



Murdoch
UNIVERSITY

MURDOCH RESEARCH REPOSITORY

*This is the author's final version of the work, as accepted for publication following peer review but without the publisher's layout or pagination.
The definitive version is available at :*

http://dx.doi.org/10.1007/3-540-45452-7_8

*Franěk, F., Jiang, J., Lu, W. and Smyth, W.F. (2002) Two-Pattern strings.
Lecture Notes in Computer Science, 2373 . pp. 76-84.*

<http://researchrepository.murdoch.edu.au/27872/>

Copyright: © 2002 Springer
It is posted here for your personal use. No further distribution is permitted.

Two-Pattern Strings^{*}

František Franěk¹, Jiandong Jiang^{1,2}, Weilin Lu^{1,2}, and W. F. Smyth^{1,3}

¹ Algorithms Research Group, Department of Computing & Software
McMaster University, Hamilton, Ontario, Canada L8S 4K1
smyth@mcmaster.ca
www.cas.mcmaster.ca/cas/research/groups.shtml

² Toronto Laboratories, IBM Canada, 8200 Warden Avenue,
Markham, Ontario, Canada L6G 1C7

³ School of Computing, Curtin University, GPO Box U-1987
Perth WA 6845, Australia

January 29, 2002

Abstract. This paper introduces a new class of strings on $\{a, b\}$, called *two-pattern strings*, that constitute a substantial generalization of Sturmian strings while at the same time sharing many of their nice properties. In particular, we show that, in common with Sturmian strings, only time linear in the string length is required to recognize a two-pattern string as well as to compute all of its repetitions. We also show that two-pattern strings occur in some sense frequently in the class of all strings on $\{a, b\}$.

1 Introduction

In this paper we outline the results of an investigation of the properties of a new class of strings on $\{a, b\}$, derived by the successive action of a sequence of morphisms on the single letter a . All of the strings so determined are finite, and we deal with them from a computational point of view: initially, we are interested in efficient algorithms to recognize such strings and to compute the repetitions in them; then we go on to estimate their frequency of occurrence among all strings on $\{a, b\}$.

A previous paper [FKS00] specified linear-time algorithms to recognize and compute repetitions in finite substrings of Sturmian strings; the class of strings discussed here significantly extends this work.

Let \mathbf{p} and \mathbf{q} denote two distinct nonempty strings on $\{a, b\}$ such that $|\mathbf{p}| \leq \lambda$ and $|\mathbf{q}| \leq \lambda$, where λ is a finite integer called the *scope*. We call

^{*} Supported in part by grants from the Natural Sciences & Engineering Research Council of Canada.

\mathbf{p} and \mathbf{q} **patterns of scope** λ . For any pair of finite positive integers i and j such that $i < j$, consider the morphism σ that maps single letters into **blocks**:

$$a \rightarrow \mathbf{p}^i \mathbf{q}, \quad b \rightarrow \mathbf{p}^j \mathbf{q}. \quad (1)$$

We call σ an **expansion of scope** λ and observe that it is specified by a 4-tuple $[\mathbf{p}, \mathbf{q}, i, j]$. Observe also that an expansion can be applied to any (finite or infinite) string on $\{a, b\}$ to yield an expanded string

$$\mathbf{y} = \sigma(\mathbf{x}).$$

Given any two morphisms σ_1 and σ_2 , the composition $\sigma_1 \circ \sigma_2$ is therefore well defined: $\mathbf{z} = \sigma_1(\mathbf{y}) = \sigma_1(\sigma_2(\mathbf{x})) = (\sigma_1 \circ \sigma_2)(\mathbf{x})$.

Definition 1. *Suppose a positive integer λ and a finite sequence*

$$\sigma_1, \sigma_2, \dots, \sigma_k$$

of expansions of scope λ are given, where

$$\sigma_r = [\mathbf{p}_r, \mathbf{q}_r, i_r, j_r]$$

for every $r = 1, 2, \dots, k$. Then the string

$$\mathbf{x} = (\sigma_1 \circ \sigma_2 \circ \dots \circ \sigma_k)(a)$$

*is a **complete two-pattern string of scope** λ if and only if every pair $(\mathbf{p}_r, \mathbf{q}_r)$ of patterns is **suitable** (defined in Section 2).*

The definition of a suitable pair of patterns is deferred till Section 2 because it is necessarily somewhat technical. However, the main idea of a suitable pair is simple: \mathbf{p} and \mathbf{q} should be dissimilar enough that they can be efficiently distinguished from each other by an algorithm that recognizes complete two-pattern strings.

We can easily provide examples of complete two-pattern strings. If we suppose that $\lambda = 3$ and $\sigma_1 = [ab, ba, 2, 3]$, $\sigma_2 = [abb, aa, 1, 4]$, the following strings are all complete two-pattern strings of scope 3:

$$\begin{aligned} \sigma_1(a) &= (ab)^2ba; \\ (\sigma_1 \circ \sigma_1)(a) &= (ab)^2ba(ab)^3ba(ab)^2ba(ab)^3ba(ab)^3ba(ab)^2ba; \\ (\sigma_1 \circ \sigma_2)(a) &= (ab)^2ba(ab)^3ba(ab)^3ba(ab)^2ba(ab)^2ba; \\ (\sigma_2 \circ \sigma_1)(a) &= (abb)aa(abb)^4aa(abb)aa(abb)^4aa(abb)^4aa(abb)aa. \end{aligned}$$

Observe further that when the scope $\lambda = 1$, the choice of \mathbf{p} and \mathbf{q} is restricted to

$$(\mathbf{p}, \mathbf{q}) = (a, b) \text{ or } (\mathbf{p}, \mathbf{q}) = (b, a). \quad (2)$$

If the further restriction is imposed that $j = i+1$, then all the strings generated by any finite sequence of expansions are finite substrings of Sturmian strings; in fact, in the terminology of [FKS00], these strings are exactly the set of “block-complete” finite substrings of Sturmian strings. We note that every complete two-pattern string of scope λ is also a complete two-pattern string of scope $\lambda+1$; thus in particular every block-complete finite substring of a Sturmian string is a complete two-pattern string.

As noted above, our initial interest in complete two-pattern strings is computational, following similar studies of Fibonacci [IMS97,FS99,KK00] and Sturmian [BF84,FKS00] strings. We pose two sets of questions:

- (Q1) What is the complexity of determining whether or not a given string $\mathbf{x} = \mathbf{x}[1..n]$ is a fragment of a complete two-pattern string? Can an efficient algorithm be found to make this determination for every \mathbf{x} ?
- (Q2) Given a fragment \mathbf{x} of a complete two-pattern string, can an algorithm be found that computes all the repetitions in \mathbf{x} in linear time?

Since complete two-pattern strings constitute a much more general class of strings than block-complete finite substrings of Sturmian strings, the following questions also become of interest:

- (Q3) What is the frequency of occurrence of fragments \mathbf{x} of complete two-pattern strings among all strings on $\{a, b\}$ of length n ? What is the asymptotic frequency of occurrence of complete two-pattern strings among all infinite strings on $\{a, b\}$?

In this paper we provide a partial answer to (Q1) by outlining an algorithm that in $\Theta(n)$ time determines whether or not a given string $\mathbf{x}[1..n]$ is complete two-pattern. Similar to the recognition algorithm in [FKS00], this algorithm outputs the sequence of expansions (1) by which a is transformed into \mathbf{x} — or more precisely, the sequence of *reductions*

$$\mathbf{p}^i \mathbf{q} \rightarrow a, \quad \mathbf{p}^j \mathbf{q} \rightarrow b \quad (3)$$

by which \mathbf{x} is reduced to a . This sequence provides a complete specification of \mathbf{x} . Since by (1) each reduction decreases string length by a factor that exceeds

$$i|\mathbf{p}|+|\mathbf{q}| \geq 2, \quad (4)$$

the recognition algorithm thus yields as a byproduct a potential data compression technique for complete two-pattern strings \boldsymbol{x} .

The reduction sequence is then used to provide partial answers to (Q2) and (Q3). Before going on to discuss these questions in more detail, we pause to provide an introduction and context for them, as well as an outline of the main results.

In dealing with (Q1), we need to cope with the possibility that at any stage of the reduction of \boldsymbol{x} , there may be more than one reduction satisfying (3): it then becomes possible that one of these reductions is a part of a sequence that reduces \boldsymbol{x} to a , while another one is not. As long as this possibility exists, any recognition algorithm would be obliged to include provision for backtracking, leading possibly to an execution time exponential in the number of reductions. Our main result in this connection is to show however that backtracking is not required, and that therefore the algorithm that recognizes complete two-pattern strings requires only $\Theta(n)$ time.

Of course this does not yet fully solve (Q1). We conjecture that, just as for Sturmian strings [FKS00], there exists a $\Theta(n)$ -time algorithm to determine whether or not a string is a *fragment* of a complete two-pattern string of scope λ . If such an algorithm were found, it would greatly extend the class of strings that could be efficiently compressed using reductions, or whose repetitions could be efficiently computed.

The view may be taken that interest in (Q2) has been superseded by other work. It has recently become clear that, as a result of research extending over a period of a quarter-century, the repetitions in any string $\boldsymbol{x}[1..n]$ on an *indexed* alphabet — that is, an alphabet of size $\alpha \in O(n)$ that maps onto the integers $1..\alpha$ — can be computed in $\Theta(n)$ time. The main steps in this development are as follows:

- an algorithm to compute the suffix tree of \boldsymbol{x} in $\Theta(n)$ time [F97];
- an algorithm to compute the s -factorization of \boldsymbol{x} , given the suffix tree of \boldsymbol{x} , in $\Theta(n)$ time [LZ76,ZL77];
- the identification of “maximal periodicities” or “runs” as a suitable encoding of repetitions in strings, and the computation of the leftmost occurrence of every distinct run in \boldsymbol{x} in $\Theta(n)$ time, based on the s -factorization [M89];
- the proof that the number of runs in any string is $O(n)$, and the extension of the algorithm [M89] to compute all occurrences of every run in \boldsymbol{x} in $\Theta(n)$ time, still based on the s -factorization [KK00].

Impressive as this intellectual edifice is, it nevertheless appears, at least in the context of strings on the alphabet $\{a, b\}$, to be rather indirect

in its approach, perhaps involving more sophistication than is really required. Indeed, it is not clear that the $\Theta(n)$ -time algorithm given in [F97] is preferable in practice to classical $O(n \log n)$ -time algorithms for suffix-tree construction. Further, the very long and technical proof that number of runs is linear in string length shows that a constant of proportionality exists, but provides no information about its size; at the same time, computer experiments described in [KK00] provide convincing evidence that the maximum number of runs in any string is at most n , and that this maximum occurs in strings on $\{a, b\}^*$. Thus, in a sense, the existing theory serves to remind us of how little, rather than how much, we know of periodicity in strings, perhaps especially those on $\{a, b\}$.

For (Q2) we adopt a more direct approach, an extension of the methodology used in [FKS00] for Sturmian strings. Making use of the reduction sequence computed by the recognition algorithm, we show how to compute all the runs in complete two-pattern strings $\mathbf{x}[1..n]$ in $\Theta(n)$ time. Essentially, we show that if \mathbf{y} is derived from \mathbf{x} by a reduction (3), then the nontrivial runs in \mathbf{x} can be computed directly from certain special configurations occurring in \mathbf{y} ; thus, over the whole reduction sequence, the runs in \mathbf{x} can be computed on a step-by-step basis, from one reduction to the next. It is the special configurations that are of interest here, since they provide insight into the way in which repetitions are formed.

Finally, we report on progress with (Q3), in estimating the frequency of occurrence of complete two-pattern strings among all strings on $\{a, b\}$. We claim that for λ sufficiently large with respect to n , complete two-pattern strings are dense in the set of all strings. We claim also that for some values of k and fixed λ , the number of distinct strings of length k (the *complexity*) can exceed $2k$ — can in fact even be exponential in k .

Sections 2-4 deal with questions (Q1)-(Q3) respectively.

2 Recognizing Two-Pattern Strings in Linear Time

Before proceeding with our development, we need to provide a definition of the term “suitable pair” mentioned in Section 1:

Definition 2. *A string \mathbf{q} is said to be **p-regular** if and only if there are strings $\mathbf{u} \neq \varepsilon$, \mathbf{v} together with nonnegative integers n_1, \dots, n_k , $k \geq 1$, and r such that*

- \mathbf{p} is neither a prefix nor a suffix of \mathbf{u} ;
- \mathbf{p} is neither a prefix nor a suffix of \mathbf{v} ;
- there are at most two integer values m_1 and m_2 such that for each $i \in 1..k$, $n_i = m_1$ or $n_i = m_2$, i.e. $|\{n_i : i \in 1..k\}| \leq 2$;

is by (4) at most $\lceil \log_2 n \rceil$ and may be much smaller. In our example, $n = 124$ and sequence length $3 = \log_{4.39} 124$.

We now introduce formally an idea mentioned in the introduction: a canonical reduction that is identified by patterns that are somehow “shortest”:

Definition 4. A reduction $\rho = [\mathbf{p}, \mathbf{q}, i, j]$ of a binary string \mathbf{x} using patterns of scope λ is **λ -canonical** if and only if for every reduction $\rho_1 = [\mathbf{p}_1, \mathbf{q}_1, i_1, j_1]$ of \mathbf{x} using patterns of scope λ :

- (a) either $|\mathbf{p}| < |\mathbf{p}_1|$, or $|\mathbf{p}| = |\mathbf{p}_1|$ and $|\mathbf{q}| < |\mathbf{q}_1|$, or $|\mathbf{p}| = |\mathbf{p}_1|$ and $|\mathbf{q}| = |\mathbf{q}_1|$.
- (b) $\mathbf{x} = \mathbf{p}_1^{i_1} \mathbf{q}_1$ implies $\mathbf{x} = \mathbf{p}^i \mathbf{q}$.

It is then possible to prove that it suffices to reduce \mathbf{x} using a sequence of canonical reductions:

Theorem 1. \mathbf{x} is a complete two-pattern string of scope λ if and only if there is a sequence of λ -canonical reductions $\{\rho_1, \rho_2, \dots, \rho_n\}$ reducing \mathbf{x} to a string a . \square

We omit the fairly predictable details of the algorithm REC that is based on this theorem. We state however the main result:

Theorem 2. For any $\lambda \geq 1$, the recognition algorithm REC determines in $O(2\lambda^8 |\mathbf{x}|)$ steps whether or not \mathbf{x} is a complete two-pattern string of scope λ , and if so, the algorithm outputs the λ -canonical reduction sequence of \mathbf{x} . \square

3 Computing the Repetitions in Linear Time

We describe here the main ideas that permit the repetitions in a complete two-pattern string to be computed in linear time. Just as for Sturmian strings [FKS00], it turns out that the repetitions that occur in an expansion $\mathbf{y} = \sigma(\mathbf{x})$ of a two-pattern string \mathbf{x} are formed as a result of the application of σ to certain well-defined configurations in \mathbf{x} . Thus, with the help of the expansion (reduction) sequence that determines a complete two-pattern string, it is possible to track and output the repetitions as they are formed by each expansion in the sequence. As mentioned in the introduction, a crucial factor that ensures the efficiency of this process is the encoding of the repetitions as runs.

Definition 5. A *repetition* in $\mathbf{x} = \mathbf{x}[1..n]$ [C81] is a triple (i, p, r) of positive integers, where $i < n$, $r > 1$,

$$\mathbf{x}[i..i+rp-1] = \mathbf{x}[i..i+p-1]^r,$$

and $\mathbf{x}[i+rp..i+(r+1)p-1] \neq \mathbf{x}[i..i+p-1]$. The **period** of the repetition is p and its **generator** is $\mathbf{x}[i..i+p-1]$.

Crochemore [C81] showed that Fibonacci strings, a special case of two-pattern strings, contain $\Omega(n \log n)$ repetitions. To avoid $\Omega(n \log n)$ processing just for output, we therefore introduce:

Definition 6. A *run* in $\mathbf{x}[1..n]$ [M89] is a 4-tuple (i, p, r, t) where

$$(i, p, r), (i+1, p, r), \dots, (i+t, p, r)$$

are all repetitions, while $(i-1, p, r)$ and $(i+t+1, p, r)$ are not. The **period** and **generator** are defined as for (i, p, r) .

It is easy to prove that for any constant κ , all the runs in \mathbf{x} whose period $p \leq \kappa$ can be output in at most $c_\kappa n$ steps, where c_κ is a constant whose value depends only on κ . In particular, this result is true for the choice $\kappa = 3\lambda$. For $p > 3\lambda$, we require the following lemma, the main result of this section:

Lemma 1. For $|\mathbf{p}| \leq \lambda$, $|\mathbf{q}| \leq \lambda$, let $\sigma = [\mathbf{p}, \mathbf{q}, i, j]$ be an expansion of \mathbf{x} , so that $\mathbf{y} = \sigma(\mathbf{x})$. Then for every run R in \mathbf{y} whose period $p > 3\lambda$, one of the following holds:

- R is an expansion under σ of a run in \mathbf{x} ,
- R is determined by a square \mathbf{u}^2 in \mathbf{x} that is derived from a substring of \mathbf{x} of one of the following forms:

$$aa, ab, ba, bb, avbva, bvavb, bvaavb, avbv, bvav,$$

for any nonempty substring v . \square

The details of “deriving” \mathbf{u}^2 and of using it to “determine” R are lengthy and complicated, to be found at web site

<http://www.cas.mcmaster.ca/~franek/>

The analysis found there enables us to claim that

Theorem 3. There exists an algorithm *RUN* such that for every integer $\lambda \geq 1$, *RUN* computes all the runs in every complete two-pattern string $\mathbf{x} = \mathbf{x}[1..n]$ of scope λ , based on the reduction sequence of \mathbf{x} , in at most $c_\lambda n$ steps, where c_λ is a constant whose value depends only on λ . \square

4 Frequency of Two-Pattern Strings

In this section we first present results showing that infinite two-pattern strings (that is, two-pattern strings formed from an infinite sequence of expansions) have complexity at least $k+1$, while for some values of k the complexity can even be exponential in k . Since Sturmian strings have complexity $k+1$, we can accordingly claim that two-pattern strings are in some sense more frequent among all strings on $\{a, b\}$ than Sturmian strings are. Nevertheless, for fixed λ the relative frequency of two-pattern strings approaches zero as string length approaches infinity.

Another point of view is also of interest. We find that if we consider only those values of n that are close to λ , then two-pattern strings occur frequently among all strings of length n . In other words, for sufficiently large λ , a large proportion of strings on $\{a, b\}$ turn out to be two-pattern strings.

We begin by recalling the notation $\text{LCP}(\mathbf{u}, \mathbf{v})$ and $\text{LCS}(\mathbf{u}, \mathbf{v})$ for arbitrary strings \mathbf{u} and \mathbf{v} : longest common prefix and longest common suffix, respectively. It is then convenient to define, for any integer $m \geq 0$,

$$\Delta_m = \Delta_m(\mathbf{u}, \mathbf{v}) = m|\mathbf{p}| + |\text{LCP}(\mathbf{u}, \mathbf{v})| + |\text{LCS}(\mathbf{u}, \mathbf{v})|.$$

Using this notation, we state:

Theorem 4. *Let \mathbf{x} be an infinite two-pattern string of scope λ reducible by $[\mathbf{p}, \mathbf{q}, i, j]$. Then the complexity $C_k = C_k(\mathbf{x})$ satisfies*

- (a) $C_k \geq k+1$ when $|\mathbf{p}| \leq k \leq \Delta_i$;
- (b) $C_k \geq 2k - \Delta_i$ when $\Delta_i + 1 \leq k \leq \Delta_{j-1} + 1$;
- (c) $C_k \geq k+1 + (j-i-1)|\mathbf{p}|$ when $k \geq \Delta_{j-1} + 2$;

where $\Delta_m = \Delta_m(\mathbf{p}, \mathbf{q})$. \square

This result provides lower bounds on the complexity C_k that are in fact sharp. We have also established upper bounds on C_k that are however not sharp. To show that the complexity can for some values of k be much larger than $k+1$, consider an infinite two-pattern string \mathbf{x} with substring \mathbf{pqp} , where $\mathbf{p} = \text{aaaabbbb}$ and $\mathbf{q} = \text{aababbab}$ are a suitable pair of patterns. It is easy to check that in the substring \mathbf{pqp} there are 2^4 substrings of length 4 — for $k = 4$ the complexity is exponential in k .

In order to state our final results, we introduce the constant [GO81]

$$\phi = \lim_{n \rightarrow \infty} P(n)/2^n \approx 0.26778684,$$

where $P(n)$ is the number of primitive strings (with no nonempty border) of length n on $\{a, b\}$. The frequency of occurrence of two-pattern strings for λ large with respect to n can then be estimated in terms of ϕ :

Theorem 5. For $n \geq 2$, let $f(n)$ denote the frequency of occurrence among all strings of length n of two-pattern strings $\mathbf{x}[1..n] = \mathbf{p}\mathbf{q}$, where \mathbf{p}, \mathbf{q} is a suitable pair of scope $\lambda \geq \lceil n/2 \rceil$. Then $f(n) \geq \phi/2$. \square

Theorem 6. For $n \geq 4$, let $f(n)$ denote the frequency of occurrence among all strings of length n of two-pattern strings $\mathbf{x}[1..n] = \mathbf{p}^i \mathbf{q}$, where $i \geq 1$ and \mathbf{p}, \mathbf{q} is a suitable pair of scope $\lambda \geq n-2$. Then $f(n) \geq 15\phi/16$. \square

References

- [BF84] M. Boshernitzan & Aviezri S. Fraenkel, **A linear algorithm for nonhomogeneous spectra of numbers**, *J. Algorithms* 5 (1984) 187-198.
- [C81] Maxime Crochemore, **An optimal algorithm for computing the repetitions in a word**, *IPL* 12-5 (1981) 244-250.
- [F97] Martin Farach, **Optimal suffix tree construction with large alphabets**, *Proc. 38th Annual IEEE Symp. FOCS* (1997) 137-143.
- [FS99] Aviezri S. Fraenkel & R. Jamie Simpson, **The exact number of squares in Fibonacci words**, *TCS* 218-1 (1999) 83-94.
- [FKS00] František Franěk, Ayşe Karaman & W. F. Smyth, **Repetitions in Sturmian strings**, *TCS* 249-2 (2000) 289-303.
- [GO81] Leo J. Guibas & Andrew M. Odlyzko, **Periods in strings**, *J. Combinatorial Theory, Series A* 30 (1981) 19-42.
- [IMS97] Costas S. Iliopoulos, Dennis Moore & W. F. Smyth, **A characterization of the squares in a Fibonacci string**, *TCS* 172 (1997) 281-291.
- [KK00] Roman Kolpakov & Gregory Kucherov, **On maximal repetitions in words**, *J. Discrete Algorithms* 1 (2000) 159-186.
- [LZ76] Abraham Lempel & Jacob Ziv, **On the complexity of finite sequences**, *IEEE Trans. Information Theory* 22 (1976) 75-81.
- [M89] Michael G. Main, **Detecting leftmost maximal periodicities**, *Discrete Applied Maths.* 25 (1989) 145-153.
- [ZL77] Jacob Ziv & Abraham Lempel, **A universal algorithm for sequential data compression**, *IEEE Trans. Information Theory* 23 (1977) 337-343.