[\frac{i-2}{3} + 1]$  is the rank of suffix  $T[i + 2, n)$ .

In the Thrill algorithm code we construct the tuples by zipping  $R_1$ ,  $R_2$ , and triple groups from  $T$  together (line 25–29). The Zip  $Z'$  (line 28) delivers  $(c_0, c_1, c_2, r_1, r_2)$  for each index  $i = 0 \bmod 3$ . To construct the tuples in  $S_i$  two adjacent tuples need to be used because  $S_2$ 's element are taken from the next tuple. This can be done in Thrill using a Window operation of size 2. Thus to construct  $S_0$ ,  $S_1$ , and  $S_2$ , we take  $(c_0, c_1, c_2, r_1, r_2)$  for each index  $i = 0 \bmod 3$  and  $(\bar{c}_0, \bar{c}_1, \bar{c}_2, \bar{r}_1, \bar{r}_2)$  for the next index  $i \bmod 3 + 3$ , and output  $(3i + 0, c_0, c_1, r_1, r_2)$  for  $S_0$ ,  $(3i + 1, c_0, c_1, r_1, c_2, r_2)$  for  $S_1$ , and  $(3i + 2, c_2, r_2, \bar{c}_0, \bar{r}_1)$  for  $S_2$ , as described above (lines 30–32). The three arrays are then sorted (lines 33–35) and merged, whereby the comparison functions compares two representatives character-wise until a rank is found. The difference cover property guarantees that such a rank is found for every pair  $S_i, S_j$  during the Merge (lines 36–44).

The difference cover algorithm DC3 generalizes to DCX using a difference cover  $D$  for any ground set size  $X > 3$ . DCX constructs a recursive subproblem of size  $|D|/X$ , has at most  $\log_X |T|$  recursion levels and only one recursion branch. At every level of the recursion, only work with sorting complexity is needed, and a straightforward application of the Master theorem shows that the whole algorithms has the same complexity due to the small recursive subproblem. In the RAM model and with integer alphabets one can use radix sort in each level and thus DCX has linear running time. For our distributed model, DCX has the same complexity as sorting and merging.

Due to the subproblem size  $|D|/X$  is it best to use the largest  $X$  for a specific difference cover size. For  $|D| = 2$  this is  $X = 3$ , aka DC3. For difference covers of size three, the largest  $X = 7$  which yields DC7 with  $D_7 = \{0, 1, 3\}$ . And for difference covers of size four, the largest  $X = 13$ . Weese [23] showed that DC7 is optimal regarding the number of I/Os in an external memory model assuming index types are four times the byte size of characters. Due to these previous results we also implemented DC7 in Thrill.

Most of the previous discussion on DC3 can easily be extended to DC7: sort by seven characters instead of three, construct  $T_R = T_0 \oplus T_1 \oplus T_3$  in case not all character tuples are unique, and have Step 1 deliver  $R_0$ ,  $R_1$ , and  $R_3$  containing the ranks of all suffixes  $i \in D_7$ . We included the Thrill algorithm code for DC7 in Algorithms 11–13.

The key to implementing DC7 is in the construction of the tuple contents of the seven arrays  $S_0, \dots, S_6$  from  $R_0$ ,  $R_1$ ,  $R_3$ , and characters from  $T$ . Figure 3 shows a schematic to illustrate the underlying construction. For each index  $i$  there are three indexes  $(i + k_0 \bmod 7)$ ,  $(i + k_1 \bmod 7)$ , and  $(i + k_3 \bmod 7)$  in the difference cover  $D_7$ . The offsets depend on  $j = i \bmod 3$  for some index  $i$ , which classifies the suffix into  $S_j$ . The tuples in the arrays must contain all characters up to (but excluding) the



$S_0$	$c_0$	$c_1/r_1$	$r_2$
$S_1$	$c_1/r_1$	$r_2$	
$S_2$	$c_2/r_2$	$\bar{c}_0$	$\bar{r}_1$

for DC3 with  $D_3 = \{1, 2\}$

$S_0$	$c_0/r_0$	$c_1/r_1$	$c_2$	$r_3$			
$S_1$	$c_1/r_1$	$c_2$	$c_3/r_3$	$c_4$	$c_5$	$c_6$	$\bar{r}_0$
$S_2$	$c_2$	$c_3/r_3$	$c_4$	$c_5$	$c_6$	$\bar{c}_0/\bar{r}_0$	$\bar{r}_1$
$S_3$	$c_3/r_3$	$c_4$	$c_5$	$c_6$	$\bar{c}_0/\bar{r}_0$	$\bar{r}_1$	
$S_4$	$c_4$	$c_5$	$c_6$	$\bar{c}_0/\bar{r}_0$	$\bar{c}_1/\bar{r}_1$	$\bar{c}_2$	$\bar{r}_3$
$S_5$	$c_5$	$c_6$	$\bar{c}_0/\bar{r}_0$	$\bar{c}_1/\bar{r}_1$	$\bar{c}_2$	$\bar{r}_3$	
$S_6$	$c_6$	$\bar{c}_0/\bar{r}_0$	$\bar{c}_1/\bar{r}_1$	$\bar{c}_2$	$\bar{r}_3$		

for DC7 with  $D_7 = \{0, 1, 3\}$

**Fig. 3.** Construction of tuples in arrays  $S_i$  to represent suffixes in DC3 and DC7.

last known rank, since this is the information needed for the comparison function to perform character-wise comparisons up to the next known rank. The components of the tuples in  $S_0, \dots, S_6$  visualized in Figure 3 are selected in Algorithm 12 from  $Z$  via seven Map operations (lines 7–13). They are then sorted by characters up to the next known rank (lines 14–20) and then merged using CompareDC7 (Algorithm 13), which compares tuples character-wise up to the next known rank from all possible  $S_i/S_j$  pairs.

In our Thrill implementation, CompareDC7 is not rolled out as shown in the figure. Instead a lookup tables is used to determine how many characters and which of the included ranks need to be compared. Surprisingly, this more complex code was faster in our preliminary experiments, possibly due to the larger cost of decoding the instructions is the large unrolled comparison function.

### 3 Conclusion

We presented the implementation of five different suffix array construction algorithms in Thrill showing that the small set of algorithmic primitives provided by Thrill is sufficient to express the algorithms within the framework.

Our preliminary experimental results show that algorithms implemented in Thrill are competitive to hand-coded MPI implementations. By using the Thrill framework we gain additional benefits like future improvements of the algorithmic primitives in Thrill, and possibly even fault tolerance. Furthermore, Thrill already has automatic external memory support, hence our implementations are the first distributed external memory suffix array construction algorithms.

In a future version of this paper, we are going to add experimental results which detail the performance of our algorithms implemented in Thrill against their counterparts using MPI.

---

**Algorithm 11: DC7 Algorithm in Thrill (part one).**


---

```

1 function DC7PartOne( $T \in \text{DIA}(\Sigma)$ )
2    $T_7 := T.\text{FlatWindow}_7((i, [c_0, c_1, \dots, c_6]) \mapsto \text{MakeTuples}(i, c_0, c_1, \dots, c_6))$ 
3   with function MakeTuples ( $i \in \mathbb{N}_0, c_0, c_1, \dots, c_6 \in \Sigma$ )
4     if  $i \in D_7$  then emit  $(i, c_0, c_1, \dots, c_6)$  // Make tuples in difference cover.
5    $S := T_7.\text{Sort}((i, c_0, c_1, \dots, c_6)$  by  $(c_0, c_1, \dots, c_6)$ ) // Sort tuples lexicographically.
6    $I_S := S.\text{Map}((i, c_0, c_1, \dots, c_7) \mapsto i)$  // Extract sorted indices.
7    $N' := S.\text{FlatWindow}_2((i, [p_0, p_1]) \mapsto \text{CmpTuple}(i, p_0, p_1))$  // Compare tuples.
8   with function CmpTuple( $i \in \mathbb{N}_0, p_0 = (c_0, c_1, \dots, c_6), p_1 = (c'_0, c'_1, \dots, c'_6)$ )
9     if  $i = 0$  then emit 0 // Emit one sentinel for index 0.
10    if  $(c_0, c_1, \dots, c_6) = (c'_0, c'_1, \dots, c'_6)$  then emit 0 // Emit 0 or 1 depending on
11    else emit 1 // whether the previous tuple is equal.
12   $N := N'.\text{PrefixSum}()$  // Use prefix sum to calculate names.
13   $n_{\text{sub}} = \lceil 3|T|/7 \rceil, n_{\text{mod}0} = \lfloor |T|/7 \rfloor$  // Size of recursive problem and mod 0,
14   $n_{\text{mod}1} = \lfloor |T|/7 \rfloor, n_{\text{mod}01} = n_{\text{mod}0} + n_{\text{mod}1}$  // mod 1 and both parts of  $T_R$ .
15  if  $N.\text{Max}() + 1 = n_{\text{sub}}$  then // If duplicate names exist, sort names back to
16     $T'_R := \text{Zip}([I_S, N], (i, n) \mapsto (i, n)).\text{Sort}((i, n)$  by  $(i \bmod 7, i \text{ div } 7)$ ) // string order
17     $\text{SA}_R := \text{DC7}(T'_R.\text{Map}((i, n) \mapsto n))$  // as  $T_0 \oplus T_1 \oplus T_3$  and call suffix sorter.
18     $I'_R := \text{SA}_R.\text{ZipWithIndex}((i, r) \mapsto (i, r))$  // Invert resulting suffix array, but
19     $I_R := I'_R.\text{Sort}((i, r)$  by  $(\text{InterleavedRank}(i, i))$  // interleaved ISA for better locality
20    with function InterleavedRank( $i \in \mathbb{N}_0$ )
21      return (if  $i < n_{\text{mod}0}$  then  $i$  else if  $i < n_{\text{mod}01}$  then  $i - n_{\text{mod}0}$  else  $i - n_{\text{mod}01}$ )
22     $R_0 := I_R.\text{Filter}((i, r) \mapsto i < n_{\text{mod}0}).\text{Map}((i, r) \mapsto r + 1)$  // after separating
23     $R_1 := I_R.\text{Filter}((i, r) \mapsto i \geq n_{\text{mod}0}$  and  $i < n_{\text{mod}01}).\text{Map}((i, r) \mapsto r + 1)$  // ISA into
24     $R_3 := I_R.\text{Filter}((i, r) \mapsto i \geq n_{\text{mod}01}).\text{Map}((i, r) \mapsto r + 1)$  //  $R_0, R_1,$  and  $R_3$ .
25  else // Else, if all names/tuples are unique, then  $I_S$  is already the suffix array.
26     $R := I_S.\text{ZipWithIndex}((i, r) \mapsto (i, r))$  // Invert it to get ISA, but
27     $I_R := R.\text{Sort}((i, r)$  by  $(i \text{ div } 7, i)$ ) // interleave ISA for
28     $R_0 := I_R.\text{Filter}((i, r) \mapsto i = 0 \bmod 7).\text{Map}((i, r) \mapsto r + 1)$  // better locality
29     $R_1 := I_R.\text{Filter}((i, r) \mapsto i = 1 \bmod 7).\text{Map}((i, r) \mapsto r + 1)$  // after separating it
30     $R_3 := I_R.\text{Filter}((i, r) \mapsto i = 3 \bmod 7).\text{Map}((i, r) \mapsto r + 1)$  // into  $R_0, R_1,$  and  $R_3$ .
31  return DC7PartTwo( $T, R_0, R_1, R_3$ )

```

---

---

**Algorithm 12:** DC7 Algorithm in Thrill (part two).

---

```

1 function DC7PartTwo( $T \in \text{DIA}\langle \Sigma \rangle$ ,  $R_0, R_1, R_3 \in \text{DIA}\langle \mathbb{N}_0 \rangle$ )
2    $\hat{T}_7 := T.\text{FlatWindow}_7((i, [c_0, c_1, \dots, c_6]) \mapsto \text{MakeTuples}(i, c_0, c_1, \dots, c_6))$ 
3   with function MakeTuples ( $i \in \mathbb{N}_0$ ,  $c_0, c_1, \dots, c_6 \in \Sigma$ )           // Prepare Zip with all
4   | if not  $i \notin D_7$  then emit ( $i, c_0, c_1, \dots, c_6$ )           // triples  $i \notin D_7$ .
5    $Z' := \text{Zip}([\hat{T}_7, R_0, R_1, R_3], ((i, c_0, \dots, c_6), r_0, r_1, r_3) \mapsto (c_0, \dots, c_6, r_0, r_1, r_3))$  // Pull
6    $Z := Z'.\text{Window}_2((i, [(z_1, z_2)]) \mapsto (i, z_1, z_2))$            // chars and ranks using Zip from
7    $S'_0 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // four arrays
8   |  $\mapsto (7i + 0, c_0, r_0, c_1, r_1, c_2, r_2))$ 
9    $S'_1 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // to make
10  |  $\mapsto (7i + 1, c_1, r_1, c_2, r_2, c_3, r_3, c_4, r_4, c_5, r_5, c_6, r_6))$ 
11   $S'_2 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // arrays of
12  |  $\mapsto (7i + 2, c_2, r_2, c_3, r_3, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$ 
13   $S'_3 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // representatives
14  |  $\mapsto (7i + 3, c_3, r_3, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$ 
15   $S'_4 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // for each
16  |  $\mapsto (7i + 4, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$ 
17   $S'_5 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$  // suffix class.
18  |  $\mapsto (7i + 5, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$ 
19   $S'_6 := Z.\text{Map}((i, (c_0, \dots, c_6, r_0, r_1, r_3), (\bar{c}_0, \dots, \bar{c}_6, \bar{r}_0, \bar{r}_1, \bar{r}_3))$ 
20  |  $\mapsto (7i + 6, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$ 
21   $S_0 := S'_0.\text{Sort}((i, c_0, r_0, c_1, r_1, c_2, r_2)$  // Sort representatives
22  |  $\text{by } (r_0))$ 
23   $S_1 := S'_1.\text{Sort}((i, c_1, r_1, c_2, r_2, c_3, r_3, c_4, r_4, c_5, r_5, c_6, \bar{r}_0)$  // character-wise up to
24  |  $\text{by } (r_1))$ 
25   $S_2 := S'_2.\text{Sort}((i, c_2, r_2, c_3, r_3, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$  // next rank, and merge
26  |  $\text{by } (c_2, r_3))$ 
27   $S_3 := S'_3.\text{Sort}((i, c_3, r_3, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$  // sorted representatives
28  |  $\text{by } (r_3))$ 
29   $S_4 := S'_4.\text{Sort}((i, c_4, r_4, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$  // to deliver the
30  |  $\text{by } (c_4, c_5, c_6, \bar{r}_0))$ 
31   $S_5 := S'_5.\text{Sort}((i, c_5, r_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$  // final suffix array.
32  |  $\text{by } (c_5, c_6, \bar{r}_0))$ 
33   $S_6 := S'_6.\text{Sort}((i, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_2)$  // See Algorithm 13
34  |  $\text{by } (c_6, \bar{r}_0))$ 
35  return Merge( $[S_0, S_1, \dots, S_6]$ , CompareDC7).Map( $(i, \dots) \mapsto i$ ) // for CompareDC7.

```

---

---

**Algorithm 13: Full Comparison Function in DC7.**


---

```

1 function CompareDC7( $z_1, z_2$ )
2    $(r_0)$  <  $(r'_0)$  if  $z_1 \in S_0, z_2 \in S_0$ ,
3    $(r_0)$  <  $(r'_1)$  if  $z_1 \in S_0, z_2 \in S_1$ ,
4    $(c_0, r_1)$  <  $(c'_2, r'_3)$  if  $z_1 \in S_0, z_2 \in S_2$ ,
5    $(r_0)$  <  $(r'_3)$  if  $z_1 \in S_0, z_2 \in S_3$ ,
6    $(c_0, c_1, c_2, r_3)$  <  $(c'_4, c'_5, c'_6, \bar{r}'_0)$  if  $z_1 \in S_0, z_2 \in S_4$ ,
7    $(c_0, c_1, c_2, r_3)$  <  $(c'_5, c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_0, z_2 \in S_5$ ,
8    $(c_0, r_1)$  <  $(c'_6, \bar{r}'_0)$  if  $z_1 \in S_0, z_2 \in S_6$ ,
9    $(r_1)$  <  $(r'_1)$  if  $z_1 \in S_1, z_2 \in S_1$ ,
10   $(c_1, c_2, c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_2, c'_3, c'_4, c'_5, c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_1, z_2 \in S_2$ ,
11   $(r_1)$  <  $(r'_3)$  if  $z_1 \in S_1, z_2 \in S_3$ ,
12   $(c_1, c_2, c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_4, c'_5, c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_1, z_2 \in S_4$ ,
13   $(c_1, c_2, r_3)$  <  $(c'_5, c'_6, \bar{r}'_0)$  if  $z_1 \in S_1, z_2 \in S_5$ ,
14   $(c_1, c_2, r_3)$  <  $(c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_1, z_2 \in S_6$ ,
15   $(c_2, r_3)$  <  $(c'_2, r'_3)$  if  $z_1 \in S_2, z_2 \in S_2$ ,
16   $(c_2, c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c_3, c_4, c_5, c_6, \bar{c}_0, \bar{r}_1)$  if  $z_1 \in S_2, z_2 \in S_3$ ,
17   $(c_2, c_3, c_4, c_5, c_6, \bar{c}_0, \bar{r}_1)$  <  $(c'_4, c'_5, c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_2, z_2 \in S_4$ ,
18   $(c_2, c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_5, c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_2, z_2 \in S_5$ ,
19   $(c_2, r_3)$  <  $(c'_6, \bar{r}'_0)$  if  $z_1 \in S_2, z_2 \in S_6$ ,
20   $(r_3)$  <  $(r'_3)$  if  $z_1 \in S_3, z_2 \in S_3$ ,
21   $(c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_4, c'_5, c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_3, z_2 \in S_4$ ,
22   $(c_3, c_4, c_5, c_6, \bar{c}_0, \bar{r}_1)$  <  $(c'_5, c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_3, z_2 \in S_5$ ,
23   $(c_3, c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_3, z_2 \in S_6$ ,
24   $(c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_4, c'_5, c'_6, \bar{r}'_0)$  if  $z_1 \in S_4, z_2 \in S_4$ ,
25   $(c_4, c_5, c_6, \bar{r}_0)$  <  $(c'_5, c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_4, z_2 \in S_5$ ,
26   $(c_4, c_5, c_6, \bar{c}_0, \bar{r}_0)$  <  $(c'_6, \bar{c}'_0, \bar{c}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_1 \in S_4, z_2 \in S_6$ ,
27   $(c_5, c_6, \bar{r}_0)$  <  $(c'_5, c'_6, \bar{r}'_0)$  if  $z_1 \in S_5, z_2 \in S_5$ ,
28   $(c_6, \bar{c}_0, \bar{r}_1)$  <  $(c'_6, \bar{c}'_0, \bar{r}'_1)$  if  $z_1 \in S_5, z_2 \in S_6$ ,
29   $(c_6, \bar{r}_0)$  <  $(c'_6, \bar{r}'_0)$  if  $z_1 \in S_6, z_2 \in S_6$ ,
30  and symmetrically for  $z_1 \in S_i, z_2 \in S_j$  if  $i > j$ ,
31  with  $z_1 = (i, c_0, r_0, c_1, r_1, c_2, r_3)$  if  $z_1 \in S_0$ ,
32   $z_2 = (i', c'_0, r'_0, c'_1, r'_1, c'_2, r'_3)$  if  $z_2 \in S_0$ ,
33   $z_1 = (i, c_1, r_1, c_2, c_3, r_3, c_4, c_5, c_6, \bar{r}_0)$  if  $z_1 \in S_1$ ,
34   $z_2 = (i', c'_1, r'_1, c'_2, c'_3, r'_3, c'_4, c'_5, c'_6, \bar{r}'_0)$  if  $z_2 \in S_1$ ,
35   $z_1 = (i, c_2, c_3, r_3, c_4, c_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$  if  $z_1 \in S_2$ ,
36   $z_2 = (i', c'_2, c'_3, r'_3, c'_4, c'_5, c'_6, \bar{c}'_0, \bar{r}'_0, \bar{r}'_1)$  if  $z_2 \in S_2$ ,
37   $z_1 = (i, c_3, r_3, c_4, c_5, c_6, \bar{c}_0, \bar{r}_0, \bar{r}_1)$  if  $z_1 \in S_3$ ,
38   $z_2 = (i', c'_3, r'_3, c'_4, c'_5, c'_6, \bar{c}'_0, \bar{r}'_0, \bar{r}'_1)$  if  $z_2 \in S_3$ ,
39   $z_1 = (i, c_4, c_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_3)$  if  $z_1 \in S_4$ ,
40   $z_2 = (i, c'_4, c'_5, c'_6, \bar{c}'_0, \bar{r}'_0, \bar{c}'_1, \bar{r}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_2 \in S_4$ ,
41   $z_1 = (i, c_5, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_3)$  if  $z_1 \in S_5$ ,
42   $z_2 = (i', c'_5, c'_6, \bar{c}'_0, \bar{r}'_0, \bar{c}'_1, \bar{r}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_2 \in S_5$ ,
43   $z_1 = (i, c_6, \bar{c}_0, \bar{r}_0, \bar{c}_1, \bar{r}_1, \bar{c}_2, \bar{r}_3)$  if  $z_1 \in S_6$ ,
44   $z_2 = (i', c'_6, \bar{c}'_0, \bar{r}'_0, \bar{c}'_1, \bar{r}'_1, \bar{c}'_2, \bar{r}'_3)$  if  $z_2 \in S_6$ .

```

---




Fig. 4. DIA data-flow graph of DC7 with no recursion.

We will also show how compressed indexes like the FM-index can be efficiently constructed using the Thrill framework. Additionally, we want to extend the existing algorithms with longest common prefix (LCP) array construction and the DCX algorithms with discarding tuples [19] similar to the technique we applied to the prefix doubling algorithms.

## References

1. Uwe Baier. Linear-time suffix sorting - A new approach for suffix array construction. In Roberto Grossi and Moshe Lewenstein, editors, *Annual Symposium on Combinatorial Pattern Matching (CPM)*, volume 54 of *LIPICs*, pages 23:1–23:12. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.
2. Timo Bingmann, Michael Axtmann, Emanuel Jöbstl, Sebastian Lamm, Huyen Chau Nguyen, Alexander Noe, Sebastian Schlag, Matthias Stumpp, Tobias Sturm, and Peter Sanders. Thrill: High-performance algorithmic distributed batch data processing with C++. *arXiv preprint arXiv:1608.05634*, 2016.
3. Timo Bingmann, Johannes Fischer, and Vitaly Osipov. Inducing suffix and LCP arrays in external memory. In *Proceedings of the Meeting on Algorithm Engineering & Experiments (ALENEX)*, pages 88–102. SIAM, 2013.
4. Roman Dementiev, Juha Kärkkäinen, Jens Mehnert, and Peter Sanders. Better external memory suffix array construction. *ACM Journal of Experimental Algorithmics (JEA)*, 12:3.4:1–3.4:24, 2008.
5. Jasbir Dhaliwal, Simon J. Puglisi, and Andrew Turpin. Trends in suffix sorting: A survey of low memory algorithms. In Mark Reynolds and Bruce H. Thomas, editors, *Australasian Computer Science Conference (ACSC)*, volume 122 of *CRPIT*, pages 91–98. Australian Computer Society, 2012.
6. Patrick Flick and Srinivas Aluru. 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 (SC)*, page 16. ACM, 2015.
7. Gaston H. Gonnet, Ricardo A. Baeza-Yates, and Tim Snider. New indices for text: PAT trees and PAT arrays. In *Information Retrieval: Data Structures and Algorithms*, chapter 3, pages 66–82. Prentice-Hall, 1992.
8. Juha Kärkkäinen and Dominik Kempa. Engineering a lightweight external memory suffix array construction algorithm. In *International Conference on Algorithms for Big Data (ICABD)*, volume 1146 of *CEUR Workshop Proceedings*, pages 53–60. CEUR-WS.org, 2014.
9. Juha Kärkkäinen, Dominik Kempa, and Simon J Puglisi. Parallel external memory suffix sorting. In *Combinatorial Pattern Matching (CPM)*, pages 329–342. Springer, 2015.
10. Juha Kärkkäinen and Peter Sanders. Simple linear work suffix array construction. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2719 of *LNCS*, pages 943–955. Springer, 2003.
11. Juha Kärkkäinen, Peter Sanders, and Stefan Burkhardt. Linear work suffix array construction. *Journal of the ACM (JACM)*, 53(6):918–936, 2006.
12. Fabian Kulla and Peter Sanders. Scalable parallel suffix array construction. *Parallel Computing*, 33(9):605–612, 2007.

13. Julian Labeit, Julian Shun, and Guy E. Blelloch. *Parallel Lightweight Wavelet Tree, Suffix Array and FM-Index Construction*. PhD thesis, BS Thesis. Karlsruhe, Germany, 2015.
14. N. Jesper Larsson and Kunihiko Sadakane. Faster suffix sorting. Technical Report LUCS-TR:99-214, LUNDFD6/(NFCS-3140)/1–20/(1999), Department of Computer Science, Lund University, Lund, Sweden, 1999.
15. Udi Manber and Gene Myers. Suffix arrays: A new method for on-line string searches. *SIAM Journal on Computing*, 22(5):935–948, 1993.
16. Essam Mansour, Amin Allam, Spiros Skiadopoulos, and Panos Kalnis. ERA: efficient serial and parallel suffix tree construction for very long strings. *Proceedings of the VLDB Endowment*, 5(1):49–60, 2011.
17. Yuta Mori. DivSufSort, 2006. <https://github.com/y-256/libdivsufsort>.
18. Ge Nong, Sen Zhang, and Wai Hong Chan. Linear suffix array construction by almost pure induced-sorting. In *Data Compression Conference (DCC)*, pages 193–202. IEEE, 2009.
19. Simon J. Puglisi, William F. Smyth, and Andrew Turpin. The performance of linear time suffix sorting algorithms. In *Data Compression Conference (DCC)*, pages 358–367. IEEE, 2005.
20. Simon J. Puglisi, William F. Smyth, and Andrew Turpin. A taxonomy of suffix array construction algorithms. *ACM Comp. Surveys*, 39(2), 2007.
21. Jouni Sirén. Burrows-wheeler transform for terabases. In *Data Compression Conference (DCC)*, pages 211–220, 2016.
22. Heng Wang, Shaoliang Peng, Yutong Lu, Chengkun Wu, Jiajun Wen, Jie Liu, and Xiaoqian Zhu. BWTCP: A parallel method for constructing BWT in large collection of genomic reads. In *International Supercomputing Conference (ISC)*, pages 171–178. Springer, 2015.
23. David Weese. Entwurf und Implementierung eines generischen Substring-Index. Master’s thesis, Humboldt University Berlin, May 2006. <http://www.seqan.de/publications/weese06.pdf>.