International Journal of Foundations of Computer Science
© World Scientific Publishing Company

SORTING SUFFIXES OF TWO-PATTERN STRINGS.

FRANTISEK FRANEK and WILLIAM F. SMYTH*

Algorithms Research Group,
Dept. of Computing & Software,
McMaster University,
Hamilton, Ontario,
Canada L8S 4L7.

ABSTRACT

Recently, several authors presented linear recursive algorithms for sorting suffixes of a string. All these algorithms employ a similar three-step approach, based on an initial division of the suffixes of x into two sets: in step 1 sort the first set using recursive reduction of the problem, in step 2 determine the order of the suffixes in the second set based on the order of the suffixes in the first set, and in step 3 merge the two sets together. To optimize such an algorithm either for space or time, it may not be sufficient to optimize one of the three steps, since in doing so, one might increase the resources required for the others to an unacceptable extent.

Franek, Lu, and Smyth introduced two-pattern strings as a generalization of Sturmian strings. Like Sturmian strings, two-pattern strings are generated by iterated morphisms, but they exhibit a much richer structure.

In this paper we show that the suffixes of two-pattern strings can be sorted in linear time using a variant of the three step approach outlined above. It turns out that, given the order of the suffixes in a two-pattern string, one can almost directly list in linear time all the suffixes of its expansion under a two-pattern morphism.

Keywords: two-pattern string, suffix, suffix tree, suffix array, sorted suffixes

1. Introduction

Ever since Manber and Myers in [9] introduced suffix arrays as data structures comparable to suffix trees for most pattern matching tasks in strings, yet requiring significantly less memory, the search was on for a linear time algorithm for their construction. Such an algorithm for suffix tree construction had been known since 1997 [1]. In 2003 to our knowledge three different groups of researchers independently proposed linear recursive algorithms to sort string suffixes: [6, 7, 8]. Though different, all three algorithms employ three steps, based on a separation of the suffixes into two sets. In step 1 the first set is ordered using recursive reduction of the problem, in step 2 the suffixes of the second set are sorted based on the order of the suffixes in the first set, and in step 3 both ordered sets are merged together.

*also Department of Computing, Curtin University, Perth WA 6845, Australia.

The fact that all three algorithms follow this basic approach, yet use a completely different separation into sets, a different way of ordering the second set based on the first set, and a different merge technique, points to some common fundamental aspect of these algorithms. To optimize such an algorithm either for space or time, it may not be sufficient to optimize one of the three steps, since in doing so, one might increase the resources required for the others to an unacceptable extent.

Two-pattern strings were introduced in [2] as a generalization of Sturmian strings. Like Sturmian strings, two-pattern strings are generated by iterated morphisms, but they exhibit a much richer structure. It was shown in [3] that the iterated construction of these strings could be used to compute all the repetitions and near-repetitions in time linear in string length.

This paper was motivated by our investigation of the three different linear suffix sorting algorithms discussed above and our desire to fully understand the underlying phenomena. Thus, we investigated whether the recursive nature of two-pattern strings could be used in sorting of the suffixes in the approach of the three algorithms mentioned. As it turned out, the "natural" recursive reduction of two-pattern strings can be used for step 1, and then steps 2 and 3 can be simplified into a single step: from having the suffixes of the reduced string ordered, one can almost directly list the suffixes of the two-pattern string in the right order.

For the sake of completeness, let us recall the definition of a two-pattern string (see [2]), including all supporting definitions. Throughout this paper, a **binary string** means a string over the alphabet $\{a, b\}$.

Definition 1 *A binary string* $\boldsymbol{q}$ *is said to be* $\boldsymbol{p}$**-regular** *if and only if* $\boldsymbol{q} = \boldsymbol{upvu}$ *for some choice of (possibly empty) substrings* $\boldsymbol{u}$ *and* $\boldsymbol{v}$.

Definition 2 *An ordered pair* $(\boldsymbol{p}, \boldsymbol{q})$ *of nonempty binary strings is said to be* **suitable** *if and only if*

- $\boldsymbol{p}$ *is* **primitive** *(that is,* $\boldsymbol{p}$ *has no nonempty border);*
- $\boldsymbol{p}$ *is not a suffix of* $\boldsymbol{q}$*;*
- $\boldsymbol{q}$ *is neither a prefix nor a suffix of* $\boldsymbol{p}$*;*
- $\boldsymbol{q}$ *is not* $\boldsymbol{p}$*-regular.*

Note: *Since a two-pattern string is a concatenation of blocks* $\boldsymbol{p}^i\boldsymbol{q}$ *and* $\boldsymbol{p}^j\boldsymbol{q}$*, the above two definitions make sure that* $\boldsymbol{p}$ *and* $\boldsymbol{q}$ *are dissimilar enough to be recognized efficiently.*

Definition 3 $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ *is an* **expansion of scope** $\boldsymbol{\lambda}$*, if* $(\boldsymbol{p}, \boldsymbol{q})$ *is suitable,* $|\boldsymbol{p}| \leq \lambda$, $|\boldsymbol{q}| \leq \lambda$, $1 \leq i, j$, $i \neq j$ *are integers, and* λ *is an integer* ≥ 1.

Note: *An expansion* $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]$ *is applied to a binary string* $\boldsymbol{x}$ *in the following fashion: each occurrence of* a *in* $\boldsymbol{x}$ *is replaced by* $\boldsymbol{p}^i\boldsymbol{q}$ *and each occurrence of* b *by* $\boldsymbol{p}^j\boldsymbol{q}$*. The resulting string is denoted as* $\sigma(\boldsymbol{x})$*. We define* $\sigma(\varepsilon) = \varepsilon$*. The composition of two expansions* $\boldsymbol{\sigma}_1$ *and* $\boldsymbol{\sigma}_2$*,* $(\boldsymbol{\sigma}_1 \circ \boldsymbol{\sigma}_2)(x)$ *is defined by* $(\boldsymbol{\sigma}_1 \circ \boldsymbol{\sigma}_2)(x) = \boldsymbol{\sigma}_1(\boldsymbol{\sigma}_2(x))$*. The role of the scope* λ *is to limit the size of* $\boldsymbol{p}$ *and* $\boldsymbol{q}$ *that can be used in the following definition.*

Definition 4 *A binary string* $\boldsymbol{x}$ *is a* **two-pattern string of scope** $\boldsymbol{\lambda}$ *if there exists a sequence* $\{\sigma_1, \sigma_2, \ldots, \sigma_m\}$ *of expansions of scope* λ *so that* $\boldsymbol{x} = \sigma_1 \circ \cdots \circ \sigma_m(a)$.
It was mentioned at the end of [2] that if the definition of $\boldsymbol{p}$-regularity were made more restrictive, a larger class of complete two-pattern strings could be obtained. The more restrictive definition, sufficient to give two-pattern strings all their desired properties, contained a few typographical errors as it was given in [2], and so we provide a corrected definition here:
Definition 5 *A binary string* $\boldsymbol{q}$ *is said to be* $\boldsymbol{p}$**-regular** *(*$\boldsymbol{p}$ *a binary string) if and only if there exist (possibly empty) strings* $\boldsymbol{u}, \boldsymbol{v}$ *together with nonnegative integers* $n_1, n_2, \ldots, n_k$, $k \geq 1$, $r \geq 0$, *such that*

- *the integers* n_i *assume at most two distinct values — that is,*

$$\left|\{n_i : \ i \in 1..k\}\right| \leq 2;$$

- $\boldsymbol{q} = (\boldsymbol{u}\boldsymbol{p}^r\boldsymbol{v}\boldsymbol{p}^{n_1})(\boldsymbol{u}\boldsymbol{p}^r\boldsymbol{v}\boldsymbol{p}^{n_2})\cdots(\boldsymbol{u}\boldsymbol{p}^r\boldsymbol{v}\boldsymbol{p}^{n_k})\boldsymbol{u}$ *for some* $\boldsymbol{u}$, $\boldsymbol{v}$, $r \geq 0$, *where* $\boldsymbol{v} = \varepsilon$ *if* $r = 0$.

Note: *the definition 5 can be used to replace the definition 1. In fact, all the proofs accompanying this paper are compliant with the more restrictive definition 5.*

Certain finite fragments of the well-known Fibonacci string and the equally well-known Sturmian strings are in fact two-pattern strings of scope $\lambda = 1$ (see [2]).

Here are a few simple examples of two-pattern strings:

1. a, now apply $\boldsymbol{\sigma_2} = [ab, ba, 2, 3]$ to it, we get
2. $\boldsymbol{\sigma_2}(a) = ababba$, now apply $\boldsymbol{\sigma_1} = [abb, aa, 1, 4]$ to it, we get
3. $\boldsymbol{\sigma_1}(\boldsymbol{\sigma_2}(ababba)) = abbaa(abb)^4aaabbaa(abb)^4aa(abb)^4aaabbaa$.

Strings 1, 2, and 3 are all two-pattern strings of scope 3 (string 2 is in fact of scope 2, and string 1 is in fact of scope 1).

It was shown in [2] that complete two-pattern strings can be recognized in linear time: the recognition algorithm outputs an essentially unique sequence of expansions to construct the string from a. So in the following we can assume that not only do we have a complete two-pattern string, but also the sequence of expansions that iteratively generates the string.

In the next section we describe the principles underlying the algorithm for sorting suffixes of a two-pattern string. In Section 3 we provide an overview of the algorithm itself, while Section 5 we list some of the main lemmas on which the algorithm is based. We conclude with Section 6.

2. The Principles Underlying the Algorithm

For the sake of clarity and brevity, we introduce several symbols: we use the symbol $\boldsymbol{u} < \boldsymbol{v}$ for strings $\boldsymbol{u}$, $\boldsymbol{v}$ to express that $\boldsymbol{u}$ is lexicographically smaller than $\boldsymbol{v}$. We use the symbol $\prec$ in $\boldsymbol{u} \prec \boldsymbol{v}$ (or $\succ$ in $\boldsymbol{u} \succ \boldsymbol{v}$) to express the fact that $\boldsymbol{u} < \boldsymbol{v}$ yet $\boldsymbol{u}$ is not a prefix of $\boldsymbol{v}$ (or $\boldsymbol{v} < \boldsymbol{u}$ yet $\boldsymbol{v}$ is not a prefix of $\boldsymbol{u}$). Note that $\boldsymbol{u} < \boldsymbol{v}$

iff ($\boldsymbol{u} \prec \boldsymbol{v}$ or $\boldsymbol{u}$ is a prefix of $\boldsymbol{v}$). We use the symbol $\boldsymbol{u} \asymp \boldsymbol{v}$ to indicate that either $\boldsymbol{u} \prec \boldsymbol{v}$ or $\boldsymbol{u} \succ \boldsymbol{v}$.

For a binary string $\boldsymbol{u}$, we will use $\overline{\boldsymbol{u}}$ to denote its ones-complement; that is, the string formed by interchanging a's and b's in $\boldsymbol{u}$.

In accordance with [2], if $\boldsymbol{x}$, $\boldsymbol{y}$ are complete two-pattern strings, σ an expansion, and $\boldsymbol{y} = \sigma(\boldsymbol{x})$, then the occurrences of copies of $\boldsymbol{p}$ and copies of $\boldsymbol{q}$ in the concatenation of blocks $\boldsymbol{p}^i\boldsymbol{q}$ and $\boldsymbol{p}^j\boldsymbol{q}$ as defined by $\sigma(\boldsymbol{x})$ are called **restrained** copies. Any other occurrence of $\boldsymbol{p}$ or $\boldsymbol{q}$ is referred to as **free**. A consecutive sequence of restrained copies of $\boldsymbol{p}$'s and/or $\boldsymbol{q}$'s will also be referred to as a **restrained configuration** or a **restrained substring** of $\boldsymbol{y}$.

Throughout the following discussion we assume that the scope λ is fixed and that $\boldsymbol{y} = \sigma(\boldsymbol{x})$, where $\boldsymbol{x}$ is a complete two-pattern string of scope λ and $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ an expansion of scope λ. Moreover we assume that all suffixes of $\boldsymbol{x}$ are lexicographically sorted: $\boldsymbol{\rho}_1 < \cdots < \boldsymbol{\rho}_{|\boldsymbol{x}|}$. We then describe how to order the suffixes of $\boldsymbol{y}$. We may assume further that $\boldsymbol{q} < \boldsymbol{p}$. If it were not the case, according to Lemma 2 (see section 5 below), $\overline{\boldsymbol{q}} < \overline{\boldsymbol{p}}$, we sort all of the suffixes of $\overline{\boldsymbol{y}} = \overline{\sigma}(\boldsymbol{x})$, where $\overline{\sigma} = [\overline{\boldsymbol{p}}, \overline{\boldsymbol{q}}, i, j]_\lambda$, and then we can list all suffixes of $\boldsymbol{y}$ in proper order efficiently in linear time.

Since we are assuming $\boldsymbol{q} < \boldsymbol{p}$, according to Lemma 1 (see section 5 below), for any suffixes $\boldsymbol{\rho}_1$, $\boldsymbol{\rho}_2$ of $\boldsymbol{x}$, if $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$, then $\sigma(\boldsymbol{\rho}_1) < \sigma(\boldsymbol{\rho}_2)$. In simple terms, the assumption $\boldsymbol{q} < \boldsymbol{p}$ makes all expansions to preserve the order of suffixes.

We put all the suffixes of $\boldsymbol{y}$ into disjoint buckets of five types $\boldsymbol{A}$–$\boldsymbol{E}$. Their definitions follow (*note that the expansion $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ is fixed*):

- For every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$ and for every integer k, $0 < k < i$, $\boldsymbol{A}_{\boldsymbol{\delta},k} = \{\boldsymbol{\delta}\boldsymbol{p}^k\boldsymbol{q}\sigma(\boldsymbol{\rho}) :\}\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$ or $\boldsymbol{\rho} = \varepsilon\}$;
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$ that is also a suffix of $\boldsymbol{q}$, $\boldsymbol{A}_{\boldsymbol{\delta},i} = \{\boldsymbol{\delta}\boldsymbol{p}^i\boldsymbol{q}\sigma(\boldsymbol{\rho}) :$ $\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$ or $\boldsymbol{\rho} = \varepsilon\}$;
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$ that is not a suffix of $\boldsymbol{q}$, $\boldsymbol{A}_{\boldsymbol{\delta},i} = \{\boldsymbol{\delta}\boldsymbol{p}^i\boldsymbol{q}\sigma(\boldsymbol{\rho}) :$ $b\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$, $\boldsymbol{\rho}$ can be empty$\}$;
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$ and for every integer k, $i < k < j$, $\boldsymbol{A}_{\boldsymbol{\delta},k} = \{\boldsymbol{\delta}\boldsymbol{p}^k\boldsymbol{q}\sigma(\boldsymbol{\rho}) :$ $b\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$, $\boldsymbol{\rho}$ can be empty$\}$.
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$, $\boldsymbol{B}_{\boldsymbol{\delta}} = \{\boldsymbol{\delta}\boldsymbol{q}\sigma(\boldsymbol{\rho}) : \boldsymbol{\rho}$ is a proper nontrivial suffix of $\boldsymbol{x}\}$;
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{q}$ that is not a suffix of $\boldsymbol{p}$, $\boldsymbol{C}_{\boldsymbol{\delta}} = \{\boldsymbol{\delta}\boldsymbol{p}^i\boldsymbol{q}\sigma(\boldsymbol{\rho}) :$ $a\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$, $\boldsymbol{\rho}$ can be empty$\}$;
- for every nontrivial suffix $\boldsymbol{\delta}$ of $\boldsymbol{q}$, $\boldsymbol{D}_{\boldsymbol{\delta}} = \{\boldsymbol{\delta}\boldsymbol{p}^j\boldsymbol{q}\sigma(\boldsymbol{\rho}) :$ $b\boldsymbol{\rho}$ is a proper suffix of $\boldsymbol{x}$, $\boldsymbol{\rho}$ can be empty$\}$;
- $\boldsymbol{E} = \{\boldsymbol{\delta}\boldsymbol{q} :$ $\boldsymbol{\delta}$ is a nontrivial suffix of $\boldsymbol{p}\} \cup \{\boldsymbol{\delta} :$ $\boldsymbol{\delta}$ is a nontrivial suffix of $\boldsymbol{q}\}$.

(where the term *proper suffix* refers to a suffix that is not equal to the whole string and the term *trivial suffix* refers to the empty suffix).
It is straightforward to check that any suffix of $\boldsymbol{y}$ belongs to one of the buckets $\boldsymbol{A}$–$\boldsymbol{E}$ (for proof see either [5] or [4]). We are going to order the suffixes in buckets $\boldsymbol{A}$–$\boldsymbol{D}$ based on the ordering of the suffixes for $\boldsymbol{x}$ (Step 1), then merge in the suffixes from $\boldsymbol{E}$ (Steps 2 & 3); since $|\boldsymbol{E}| \leq 2\lambda$, this will not destroy the linearity of the algorithm. Note that the order within each bucket is determined by the order of suffixes of $\boldsymbol{x}$:

in the bucket $\boldsymbol{A}_{\boldsymbol{\delta},k}$: $\boldsymbol{\delta p}^k\boldsymbol{q}\sigma(\boldsymbol{\rho}_1) < \boldsymbol{\delta p}^k\boldsymbol{q}\sigma(\boldsymbol{\rho}_2)$ if $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$;
in the bucket $\boldsymbol{B}_{\boldsymbol{\delta}}$: $\boldsymbol{\delta q}\sigma(\boldsymbol{\rho}_1) < \boldsymbol{\delta q}\sigma(\boldsymbol{\rho}_2)$ if $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$;
in the bucket $\boldsymbol{C}_{\boldsymbol{\delta}}$: $\boldsymbol{\delta p}^i\boldsymbol{q}\sigma(\boldsymbol{\rho}_1) < \boldsymbol{\delta p}^i\boldsymbol{q}\sigma(\boldsymbol{\rho}_2)$ if $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$;
and in the bucket $\boldsymbol{D}_{\boldsymbol{\delta}}$: $\boldsymbol{\delta p}^j\boldsymbol{q}\sigma(\boldsymbol{\rho}_1) < \boldsymbol{\delta p}^j\boldsymbol{q}\sigma(\boldsymbol{\rho}_2)$ if $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$.

Thus, it is straightforward to list the suffixes in each bucket in the correct order, given the order of the suffixes of $\boldsymbol{x}$.

We make use of the following notation: if X, Y are sets of suffixes of $\boldsymbol{y}$, we write $X \ll Y$ iff $(\forall x \in X)(\forall y \in Y)(x < y)$. The major observation our algorithm is based on is that the buckets are linearly ordered by $\ll$; that is, pairwise orderings can be made between bucket pairs of types

$$\boldsymbol{AA}, \boldsymbol{AB}, \boldsymbol{AC}, \boldsymbol{AD}, \boldsymbol{BB}, \boldsymbol{BC}, \boldsymbol{BD}, \boldsymbol{CC}, \boldsymbol{CD}, \boldsymbol{DD}, \tag{1}$$

based on five mutually exclusive (and exhaustive) conditions on any pair $\boldsymbol{\delta}_1$, $\boldsymbol{\delta}_2$ of suffixes of $\boldsymbol{p}$ and/or $\boldsymbol{q}$:

(C1) $\boldsymbol{\delta}_1 \prec \boldsymbol{\delta}_2$;

(C2) $\boldsymbol{\delta}_1 \succ \boldsymbol{\delta}_2$;

(C3) $\boldsymbol{\delta}_1$ is a proper prefix of $\boldsymbol{\delta}_2$;

(C4) $\boldsymbol{\delta}_2$ is a proper prefix of $\boldsymbol{\delta}_1$;

(C5) $\boldsymbol{\delta}_1 = \boldsymbol{\delta}_2 = \boldsymbol{\delta}$.

Observe that, given $\boldsymbol{\delta}_1$ and $\boldsymbol{\delta}_2$, to determine which of these conditions holds requires at most λ letter comparisons (since $|\boldsymbol{\delta}_1| \leq \lambda$, $|\boldsymbol{\delta}_2| \leq \lambda$).

Thus, for example, two $\boldsymbol{A}$ buckets can be compared as follows:

(C1) $\boldsymbol{A}_{\boldsymbol{\delta}_1,k_1} \ll \boldsymbol{A}_{\boldsymbol{\delta}_2,k_2}$.

(C2) $\boldsymbol{A}_{\boldsymbol{\delta}_2,k_2} \ll \boldsymbol{A}_{\boldsymbol{\delta}_1,k_1}$.

(C3) Let $\boldsymbol{\delta}_2 = \boldsymbol{\delta}_1\boldsymbol{\delta}_1'$ for some nonempty $\boldsymbol{\delta}_1'$:

(a) if $\boldsymbol{\delta}_1' \prec \boldsymbol{p}$, then $\boldsymbol{A}_{\boldsymbol{\delta}_2,k_2} \ll \boldsymbol{A}_{\boldsymbol{\delta}_1,k_1}$;
(b) otherwise, $\boldsymbol{A}_{\boldsymbol{\delta}_1,k_1} \ll \boldsymbol{A}_{\boldsymbol{\delta}_2,k_2}$.

(C4) Let $\boldsymbol{\delta}_1 = \boldsymbol{\delta}_2\boldsymbol{\delta}_2'$ for some nonempty $\boldsymbol{\delta}_2'$:

(a) If $\boldsymbol{\delta}_2' \prec \boldsymbol{p}$, then $\boldsymbol{A}_{\boldsymbol{\delta}_1,k_1} \ll \boldsymbol{A}_{\boldsymbol{\delta}_2,k_2}$;

(b) otherwise, $\boldsymbol{A}_{\boldsymbol{\delta}_2,k_2} \ll \boldsymbol{A}_{\boldsymbol{\delta}_1,k_1}$.

(C5) (a) If $k_1 < k_2$, then $\boldsymbol{A}_{\boldsymbol{\delta},k_1} \ll \boldsymbol{A}_{\boldsymbol{\delta},k_2}$;
(b) if $k_1 = k_2$, then $\boldsymbol{A}_{\boldsymbol{\delta},k_1} = \boldsymbol{A}_{\boldsymbol{\delta},k_2}$;
(c) if $k_1 > k_2$, then $\boldsymbol{A}_{\boldsymbol{\delta},k_2} \ll \boldsymbol{A}_{\boldsymbol{\delta},k_1}$.

It is not very hard to prove that this ordering is correct. The demonstration for cases (C1), (C2) and (C5) is straightforward. For (C3), observe that we are comparing $\boldsymbol{\delta}_1\boldsymbol{p}^{k_1}\boldsymbol{q}\cdots$ with $\boldsymbol{\delta}_2\boldsymbol{p}^{k_2}\boldsymbol{q}\cdots$, hence $\boldsymbol{p}^{k_1}\boldsymbol{q}\cdots$ with $\boldsymbol{\delta}'_1\boldsymbol{p}^{k_2}\boldsymbol{q}\cdots$. Since $\boldsymbol{\delta}'_1$ is a suffix of $\boldsymbol{\delta}_2$, it is also a suffix of $\boldsymbol{p}$ and so cannot be a prefix of $\boldsymbol{p}$. It follows that either $\boldsymbol{\delta}'_1 \prec \boldsymbol{p}$ or $\boldsymbol{\delta}'_1 \succ \boldsymbol{p}$, and the result follows. The proof for (C4) is exactly analogous.

Furthermore the $\boldsymbol{AA}$ ordering is efficient, since the cases (a) and (b) in (C3) and (C4) can be processed in at most λ constant-time steps in addition to the λ steps that may be required to identify which condition holds: thus a total of at most 2λ steps altogether.

The results for the other pairs listed in (1) are similar: the details vary slightly from one case to another. The main result is that any of the pairs can be processed in at most 3λ steps, a constant. To avoid distracting the reader with unnecessary and uninteresting detail, we do not include the other cases here. For those details, please see either [4] or [5].

3. The High-Level Logic of the Algorithm

We describe only the recursive step (Step 1) that takes us from $\boldsymbol{x}$ and its sorted suffixes to the corresponding sorted suffixes of $\boldsymbol{y} = \sigma(\boldsymbol{x})$, where $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$. Recall that we assume $\boldsymbol{q} < \boldsymbol{p}$.

1. Create names $(A, \boldsymbol{\delta})$ for every suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$. (*This requires at most λ steps. Each name will be eventually replaced by a sequence of buckets, see below.*)

2. Sort the names according to the order described in the previous section for mutual comparison of the four $\boldsymbol{A}$ buckets (of course, according to (C1)-(C4) only). (*This requires at most $2\lambda^3$ steps as we are sorting λ names and each comparison requires $\leq 2\lambda$ steps.*)

3. Replace every name $(A, \boldsymbol{\delta})$ by a sequence of names $(A, \boldsymbol{\delta}, k)$, $1 \leq k < j$. Let us call the resulting sequence BUCKETS. (*Now we have the names of $\boldsymbol{A}$ buckets in the proper order. This requires at most $|\boldsymbol{y}|$ steps as the size of* BUCKETS *is $\leq |\boldsymbol{y}|$. Each name $(A, \boldsymbol{\delta}, k)$ will eventually be replaced by a corresponding bucket $\boldsymbol{A}_{\delta,k}$, see below.*)

4. Create names $(B, \boldsymbol{\delta})$ for every suffix $\boldsymbol{\delta}$ of $\boldsymbol{p}$. (*This requires at most λ steps. Each name $(B, \boldsymbol{\delta})$ will eventually be replaced by a corresponding bucket $\boldsymbol{B}_{\boldsymbol{\delta}}$, see below.*)

5. Merge into BUCKETS all names $(B, \boldsymbol{\delta})$ according to comparisons as described in comparing $\boldsymbol{A}$ buckets to $\boldsymbol{B}$ buckets. (*This requires at most* $|\text{BUCKETS}|3\lambda^2$

steps, as we are merging in λ names and each comparison requires $\leq 3\lambda$ steps, hence at most $|\boldsymbol{y}|3\lambda^2$ steps.)

6. Create names $(C, \boldsymbol{\delta})$ for every suffix $\boldsymbol{\delta}$ of $\boldsymbol{q}$ that is not a suffix of $\boldsymbol{p}$. (*This requires at most λ^2 steps. Each name $(C, \boldsymbol{\delta})$ will eventually be replaced by a bucket $\boldsymbol{C_\delta}$, see below.*)

7. Merge into BUCKETS all names $(C, \boldsymbol{\delta})$ according to comparisons as described in comparing $\boldsymbol{A}$ buckets to $\boldsymbol{C}$ buckets and $\boldsymbol{B}$ buckets to $\boldsymbol{C}$ buckets. (*This requires at most* |BUCKETS|$3\lambda^2$ *steps, hence at most $|\boldsymbol{y}|3\lambda^2$ steps.*)

8. Create names $(D, \boldsymbol{\delta})$ for every suffix $\boldsymbol{\delta}$ of $\boldsymbol{q}$. (*This requires at most λ steps. Each name $(D, \boldsymbol{\delta})$ will eventually be replaced by a bucket $\boldsymbol{D_\delta}$, see below.*)

9. Merge into BUCKETS all names $(D, \boldsymbol{\delta})$ according to comparisons as described in comparing $\boldsymbol{A}$ buckets to $\boldsymbol{D}$ buckets, $\boldsymbol{B}$ buckets to $\boldsymbol{D}$ buckets, $\boldsymbol{C}$ buckets to $\boldsymbol{D}$ buckets. (*Now we have all required bucket names, except $\boldsymbol{E}$, in proper order. This requires at most* |BUCKETS|$3\lambda^2$ *steps, hence at most $|\boldsymbol{y}|3\lambda^2$ steps.*)

10. Traverse BUCKETS and replace each name by a sequence of suffixes according to the sequence of suffixes of $\boldsymbol{x}$. Let us call this sequence SUFFIXES. (*We turned the names into proper buckets and merged them all together in a single list. Now we have all suffixes from buckets $\boldsymbol{A}$–$\boldsymbol{D}$ in proper order. This requires at most $|\boldsymbol{y}|$ steps as the size of* SUFFIXES *is $\leq |\boldsymbol{y}|$.*)

11. Merge into SUFFIXES the suffixes from the bucket $\boldsymbol{E}$. (*This requires at most* |SUFFIXES|$4\lambda^2$ *steps, as we are merging in 2λ suffixes, each of length $\leq 2\lambda$, hence at most $|\boldsymbol{y}|4\lambda^2$ steps.*)

SUFFIXES is now a sorted list of all suffixes of $\boldsymbol{y}$ and it took less than $\alpha|\boldsymbol{y}|$ steps, where we set $\alpha = 2\lambda^3 + 14\lambda^2 + 3\lambda + 2$. Since every reduction of a complete two-pattern string at least halves its length, altogether the algorithm with all iterative steps included took less than $\boldsymbol{\alpha}n + \boldsymbol{\alpha}\frac{n}{2} + \boldsymbol{\alpha}\frac{n}{4} + \cdots < 2\boldsymbol{\alpha}n$ steps, where n is the size of the input string.

4. An example

Let $\boldsymbol{x} = aab\$$, and let $\sigma = [ba, ab, 1, 2]$. (Thus $\boldsymbol{q} = ab < \boldsymbol{p} = ba$.) Hence $\boldsymbol{y} = \sigma(\boldsymbol{x}) = baabbaabbabaab\$$.
All nontrivial suffixes of $\boldsymbol{x}$ are (listed in the lexicographic order) a, aa, and aab. All nontrivial suffixes of $\boldsymbol{p}$ are ba and a, and all nontrivial suffixes of $\boldsymbol{q}$ are ab and b. Let us list the buckets:
$\boldsymbol{A}_{ba,1} = \{babaab\sigma(\rho) : b\rho$ is a proper suffix of $\boldsymbol{x}$, ρ can be $\varepsilon\} = \{babaab\} = \{\boldsymbol{y}[9..14]\}$.
$\boldsymbol{A}_{a,1} = \{abaab\sigma(\rho) : b\rho$ is a proper suffix of $\boldsymbol{x}$, ρ can be $\varepsilon\} = \{abaab\} = \{\boldsymbol{y}[10..14]\}$.
$\boldsymbol{B}_{ba} = \{baab\sigma(\rho) : \rho$ is a proper suffix of $\boldsymbol{x}\} = \{baab\sigma(ab), baab\sigma(b)\} =$
$\{baabbaabbabaab, baabbabaab\} = \{\boldsymbol{y}[1..14], \boldsymbol{y}[5..14]\}$.
$\boldsymbol{B}_{a} = \{aab\sigma(\rho) : \rho$ is a proper suffix of $\boldsymbol{x}\} = \{aab\sigma(ab), aab\sigma(b)\} =$

$\{aabbaabbabaab, aabbabaab\} = \{\boldsymbol{y}[2..14], \boldsymbol{y}[6..14]\}$.
$\boldsymbol{C}_{ab} = \{abbaab\sigma(\rho) : a\rho$ is a proper suffix of $\boldsymbol{x}\} = \{abbaab\sigma(b)\} = \{abbaabbabaab\} = \{\boldsymbol{y}[3..14]\}$.
$\boldsymbol{C}_b = \{bbaab\sigma(\rho) : b\rho$ is a proper suffix of $\boldsymbol{x}\} = \{bbaab\sigma(b)\} = \{bbaabbabaab\} = \{\boldsymbol{y}[4..14]\}$.
$\boldsymbol{D}_{ab} = \{abbabaab\sigma(\rho) : b\rho$ is a proper suffix of $\boldsymbol{x}$, ρ can be $\varepsilon\} = \{abbabaab\} = \{\boldsymbol{y}[7..14]\}$.
$\boldsymbol{D}_b = \{bbabaab\sigma(\rho) : b\rho$ is a proper suffix of sx, ρ can be $\varepsilon\} = \{bbabaab\} = \{\boldsymbol{y}[8..14]\}$.
$\boldsymbol{E} = \{baab, aab, ab, b\} = \{\boldsymbol{y}[11..14], \boldsymbol{y}[12..14], \boldsymbol{y}[13..14], \boldsymbol{y}[14..14]\}$.

First note that we really listed all nontrivial suffixes of $\boldsymbol{y}$: $\boldsymbol{y}[1..14]$, $\boldsymbol{y}[2..14]$, ..., $\boldsymbol{y}[14..14]$. Also note that the suffixes in the buckets are listed in lexicographic order. Let us list the pairwise relationships of all buckets: $\boldsymbol{A}_{ba,1} \gg \boldsymbol{A}_{a,1}$ (by C2), $\boldsymbol{A}_{ba,1} \gg \boldsymbol{B}_{ba}$ (by C5), $\boldsymbol{A}_{ba,1} \gg \boldsymbol{B}_a$ (by C2), $\boldsymbol{A}_{ba,1} \gg \boldsymbol{C}_{ab}$ (by C2), $\boldsymbol{A}_{ba,1} \ll \boldsymbol{C}_b$ (by C4a), $\boldsymbol{B}_{ba} \gg \boldsymbol{B}_a$ (by C2), $\boldsymbol{B}_{ba} \gg \boldsymbol{C}_{ab}$ (by C2), $\boldsymbol{B}_{ba} \ll \boldsymbol{C}_b$ (by C4a), $\boldsymbol{B}_{ba} \gg \boldsymbol{D}_{ab}$ (by C2), $\boldsymbol{B}_{ba} \ll \boldsymbol{D}_b$ (by C4a), $\boldsymbol{B}_a \ll \boldsymbol{C}_{ab}$ (by C3b), $\boldsymbol{B}_a \ll \boldsymbol{C}_b$ (by C1), $\boldsymbol{B}_a \ll \boldsymbol{D}_{ab}$ (by C3b), $\boldsymbol{B}_a \ll \boldsymbol{D}_b$ (by C1), $\boldsymbol{A}_{ba,1} \gg \boldsymbol{D}_{ab}$ (by C2), $\boldsymbol{A}_{ba,1} \ll \boldsymbol{D}_b$ (by C4a), $\boldsymbol{A}_{a,1} \ll \boldsymbol{B}_{ba}$ (by C1), $\boldsymbol{A}_{a,1} \gg \boldsymbol{B}_a$ (by C5), $\boldsymbol{A}_{a,1} \ll \boldsymbol{C}_{ab}$ (by C3b), $\boldsymbol{A}_{a,1} \ll \boldsymbol{C}_b$ (by C1), $\boldsymbol{A}_{a,1} \ll \boldsymbol{D}_{ab}$ (by C3b), $\boldsymbol{A}_{a,1} \ll \boldsymbol{D}_b$ (by C1), $\boldsymbol{C}_{ab} \ll \boldsymbol{C}_b$ (by C1), $\boldsymbol{C}_{ab} \ll \boldsymbol{D}_{ab}$ (by C5), $\boldsymbol{C}_{ab} \ll \boldsymbol{D}_b$ (by C1), $\boldsymbol{C}_b \gg \boldsymbol{D}_{ab}$ (by C1), $\boldsymbol{C}_b \ll \boldsymbol{D}_b$ (by C5), $\boldsymbol{D}_{ab} \ll \boldsymbol{D}_b$ (by C1).

Now follow the 11 steps.

1. create names $(A, ba), (A, a)$
2. sort them: $(A, a), (A, ba)$ (according to (C2))
3. "refine" the names to BUCKETS=$(A, a, 1), (A, ba, 1)$
4. create names to $(B, ba), (B, a)$
5. merge them into BUCKETS=$(B, a), (A, a, 1), (B, ba), (A, ba, 1)$
6. create names to $(C, ab), (C, b)$
7. merge them into BUCKETS=$(B, a), (A, a, 1), (C, ab), (B, ba), (A, ba, 1), (C, b)$
8. create names to $(D, ba), (D, a)$
9. merge them into BUCKETS=$(B, a), (A, a, 1), (C, ab), (D, ab), (B, ba), (A, ba, 1), (C, b), (D, b)$
10. replace the names by buckets: SUFFIXES= $\boldsymbol{y}[2..14], \boldsymbol{y}[6..14]$, $\boldsymbol{y}[10..14]$, $\boldsymbol{y}[3..14]$, $\boldsymbol{y}[7..14]$, $\boldsymbol{y}[1..14], \boldsymbol{y}[5..14]$, $\boldsymbol{y}[9..14]$, $\boldsymbol{y}[4..14]$, $\boldsymbol{y}[8..14]$.
11. merge in $\boldsymbol{E}$ bucket: SUFFIXES= $\boldsymbol{y}[12..14]$, $\boldsymbol{y}[2..14], \boldsymbol{y}[6..14]$, $\boldsymbol{y}[13..14]$, $\boldsymbol{y}[10..14]$, $\boldsymbol{y}[3..14]$, $\boldsymbol{y}[7..14]$, $\boldsymbol{y}[14..14], \boldsymbol{y}[11..14]$, $\boldsymbol{y}[1..14], \boldsymbol{y}[5..14]$, $\boldsymbol{y}[9..14]$, $\boldsymbol{y}[4..14]$, $\boldsymbol{y}[8..14]$.

5. The Supporting Lemmas

For the proofs, as mentioned above, see either [4] or [5].

The first lemma establishes that the ordering of suffixes is invariant under an expansion with $\boldsymbol{q} < \boldsymbol{p}$.

Lemma 1 *Let $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ be an expansion and $\boldsymbol{q} < \boldsymbol{p}$. Let $\boldsymbol{x}$ and $\boldsymbol{y}$ be two-pattern strings of scope λ and let $\boldsymbol{y} = \sigma(\boldsymbol{x})$. Let $\boldsymbol{\rho}_1$, $\boldsymbol{\rho}_2$ be suffixes of $\boldsymbol{x}$ so that $\boldsymbol{\rho}_1 < \boldsymbol{\rho}_2$.*
Then $\sigma(\boldsymbol{\rho}_1) < \sigma(\boldsymbol{\rho}_2)$.

The next lemma tells us that after interchanging a and b in a binary string, we can efficiently list the suffixes of the complement in lexicographic order knowing the order of suffixes in the original string.

Lemma 2 *Let $\boldsymbol{\rho}_1 < \cdots < \boldsymbol{\rho}_n$ be the sequence of all suffixes of a binary string $\boldsymbol{u}$ in ascending lexicographic order. Then there is an efficient linear-time procedure to list all suffixes of $\overline{\boldsymbol{u}}$ in ascending lexicographic order.*

The next three lemmas are technical lemmas required for some of the proofs (see website referenced above) that the pairs (1) can be processed correctly in $O(3\lambda)$ time. Essentially these lemmas tell us that the ordering of restrained suffixes of $\boldsymbol{y}$ can be accomplished in at most 2λ constant-time algorithmic steps.

Lemma 3 *Let $\boldsymbol{x}$, $\boldsymbol{y}$ be two-pattern strings of scope λ, $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ an expansion, and $\boldsymbol{y} = \sigma(\boldsymbol{x})$. Let $\boldsymbol{u}$ be a non-empty binary string and let $\boldsymbol{uqp}$ be a suffix of a restrained configuration $\boldsymbol{pqp}$ of $\boldsymbol{y}$ and let $\boldsymbol{qp}$ be a restrained configuration of $\boldsymbol{y}$. Then $\boldsymbol{uqp} \asymp \boldsymbol{qp}$ and whether $\boldsymbol{uqp} \prec \boldsymbol{qp}$ or $\boldsymbol{uqp} \succ \boldsymbol{qp}$ can be determined in $\leq 2\lambda$ steps.*

Lemma 4 *Let $\boldsymbol{x}$, $\boldsymbol{y}$ be two-pattern strings of scope λ, $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ an expansion, and $\boldsymbol{y} = \sigma(\boldsymbol{x})$. Let $\boldsymbol{u}$ be a non-empty binary string and let $\boldsymbol{up}$ be a suffix of a restrained configuration $\boldsymbol{qp}$ of $\boldsymbol{y}$. Let $1 \leq k$, and let $\boldsymbol{p}^k\boldsymbol{q}$ be a restrained configuration of $\boldsymbol{y}$. Then $\boldsymbol{up} \asymp \boldsymbol{p}^k\boldsymbol{q}$ and whether $\boldsymbol{up} \prec \boldsymbol{p}^k\boldsymbol{q}$ or $\boldsymbol{up} \succ \boldsymbol{p}^k\boldsymbol{q}$ can be determined in $\leq 2\lambda$ steps.*

Lemma 5 *Let $\boldsymbol{x}$, $\boldsymbol{y}$ be two-pattern strings of scope λ, $\sigma = [\boldsymbol{p}, \boldsymbol{q}, i, j]_\lambda$ an expansion, and $\boldsymbol{y} = \sigma(\boldsymbol{x})$. Let $\boldsymbol{u}$ be a non-empty binary string and let $\boldsymbol{up}^k\boldsymbol{q}$, $1 \leq k$, be a suffix of a restrained configuration $\boldsymbol{p}^{k+1}\boldsymbol{q}$ or $\boldsymbol{qp}^k\boldsymbol{q}$ of $\boldsymbol{y}$. Let $\boldsymbol{qp}$ be a restrained configuration of $\boldsymbol{y}$. Then $\boldsymbol{up}^k\boldsymbol{q} \asymp \boldsymbol{qp}$ and whether $\boldsymbol{up}^k\boldsymbol{q} \prec \boldsymbol{qp}$ or $\boldsymbol{up}^k\boldsymbol{q} \succ \boldsymbol{qp}$ can be determined in $\leq 2\lambda$ steps.*

6. Conclusion

Even though it is known that suffixes for all strings can be sorted in linear time using recursive algorithms, our research verified that for the class of complete two-pattern strings the sorting can be done iteratively, also in linear time. The analysis shows that the approach presented here is rather straightforward, thus providing

additional evidence of how two-pattern strings are well-suited for computational processing, the main goal of this effort.

Acknowledgements

The first author would like to acknowledge the support and hospitality of the School of Computing, Curtin University, Perth, Australia during the research for this paper. The research of both authors was supported in part by their respective research grants from the Natural Sciences and Engineering Research Council of Canada.

References

1. M. Farach, "Optimal suffix tree construction with large alphabets", in *Proc. 38th Annual Symposium on Foundations of Computer Science*, **IEEE** (1997) 137–143.
2. F. Franek, W. Lu, and W. F. Smyth, "Two-pattern strings I — a recognition algorithm", *J. Discrete Algorithms*, **1** (2003) 445–460.
3. F. Franek, W. Lu, and W. F. Smyth, "Two-pattern strings II — computing all repetitions and near-repetitions", submitted to *J. Discrete Algorithms*.
4. F. Franek and W. F. Smyth, "Sorting suffixes of two-pattern strings", *Technical Report* **CAS-04-09-FF**, Dept. of Comp. & Soft., McMaster University, October 2004.
5. F. Franek and W. F. Smyth, "Web supplement for paper: Sorting suffixes of two-pattern strings", *to be found at URL*: `http://www.cas.mcmaster.ca/~franek/web-publications.html`
6. P. Ko and S. Aluru, "Space efficient linear time construction of suffix arrays", Proceedings of the 14th Annual Symposium CPM, *LNCS* **2676**, Springer (2003) 200–210.
7. D. K. Kim, J. S. Sim, H. Park, and K. Park, "Linear-time construction of suffix arrays", Proceedings of the 14th Annual Symposium CPM, *LNCS* **2676**, Springer (2003) 186–199.
8. J. Kärkkäinen and P. Sanders, "Simple linear work suffix array construction", Proceedings of the 30th International Colloquium on Automata, Languages and Programming, *LNCS* **2719**, Springer (2003) 943–955.
9. U. Manber and G. Myers, "Suffix arrays: a new method for on-line string searches", *SIAM Journal on Computing* **22** (1993) 935–948.