



On arithmetically progressed suffix arrays and related Burrows–Wheeler transforms[☆]

Jacqueline W. Daykin^{a,b,c}, Dominik Köppl^{d,*}, David Kübel^e, Florian Stober^f

^a Department of Computer Science, Aberystwyth University, UK

^b Univ Rouen Normandie, INSA Rouen Normandie, Université Le Havre Normandie, Normandie Univ, LITIS, 76000 Rouen, France

^c Department of Information Science, Stellenbosch University, South Africa

^d Department of Computer Science and Engineering, University of Yamanashi, Kōfu, Japan

^e University of Bonn, Institute of Computer Science, Germany

^f University of Stuttgart, Institute for Formal Methods of Computer Science (FMI), Germany

ARTICLE INFO

Article history:

Received 30 October 2021

Received in revised form 19 October 2023

Accepted 16 April 2024

Available online xxxx

Keywords:

Arithmetic progression

Burrows–Wheeler transform

Christoffel words

Suffix array

String combinatorics

ABSTRACT

We characterize those strings whose suffix arrays are based on arithmetic progressions, in particular, arithmetically progressed permutations where all pairs of successive entries of the permutation have the same difference modulo the respective string length. We show that an arithmetically progressed permutation P coincides with the suffix array of a unary, binary, or ternary string. We further analyze the conditions of a given P under which we can find a uniquely defined string over either a binary or ternary alphabet having P as its suffix array. For the binary case, we show its connection to lower Christoffel words, balanced words, and Fibonacci words. In addition to solving the arithmetically progressed suffix array problem, we give the shape of the Burrows–Wheeler transform of those strings solving this problem. These results give rise to numerous future research directions.

© 2024 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

1. Introduction

The integral relationship between the suffix array [1] (SA) and Burrows–Wheeler transform [2] (BWT) is explored by Adjeroh et al. [3], who also illustrate the versatility of the BWT beyond its original motivation in lossless block compression [2]. The BWT has an injective property, that is, it computes a One-to-One mapping from the input to the transform, which enables recovery of input data in lossless compression scenarios. BWT applications include compressed index structures using backward search pattern matching, multimedia information retrieval, bioinformatics sequence processing, and it is at the heart of the bzip2 suite of text compressors. By its association with the BWT, this also indicates the importance of the SA data structure and hence our interest in exploring its combinatorial properties.

These combinatorial properties can be useful when checking the performance or integrity of string algorithms or data structures on string sequences in testbeds when properties of the employed string sequences are well understood. In particular, due to current trends involving massive data sets, indexing data structures need to work in external memory (e.g., [4]), or on distributed systems (e.g. [5]). For devising a solution adaptable to these scenarios, it is crucial to test whether the computed index (consisting of the suffix array or the BWT, for instance) is correctly stored on the hard disk or on the computing nodes, respectively. This test is more cumbersome than in the case of a single machine working only

[☆] This paper is an extension of a contribution to the Prague Stringology Conference 2020, published in Daykin et al. (2020).

* Corresponding author.

E-mail addresses: jwd6@aber.ac.uk (J.W. Daykin), dkppl@yamanashi.ac.jp (D. Köppl), florian.stober@fmi.uni-stuttgart.de (F. Stober).

with its own RAM. One way to test is to compute the index for an instance, whose index shape can be easily verified. For example, one could check the validity of the computed BWT on a Fibonacci word since the shape of its BWT is known [6–8].

Other studies based on Fibonacci words are the suffix tree [9] or the Lempel–Ziv 77 (LZ77) factorization [10]. In [11], the suffix array and its inverse of each even Fibonacci word is studied as an arithmetic progression. In this study, the authors, like many at that time, did not append the artificial \$ delimiter (also known as a sentinel) to the input string, thus allowing suffixes to be prefixes of other suffixes. This small fact makes the definition of $BWT_T[i] = T[SA_T[i] - 1]$ for a string T with suffix array SA_T incompatible with the traditional BWT defined on the BWT matrix, namely the lexicographic sorting of all cyclic rotations of the string T . Note that throughout this paper we use indices starting from 1. Consequently, we expect $i \in [1..n]$ in the preceding definition and for convenience we define $T[0] := T[n]$.

For instance, $SA_{bab} = [2, 3, 1]$ with $BWT_{bab} = bab$, while $SA_{bab\$} = [4, 2, 3, 1]$ and $BWT_{bab\$} = bba\$$ with $\$ < a < b$. However, the traditional BWT constructed by reading the last characters of the lexicographically sorted cyclic rotations $[abb, bab, bba]$ of bab yields bba , which is equal to $BWT_{bab\$} = bba\$$ after removing the \$ character.

Note that not all strings are in the BWT image. An $O(n \log n)$ -time algorithm is given by Giuliani et al. [12] for identifying all the positions in a string S where a \$ can be inserted into so that S becomes the BWT image of a string ending with \$.

Despite this incompatibility in the suffix array based definition of the BWT, we can still observe a regularity for even Fibonacci words [11, Sect. 5]. Similarly, both methods for constructing the BWT are compatible when the string T consists of a Lyndon word. The authors of [11, Remark 1] also observed similar characteristics for other, more peculiar string sequences. For the general case, the \$ delimiter makes both methods equivalent, however the suffix array approach is typically preferred as it requires $O(n)$ time [13] compared to $O(n^2)$ with the BWT matrix method [2]. By utilizing combinatorial properties of the BWT, an in-place algorithm is given by Crochemore et al. [14], which avoids the need for explicit storage for the suffix sort and output arrays, and runs in $O(n^2)$ time using $O(1)$ extra memory (apart from storing the input text). Köppl et al. [15, Sect. 5.1] adapted this algorithm to compute the traditional BWT within the same space and time bounds.

Up to now, it has remained unknown whether we can formulate a class of string sequences for which we can give the shape of the suffix array as an arithmetic progression (independent of the \$ delimiter). With this article, we catch up on this question, and establish a correspondence between strings and suffix arrays generated by arithmetic progressions. Calling a permutation of integers $[1..n]$ *arithmetically progressed* if all pairs of successive entries of the permutation have the same difference modulo n , we show that an arithmetically progressed permutation coincides with the suffix array of a unary, binary or ternary string. We analyze the conditions of a given arithmetically progressed permutation P under which we can find a uniquely defined string T over either a unary, a binary, or ternary alphabet having P as its suffix array.

The simplest case is for unary alphabets: Given the unary alphabet $\Sigma := \{a\}$ and a string T of length n over Σ , $SA_T = [n, n - 1, \dots, 1]$ is an arithmetically progressed permutation with ratio $-1 \equiv n - 1 \pmod n$.

For the case of a binary alphabet $\{a, b\}$, several strings of length n exist that solve the problem. Trivially, the solutions for the unary alphabet also solve the problem for the binary alphabet. However, studying those strings of length n whose suffix array is $[n, n - 1, \dots, 1]$, there are now multiple solutions: each $T = b^r a^s$ with $r, s \in [0..n]$ such that $r + s = n$ has this suffix array. Similarly, $T = a^{n-1}b$ has the suffix array $SA_T = [1, 2, \dots, n]$, which is an arithmetically progressed permutation with ratio 1.

A non-lexicographic ordering on strings, *the V-order*, considered for a modified FM-index [16], provides a curious example for the case with ratio -1 : if S is a proper subsequence of a string T of length n , then S precedes T in V -order, written $S <_V T$. This implies that $T[n] <_V T[n - 1..n] <_V \dots <_V T[1..n]$, so that $SA_T[i] = n - i + 1$ for every $i \in [1..n]$, thus enabling trivial suffix sorting.

Clearly, any string ordering method that prioritizes minimal length within its definition will have a suffix array progression ratio of -1 . Binary Gray codes¹ have this property and are ordered so that adjacent strings differ by exactly one bit [17]. These codes, which exhibit non-lexicographic order, have numerous applications, notably in error detection schemes such as flagging unexpected changes in data for digital communication systems, and logic circuit minimization. While it seems that Gray code order has not been applied directly to order the rows in a BWT matrix, just four years after the BWT appeared in 1994 [2], Chapin and Tate applied the concept when they investigated the effect of both alphabet and string reordering on BWT-based compressibility [18]. Their string sorting technique for the BWT matrix ordered the strings in a manner analogous to reflected Gray codes, but for more general alphabets. This modification inverted the sorting order for alternating character positions with which they demonstrated improved compression.

To date the only known BWT designed specifically for binary strings is the binary block order B -BWT [19]. This binary string sorting method prioritizes length, thus also exhibiting a suffix array with progression ratio -1 .² Ordering strings of the same length, as applicable to forming the BWT matrix, is by a play on decreasing ($>$) and increasing ($<$) lexicographic ordering of the run length exponents of blocks of bits. Experimentation showed roughly equal performance to the original BWT in binary lexicographic order when comparing the number of character runs; however, there are a few instances where the B -BWT has significantly fewer runs.

¹ Originally known as *Reflected Binary Code* by the inventor Frank Gray.

² To see that the binary block order and the Gray code order are distinct: 1101 comes before 1110 in Gray code order, whereas, 1110 comes before 1101 in binary block order.

In what follows, we present a comprehensive analysis of strings whose suffix arrays are arithmetically progressed permutations (under the standard lexicographic order). In practice, such knowledge can reduce the $O(n)$ space for the suffix array to $O(1)$. Compared to the conference version [20], we here additionally analyze the shapes of BWTs of strings whose suffix arrays are arithmetically progressed. Additionally, we give applications for Christoffel words, balanced words, and meta strings. Finally, we extend our study on binary and ternary alphabets to general alphabets.

Arithmetic progressions have arisen in further applications of strings, and we conclude this section by mentioning a few. An effective approach in search-based tasks can be to reduce convolution problems, such as pattern matching with don't cares, or approximate matching, to polynomial multiplication. The practicality of complex convolutions-based bounded divide-and-conquer algorithms was investigated in [21] which demonstrated that, despite its theoretical efficiency, the convolution approach may not be practical across all problems and trade-offs. This led to devising a sequences method which was shown for some real problems to be a superior method to convolutions. In their experimentation, Amir et al. [21] included arithmetically progressed sequences. A lossless coding method distinct from Lempel–Ziv coding, known as *Arithmetic Progressions Tree coding* (APT), was proposed in [22]. The APT coding method is based on a tree-structure principle and involves finding arithmetic progressions in a given binary string. Empirical evidence showed APT coding to be asymptotically optimal up to a constant factor. Finding periodicity in real-world data often leads to insights on the structure of the data, and can be useful in prediction of future events or flagging anomalies in systems. Input to periodicity detection algorithms is typically a sequence of events, each associated with a timestamp where in most practical applications, these timestamps are imprecise – hence the study of approximate periodic arithmetic progressions. Motivated by naturally occurring periodic events in time intervals, such as load patterns on web servers, an algorithm to find the longest ϵ -relative error periodic pattern in a sequence is presented in [23] and runs in sub-cubic time for many practical situations. Here, reflecting the nature of real sequence data, the constant difference between any two consecutive numbers in the longest subsequence found allows for a tolerance of $\epsilon \in [0, 1)$ thus determining an approximate arithmetic progression.

The structure of the paper is as follows.³ In Section 2 we give the basic definitions and background, and also deal with the elementary case of a unary alphabet. We present the main results in Section 3, where we cover ternary and binary alphabets, and consider inverse permutations. In Section 4, we illustrate the binary case for Christoffel words and in particular establish that every (lower) Christoffel word has an arithmetically progressed suffix array. We go on to link the binary characterization to balanced words and Fibonacci words. A characterization of strings with larger alphabets follows in Section 5. We overview the theme concepts for meta strings in Section 6. We conclude in Section 7 and propose a list of open problems and research directions, showing there is plenty of scope for further investigation. We proceed to the foundational concepts presented in the following section, starting with the case of a unary alphabet.

2. Preliminaries

Let Σ be an alphabet with size $\sigma := |\Sigma|$. An element of Σ is called a *character*.⁴ Let Σ^+ denote the set of all nonempty finite strings over Σ . The *empty string* of length zero is denoted by ε ; we write $\Sigma^* = \Sigma^+ \cup \{\varepsilon\}$. Given an integer $n \geq 1$, a *string*⁵ of length n over Σ takes the form $T = t_1 \cdots t_n$ with each $t_i \in \Sigma$. We write $T = T[1..n]$ with $T[i] = t_i$. The length n of a string T is denoted by $|T|$. If $T = uwv$ for some strings $u, w, v \in \Sigma^*$, then u is a *prefix*, w is a *substring*, and v is a *suffix* of T ; we say u (resp. w and v) is *proper* if $u \neq T$ (resp. $w \neq T$ and $v \neq T$). We say that a string T of length n has *period* $p \in [1..n - 1]$ if $T[i] = T[i + p]$ for every $i \in [1..n - p]$ (note that we allow periods larger than $n/2$). If $T = uv$, then vu is said to be a *cyclic rotation* of T . A string T is said to be a *repetition* if and only if it has a factorization $T = T^k$ for some integer $k > 1$; otherwise, T is said to be *primitive*. A string that is both a proper prefix and a proper suffix of a string $T \neq \varepsilon$ is called a *border* of T ; a string is *border-free* if the only border it has is the empty string ε .

If Σ is a totally ordered alphabet with order $<$, then this order $<$ induces the *lexicographic ordering* $<$ on Σ^* such that $u < v$ for two strings $u, v \in \Sigma^*$ if and only if either u is a proper prefix of v , or $u = ras$, $v = rbt$ for two characters $a, b \in \Sigma$ such that $a < b$ and for some strings $r, s, t \in \Sigma^*$. In the following, we select a totally ordered alphabet Σ having three characters a, b, c with $a < b < c$.

A string T is a *Lyndon word* if it is strictly least in the lexicographic order among all its cyclic rotations [24]. For instance, $abcac$ and $aaacbaabaaacc$ are Lyndon words, while the string $aaacbaabaaac$ with border $aaac$ is not.

A reciprocal relationship exists between the suffix array of a text and its Lyndon factorization, that is, the unique factorization of the text by greedily choosing the maximal length Lyndon prefix while processing the text from start to end: the Lyndon factorization of a text can be obtained from its suffix array [25]; conversely, the suffix array of a text can be constructed iteratively from its Lyndon factorization [26].

Lyndon words have numerous applications in combinatorics and algebra, and prove challenging entities due to their non-commutativity. Additionally, Lyndon words can arise naturally in Big Data – an instance in a biological sequence

³ Compared to the conference version, we added applications to Christoffel words and balanced words, generalized our results for larger alphabets, and gave examples for indeterminate strings.

⁴ Also known as *letter* or *symbol* in the literature.

⁵ Also known as *word* in the literature.

over the DNA alphabet, with $\Sigma = \{A, C, G, T\}$ and $A < C < G < T$, is the following substring occurring in a SARS-CoV-2 genome⁶:

AAAAACAGTAAAGTACAAATAGGAGAGTACACCTTTGAAAAAGGTGACTATGGTGAT

For the rest of the article, we take a string T of length $n \geq 2$. The suffix array $SA := SA_T[1..n]$ of T is a permutation of the integers $[1..n]$ such that $T[SA[i]..n]$ is the i th lexicographically smallest suffix of T . We denote with ISA its inverse, i.e., $ISA[SA[i]] = i$. By definition, ISA is also a permutation. The string BWT with $BWT[i] = T[SA[i] - 1 \bmod n]$ is the (SA-based) BWT of T . Here we shift the usual image $[0..n - 1]$ of $x \mapsto x \bmod n$ by one to admit values in $[1..n]$. Formally, $x \bmod n := x$ if $1 \leq x \leq n$, $x - n \bmod n$ for an integer $x > n$, and $x + n \bmod n$ for $x < 1$. In particular, $n \bmod n = 0 \bmod n = n$.

The focus of this paper is on arithmetic progressions. An *arithmetic progression* is a sequence of numbers such that the differences between all two consecutive terms are of the same value: Given an arithmetic progression $\{p_i\}_{i \geq 1}$, there is an integer $k \geq 1$ such that $p_{i+1} = p_i + k$ for all $i \geq 1$. We call k the *ratio* of this arithmetic progression. Similarly to sequences, we can define permutations that are based on arithmetic progressions: An *arithmetically progressed permutation* with ratio $k \in [1..n - 1]$ is an array $P := [p_1, \dots, p_n]$ with $p_{i+1} = p_i + k \bmod n$ for all $i \in [1..n]$, where we stipulate that $p_{n+1} := p_1$.⁷ In what follows, we want to study (a) strings whose suffix arrays are arithmetically progressed permutations, and (b) the shape of these suffix arrays. For a warm-up, we completely classify the case when the alphabet is unary with the following theorem.

Theorem 2.1. *Given the unary alphabet $\{a\}$, the suffix array of a string of length n over $\{a\}$ is uniquely defined by the arithmetically progressed permutation $[n, n - 1, \dots, 1]$ with ratio $n - 1$.*

From now on, we stick to a general alphabet. While the suffix array $[n, n - 1, \dots, 1]$ in the statement of [Theorem 2.1](#) uniquely defines the input string if the alphabet is unary, this is no longer the case if we allow larger alphabet sizes: Given the arithmetically progressed permutation $P = [n, n - 1, \dots, 1]$, we want to know the number of strings from a general totally ordered alphabet $\Sigma = [1..\sigma]$ with the natural order $1 < 2 < \dots < \sigma$, having P as their suffix array. For that, we fix a string T of length n with $SA_T = P$. Let $s_j \geq 0$ be the number of occurrences of the character $j \in \Sigma$ appearing in T . Then $\sum_{j=1}^{\sigma} s_j = n$. For all $j < k$ it holds that $T[j] \geq T[k]$ because k appears before j in SA_T . Therefore, $T = \sigma^{s_{\sigma}} (\sigma - 1)^{s_{\sigma-1}} \dots 1^{s_1}$ such that the position of the characters are uniquely determined.⁸ In other words, we can reduce this problem to the classic stars and bars problem⁹ [27, Chp. II, Sect. 5] with n stars and $\sigma - 1$ bars, yielding $\binom{n+s_{\sigma}-1}{n}$ possible strings. Hence we obtain:

Theorem 2.2. *There are $\binom{n+s_{\sigma}-1}{n}$ strings of length n over an alphabet with size σ having the suffix array $[n, n - 1, \dots, 1]$.*

As described above, strings of [Theorem 2.2](#) have the form $\sigma^{s_{\sigma}} (\sigma - 1)^{s_{\sigma-1}} \dots 1^{s_1}$. The BWT based on the suffix array $[n, n - 1, \dots, 1]$ is $1^{s_1-1} 2^{s_2} \dots \sigma^{s_{\sigma}} 1$. For $s_1 \geq 2$, it does not coincide with the BWT based on the rotations since the lexicographically smallest rotation is $1^{s_1} \sigma^{s_{\sigma}} \dots 2^{s_2}$, and hence the first entry of this BWT is 2. For $s_1 = 1$, the last character ‘1’ acts as the dollar sign being unique and least among all characters, making both BWT definitions equivalent.

For the rest of the analysis, we omit the arithmetically progressed permutation $[n, n - 1, \dots, 1]$ of ratio $k = n - 1$ as this case is complete. All other permutations (including those of ratio $k = n - 1$) are covered in our following theorems whose results we summarized in [Fig. 1](#).

3. Arithmetically progressed suffix arrays

We start with the claim that each arithmetically progressed permutation coincides with the suffix array of a string on a ternary alphabet. Subsequently, given an arithmetically progressed permutation P , we show that either there is precisely one string T with $SA_T = P$ whose characters are drawn from a *ternary* alphabet, or, if there are multiple candidate strings, then there is precisely one whose characters are drawn from a *binary* alphabet. For this aim, we start with the restriction on k and n to be coprime.

3.1. Coprimality

Two integers are *coprime*¹⁰ if their greatest common divisor (gcd) is one. An *ideal* $k\mathbb{N} := \{ki\}_{i \in \mathbb{N}}$ is a subgroup of $([1..n], +)$. It *generates* $[1..n]$ if $|k\mathbb{N}| = n$, i.e., $k\mathbb{N} = [1..n]$. Fixing one element $P[1] \in k\mathbb{N}$ of an ideal $k\mathbb{N}$ generating $[1..n]$ induces an arithmetically progressed permutation $P[1..n]$ with ratio k by setting $P[i+1] \leftarrow P[i] + k$ for every $i \in [1..n-1]$. On the contrary, each arithmetically progressed permutation with ratio k induces an ideal $k\mathbb{N}$ (the induced ideals are the

⁶ https://www.ncbi.nlm.nih.gov/nuccore/NC_045512.2?report=fasta.

⁷ We can also support negative values of k : Given a negative $k < 0$, we exchange it with $k' := n - k \bmod n \in [1..n]$ and use k' instead of k .

⁸ Note that we slightly misused notation as the exponentiations of the characters being integers have to be understood as writing a character as many times as the exponent. So 1^{s_1} does not give 1 but a string of length s_1 having only 1’s as characters.

⁹ This fundamental result in probability is useful in combinatorial counting problems involving variations of distributing q identical items into r distinguishable bins. This is equivalent to the number of permutations of q stars and $r - 1$ bars and given by the binomial coefficient $\binom{q+r-1}{q}$.

¹⁰ Also known as *relatively prime* in the literature.

p_1	k	Min. σ	Properties of Strings	Reference
1		2	unique, Lyndon word	Theorem 3.9
$k + 1$		2	unique, period $(n - k)$	Theorem 3.9
n	$\neq (n - 1)$	2	unique, period $(n - k)$	Theorem 3.9
	$= (n - 1)$	1	trivially periodic	Theorem 2.2
$\notin \{1, k + 1, n\}$		3	unique	Theorem 3.2

Fig. 1. Characterization of strings whose suffix array is an arithmetic progression $P = [p_1, \dots, p_n]$ of ratio k . The choice of p_1 determines the minimum size of the alphabet and whether a string is unique, periodic or a Lyndon word. The column *Min. σ* denotes the smallest possible size σ for which there exists such a string whose characters are drawn from an alphabet Σ with $|\Sigma| = \sigma$.

Rotation	T	P	$p_1 - k - 1$	BWT _T
	1 2 3 4 5 6 7 8	1 2 3 4 5 6 7 8		mod n
(1)	b a b b a b a c	[5, 2, 7 4, 1, 6, 3 8]	7	b^4ca^3
(2)	b a b a c b a c	[2, 7, 4 1, 6, 3 8, 5]	4	$b^3c^2a^3$
(3)	a c b a c b a c	[7, 4, 1 6, 3 8, 5, 2]	1	$b^2c^3a^3$
(4)	a c b a c a c c	[4, 1, 6 3 8, 5, 2, 7]	6	bc^4a^3
(5)	a b a b b a b b	[1, 6, 3 8, 5, 2, 7, 4]	3	b^5a^3
(6)	c c a c c a c b	[6, 3 8 5, 2, 7, 4, 1]	8	c^5a^2b
(7)	c c a c b c c b	[3 8, 5 2, 7, 4, 1, 6]	5	c^5ab^2
(8)	b a b b a b b a	[8, 5, 2 7, 4, 1, 6, 3]	2	b^5a^3

Fig. 2. T of Eq. (1) for each arithmetically progressed permutation P of length $n = 8$ with ratio $k = 5$, starting with $p_1 := P[1] = k = 5$. For $h \in [2..n]$, the permutation of the h th row is the $(h - 1)$ th cyclic rotation of the permutation P in the first row. The splitting of P into the subarrays is visualized by the vertical bar (|) symbol. For (5) and (8), the alphabet is binary and the BWTs are the same. The strings of (3) and (8) are periodic with period $n - k$, since the last text position of each subarray is at most as large as $n - k = 3$ (cf. the proof of Theorem 3.9). For $i \in [1..n]$, $BWT_T[i] = T[P[i + n - k^{-1} \bmod n]] = T[P[i + 3 \bmod n]]$ with $k^{-1} = k = 5$ defined in Section 3.4.

same for two arithmetically progressed permutations that are shifted). Consequently, there is no arithmetically progressed permutation with ratio k if k and n are not coprime since in this case $\{(ki \bmod n \mid i \geq 1)\} \subsetneq [1..n]$, from which we obtain:

Lemma 3.1. *The numbers k and n must be coprime if there exists an arithmetically progressed permutation of length n with ratio k .*

3.2. Ternary alphabet

Given an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio k , we define the ternary string $T[1..n]$ by splitting P right after the values $n - k$ and $(p_1 - k - 1) \bmod n$ into the three subarrays $A, B,$ and C (one of which is possibly empty) such that $P = ABC$. Subsequently, we set

$$T[p_i] := \begin{cases} a & \text{if } p_i \in A, \text{ or} \\ b & \text{if } p_i \in B, \text{ or} \\ c & \text{if } p_i \in C. \end{cases} \tag{1}$$




Fig. 3. Setting of the proof of Theorem 3.2 with the condition $p_i + 1 = p_1 - k = p_n$. In the figure we assume that the entry $p_1 - k - 1$ appears before $n - k$ in P .

BWT matrix of	BWT matrix of	BWT matrix of
babbabac:	ccaccacb:	bbabbabb:
abacbabb	acbccacc	abbabbbb
abbabacb	accacbcc	abbbbabb
acbabbab	bccaccac	babbabbb
babacbab	cacbccac	babbbbab
babbabac	caccacbc	bbabbabb
bacbabba	cbccacca	bbabbbba
bbabacba	ccacbcc	bbbabbab
cbabbaba	ccaccacb	bbbabba

Fig. 4. BWTs defined by the lexicographic sorting of all rotations of strings whose suffix arrays are cyclic rotations. This figure shows (from left to right) the BWT matrices of the strings of Rotation (1) and (6) of Fig. 2 as well as of Case (2) from Fig. 6. Reading the last column of a BWT matrix (whose characters are italic) from top downwards yields the BWT defined on the BWT matrix. While the BWT defined on the BWT matrix and the one defined by the suffix array coincides for the strings of Eq. (1) due to Theorem 3.3, this is not the case in general for the binary strings studied in Section 3.3, where we observe that $\text{BWT}_{\text{bbabbabb}} = \text{bbbbbbaab}$ defined by the suffix array differs from bbbbbaba (the last column on the right).

Fig. 2 gives examples for induced ternary/binary strings. Fig. 3 illustrates the case, where the value $(p_1 - k - 1) \bmod n$ is assumed to be left of $n - k$ in P .

Theorem 3.2. Given an arithmetically progressed permutation $P := [p_1, \dots, p_n] \neq [n, n - 1, \dots, 1]$ with ratio k , $\text{SA}_T = P$ for T defined in Eq. (1).

Proof. Suppose we have constructed SA_T . Since $a < b < c$, according to the above assignment of T , the suffixes starting with a lexicographically precede the suffixes starting with b, which lexicographically precede the suffixes starting with c. Hence, $\text{SA}[1..|A|]$, $\text{SA}[|A| + 1..|A| + |B|]$ and $\text{SA}[|A| + |B| + 1..n]$ store the same text positions as A , B , and C , respectively. Consequently, it remains to show that the entries of each subarray (A , B or C) are also sorted appropriately. Let p_i and p_{i+1} be two neighboring entries within the same subarray. Thus, $T[p_i] = T[p_{i+1}]$ holds, and the lexicographic order of their corresponding suffixes $T[p_i..n]$ and $T[p_{i+1}..n]$ is determined by comparing the subsequent positions, starting with $T[p_i + 1]$ and $T[p_{i+1} + 1]$. Since we have $(p_{i+1} + 1) - (p_i + 1) = p_{i+1} - p_i = k$, we can recursively show that these entries remain in the same order next to each other in the suffix array until either reaching the last array entry or a subarray split, that is, (1) $p_i + 1 = p_1 - k$ or (2) $p_i = n - k$.

- (1) When $p_i + 1$ becomes $p_1 - k \bmod n = p_n$ (the last entry in SA), p_{i+1} is in the subsequent subarray of the subarray of $p_i = p_1 - k - 1$ (remember that A or B ends directly after $p_1 - k - 1$, cf. Fig. 3). Hence $T[p_i] < T[p_{i+1}]$, and $T[p_i..n] < T[p_{i+1}..n]$.
- (2) The split at the value $n - k$ ensures that when reaching $p_i = n - k$ and $p_{i+1} = n$, we can stop the comparison here as there is no character following $T[p_{i+1}]$. The split here ensures that we can compare the suffixes $T[n - k..n]$ and $T[n]$ by the characters $T[n - k] < T[n]$. If we did not split here, $T[n] = T[n - k]$, and the suffix $T[n]$ would be a prefix of $T[n - k..n]$, resulting in $T[n] < T[n - k..n]$ (which yields a contradiction unless $p_1 = n$).

To sum up, the text positions stored in each of A , B and C are in the same order as in SA_T since the $j - 1$ subsequent text positions of each consecutive pair of entries p_i and p_{i+1} are consecutive in P for the smallest integer $j \in [1..n]$ such that $p_{i+1} + jk \in \{p_1 - 1, n\}$. □




Fig. 5. Setting of Eq. (1) with the distinction whether the entry $p_1 - k - 1$ appears before (a) or after (b) $n - k$ in P , yielding a different shape of the BWT_T defined as $BWT_{T[i]} = T[P'[i]]$ with $P'[i] = SA_T[i] - 1 \pmod n$.

Knowing the suffix array of the ternary string T of Eq. (1), we can give a characterization of its BWT. We start with the observation that both BWT definitions (rotation based and suffix array based) coincide for the strings of Eq. (1) (but do not in general as highlighted in the introduction, cf. Fig. 4), and then continue with insights in how the BWT looks like.

Theorem 3.3. *Given an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio $k \neq n - 1$ and the string T of Eq. (1), the BWT of T defined on the BWT matrix coincides with the BWT of T defined on the suffix array.*

Proof. According to Theorem 3.2, $SA_T = P$, and therefore the BWT of T defined on the suffix array is given by $BWT_T[i] = T[p_i - 1 \pmod n]$. The BWT matrix is constituted of the lexicographically ordered cyclic rotations of T . The BWT based on the BWT matrix, denoted by BWT_{matrix} , is obtained by reading the last column of the BWT matrix from top downwards (see Fig. 4). Formally, let $Q[i]$ be the starting position of the lexicographically i th smallest rotation $T[Q[i]..n]T[1..Q[i] - 1]$. Then $BWT_{matrix}[i]$ is $T[Q[i] - 1 \pmod n]$ if $Q[i] > 1$, or $T[n]$ if $Q[i] = 1$. We prove the equality $P = Q$ by showing that, for all $i \in [1..n - 1]$, the rotation $R_i := T[p_i..n]T[1..p_i - 1]$ starting at $p_i = SA_T[i]$ is lexicographically smaller than the rotation $R_{i+1} := T[p_{i+1}..n]T[1..p_{i+1} - 1]$ starting at $p_{i+1} = SA_T[i + 1]$. We do that by comparing both rotations R_i and R_{i+1} characterwise:

Let j be the first position where R_i and R_{i+1} differ, i.e., $R_i[j] \neq R_{i+1}[j]$ and $R_i[q] = R_{i+1}[q]$ for every $q \in [1..j]$.

First we show that $j \neq p_n - p_i + 1 \pmod n$ by a contradiction: Assuming that $j = p_n - p_i + 1 \pmod n$, we conclude that $j \neq 1$ by the definition of $i \in [1..n - 1]$. Since k is the ratio of P , we have

$$R_i[j - 1] = T[p_i + j - 2 \pmod n] = T[p_n - 1 \pmod n] = T[p_1 - k - 1 \pmod n]$$

and $R_{i+1}[j - 1] = T[p_{i+1} - 1 \pmod n]$. By Eq. (1), $p_1 - k - 1 \pmod n$ and $p_{i+1} - 1 \pmod n$ belong to different subarrays of P , therefore $T[p_1 - k - 1] \neq T[p_{i+1} - 1]$ and $R_i[j - 1] \neq R_{i+1}[j - 1]$, contradicting the choice of j as the first position where R_i and R_{i+1} differ.

This concludes that $j \leq n$ (hence, $R_i \neq R_{i+1}$) and $j \neq p_n - p_i + 1 \pmod n$. Hence, $R_i[j] = T[p_i + j - 1 \pmod n]$ and $R_{i+1}[j] = T[p_{i+1} + j - 1 \pmod n] = T[p_i + j - 1 + k \pmod n]$ are characters given by two consecutive entries in SA_T , i.e., $SA_T[q] = p_i + j - 1 \pmod n$ and $SA_T[q + 1] = p_i + j - 1 + k \pmod n$ for a $q \in [1..n - 1]$. Thus $R_i[j] \leq R_{i+1}[j]$, and by definition of j we have $R_i[j] < R_{i+1}[j]$, leading finally to $R_i < R_{i+1}$. Hence, $Q = P$. \square

Lemma 3.4. *Let $P := [p_1, \dots, p_n]$ be an arithmetically progressed permutation with ratio $k \neq n - 1$. Further, let $T[1..n]$ be given by Eq. (1) such that $SA_T = P$ according to Theorem 3.2. Given that $p_t = p_1 - k - 1 \pmod n$ for a $t \in [1..n]$, BWT_T is given by the t th rotation of $T[SA[1]] \cdots T[SA[n]]$, i.e., $BWT_T[i] = T[P[i + t \pmod n]]$ for $i \in [1..n]$.*

Proof. Since P is an arithmetically progressed permutation with ratio k then so is the sequence $P' := [p'_1, \dots, p'_n]$ with $p'_i = p_i - 1 \pmod n$. In particular, P' is a cyclic shift of P with $p'_n = p_1 - 1 - k \pmod n$ because $p'_1 = p_1 - 1$. However, $p'_n = p_n - 1 = p_1 - 1 - k$ is a split position of one of the subarrays A , B , or C , meaning that P' starts with one of these subarrays and ends with another of them (cf. Fig. 5). Consequently, there is a t such that $p_t = p'_n$, and we have the property that BWT_T with $BWT_T[i] = T[P'[i]]$ is the t th rotation of $T[SA[1]] \cdots T[SA[n]]$. \square

We will determine the parameter $t = n - k^{-1} \pmod n$ after Eq. (3) in Section 3.4, where k^{-1} is defined such that $k \cdot k^{-1} \pmod n = 1 \pmod n$. With Lemma 3.4, we obtain the following corollary which shows that the number of runs in BWT_T for T defined in Eq. (1) is minimal:

Corollary 3.5. For an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio $k \neq n - 1$ and the string T defined by Eq. (1), BWT_T consists of exactly 2 runs if T is binary, while it consists of exactly 3 runs if T is ternary.

Theorem 3.6. Given an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio k such that $p_1 \notin \{1, k + 1, n\}$, the string T given in Eq. (1) is unique.

Proof. The only possible way to define another string T' would be to change the borders of the subarrays A , B , and C . Since $p_1 \notin \{1, n\}$, $n - k$ and n , as well as $p_1 - k - 1$ and $p_1 - 1$, are stored as a consecutive pair of text positions in P .

- If P is not split between its consecutive text positions $n - k$ and n , then $T'[n - k] = T'[n]$. Consequently, we have the contradiction $T'[n] < T'[n - k..n]$.
- If P is not split between its consecutive text positions $(p_1 - k - 1) \bmod n$ and $(p_1 - 1) \bmod n$, then $T'[p_1 - k - 1 \bmod n] = T'[p_1 - 1 \bmod n]$. Since $p_1 \neq k + 1$, and $T'[p_1 - k \bmod n] = T'[p_n] > T'[p_1]$, this leads to the contradiction $T'[(p_1 - k - 1 \bmod n)..n] > T'[(p_1 - 1 \bmod n)..n]$, cf. Fig. 3. \square

Following this analysis of the ternary case we proceed to consider binary strings. A preliminary observation is given in Fig. 2, which shows, for the cases p_1 is 1 and n in Theorem 3.6, namely Rotations (5) and (8), that a rotation of $n - k$ in the permutation gives a rotation of one in the corresponding binary strings. We formalize this observation in the following lemma, drawing a connection between binary strings whose suffix arrays are arithmetically progressed and start with 1 or n .

Lemma 3.7. Let $P := [p_1, \dots, p_n]$ be an arithmetically progressed permutation with ratio k and $p_1 = 1$ for a binary string T over $\Sigma = \{a, b\}$ with $SA_T = P$. Suppose that the number of a 's in T is m and that $T' = T[2] \dots T[n]T[1]$ is the first rotation of T . Then $SA_{T'}$ is the m th rotation of P with $SA_{T'}[1] = n$. Furthermore, $BWT_T = BWT_{T'}$.

Proof. Since $p_1 = 1$, $T[1] = a$ and $T[n] = b$. In the following, we show that $P' = SA_{T'}$ for $P' := [p'_1, \dots, p'_n] := [p_1 - 1 \bmod n, \dots, p_n - 1 \bmod n]$ with $p'_1 = p_1 - 1 = n$ (since $T'[n] = a$). For that, we show that each pair of suffixes in SA_T is kept in the same relative order in P' (excluding $SA_T[1] = 1$):

Consider two text positions $p_i, p_j \in [p_2, \dots, p_n]$ with $T[p_i..n] = u_1 \dots u_s < T[p_j..n] = v_1 \dots v_t$.

- If $u_h \neq v_h$ for the least $h \in [1.. \min\{s, t\}]$, then $u_1 \dots u_s a < v_1 \dots v_t a$.
- Otherwise, $u_1 \dots u_s$ is a proper prefix of $v_1 \dots v_t = u_1 \dots u_s v_{s+1} \dots v_t$.
 - If $v_{s+1} = b$, then $u_1 \dots u_s a < u_1 \dots u_s b v_{s+2} \dots v_t a = v_1 \dots v_t a$.
 - Otherwise ($v_{s+1} = a$), $u_1 \dots u_s a$ is a proper prefix of $v_1 \dots v_t a$, and similarly $u_1 \dots u_s a < u_1 \dots u_s a v_{s+2} \dots v_t a = v_1 \dots v_t a$.

Hence the relative order of these suffixes given by $[p_2, \dots, p_n]$ and $[p'_2, \dots, p'_n]$ is the same. In total, we have $p'_i = p_i - 1 \bmod n$ for $i \in [1..n]$, hence P' is an arithmetically progressed permutation with ratio k . Given the first m entries in P represent all suffixes of T starting with a , P' is the m th rotation of P since $p'_1 = n$ is the $(m + 1)$ th entry of P , i.e., the smallest suffix starting with b in T . Finally, since the strings T and T' are rotations of each other, their BWTs are the same. \square

Like the parameter t of Lemma 3.4, we will determine the parameter m after Eq. (3) in Section 3.4.

3.3. Binary alphabet

We start with the construction of a binary string from an arithmetically progressed permutation:

Theorem 3.8. Given an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio $k \neq n - 1$ such that $p_1 \in \{1, k + 1, n\}$, we can modify T of Eq. (1) to be a string over the binary alphabet $\{a, b\}$ with $SA_T = P$.

Proof. If $p_1 = 1$, then P is split after the occurrences of the values $n - k$ and $-k = n - k \bmod n$, which gives only two non-empty subarrays. If $p_1 = n$, P is split after the occurrence of $n - k - 1$, which implies that C is empty since $p_n = n - k$. Hence, T can be constructed with a binary alphabet in those cases, cf. Fig. 2.

For the case $p_1 = k + 1$, P is split after the occurrences of the values $n - k$ and $k + 1 - k - 1 \bmod n = n \bmod n$, so B contains only the text position n . By construction, the requirement is that the suffix $T[n]$ is smaller than all other suffixes starting with c . So instead of assigning the unique character $T[n] \leftarrow b$ like in Theorem 3.2, we can assign $T[n] \leftarrow c$, which still makes $T[n]$ the smallest suffix starting with c . We conclude this case by converting the binary alphabet $\{a, c\}$ to $\{a, b\}$. Cf. Fig. 2, where T in Rotation (6) has become $bbabbabb$ with period $n - k = 3$. \square

The main result of this section is the following theorem. There, we characterize all binary strings whose suffix arrays are arithmetically progressed permutations. More precisely, we identify which of them are unique,¹¹ periodic, or a Lyndon word.

¹¹ The exact number of these binary strings is not covered by Theorem 3.6.

Case	T								SA								p_s	s
	1	2	3	4	5	6	7	8	1	2	3	4	5	6	7	8		
(1)	b	a	b	b	a	b	b	a	[8, 5, 2 7, 4, 1, 6, 3]	2	3							
(2)	b	b	a	b	b	a	b	b	[6, 3 8, 5, 2, 7, 4, 1]	3	2							
(3)	a	b	a	b	b	a	b	b	[1, 6, 3 8, 5, 2, 7, 4]	3	3							

Fig. 6. All binary strings of length 8 whose suffix arrays are arithmetically progressed permutations with ratio $k = 5$. Theorem 3.9 characterizes these strings (and also gives the definition of p_s). Cases (1) and (3) also appear in Fig. 2 at Rotation (8) and (5), respectively, while Case (2) can be obtained from Rotation (6) by exchanging the last character with c. Cases (1) and (2) both have period $n - k = 3$, and Case (3) is a Lyndon word.




Fig. 7. Sketches of the cases of Theorem 3.9. T_s is uniquely determined if the suffix array SA of T_s is arithmetically progressed with ratio k and the first entry $SA[1] \in \{1, k + 1, n - k\}$ is given.

Theorem 3.9. Let n and $k \in [1..n - 1]$ be two coprime integers. If $k \neq n - 1$, there are exactly three binary strings of length n whose suffix arrays are arithmetically progressed permutations with ratio k . Each such solution $T_s \in \{a, b\}^+$ is characterized by

$$T_s[i] = \begin{cases} a & \text{for } i \in SA_{T_s}[1..s], \text{ or} \\ b & \text{otherwise,} \end{cases} \tag{2}$$

for all text positions $i \in [1..n]$ and an index $s \in [1..n - 1]$, called the split index.

The individual solutions are obtained by fixing the values for p_1 and p_s , the position of the lexicographically largest suffix starting with a, of $SA_{T_s} = [p_1, \dots, p_n]$:

1. $p_1 = n$ and $p_s = n - k - 1$,
2. $p_1 = k + 1$ and $p_s = n - k$, and
3. $p_1 = 1$ and $p_s = n - k$.

The string T_s has period $n - k$ in Cases 1 and 2, while T_s of Case 3 is a Lyndon word, which is not periodic by definition.

Proof. Let S be a binary string of length n , and suppose that $SA_S = P := [p_1, \dots, p_n]$ is an arithmetically progressed permutation with ratio k . Further let p_s be the position of the largest suffix of S starting with a. Then $S[p_i..n] < S[p_{i+1}..n]$ and thus $S[p_i] \leq S[p_{i+1}]$. We have $S[j] = S[j + k \bmod n]$ for all $j \in [1..n] \setminus \{p_n, p_s\}$ since

- $p_n = SA_S[n]$ is the starting position of the largest suffix ($S[p_n] = b \neq a = S[p_1] = S[p_n + k]$).
- $S[p_s..n]$ and $S[p_{s+1}..n]$ are the lexicographically largest suffix starting with a and the lexicographically smallest suffix starting with b, respectively, such that $S[p_s] = a \neq b = S[p_{s+1}]$.

To sum up, since $S[p_s..n] < S[p_{s+1}..n]$ by construction, $S[p_i..n] < S[p_{i+1}..n]$ holds for $p_i > p_{i+1}$ whenever $p_i \neq p_n$. This, together with the coprimality of n and k , determines p_s uniquely in the three cases (cf. Fig. 6 for the case that $n = 8$ and $k = 5$, and Fig. 7 for sketches of the proof):

Case $p_1 = n$: We first observe that the case $k = n - 1$ gives us $P = [n, n - 1, \dots, 1]$, and this case was already treated with Theorem 2.1. In the following, we assume $k < n - 1$, and under this assumption we have $s > 1$,

$T_s[n] = T_s[p_1] = a$ and $T_s[n - 1] = b$ (otherwise $T_s[n - 1..n]$ would be the second smallest suffix, i.e., $P[2] = n - 1$ and hence $k = n - 1$). Consequently, $T_s[n - 1..n] = T_s[p_{s+1}..n]$ is the smallest suffix starting with b , namely ba , and therefore $p_s = n - 1 - k$.

Case $p_1 \neq n$: If $p_1 \neq n$, then $T_s[n] = b$ (otherwise $T_s[n] < T_s[p_1..n]$). Therefore, $T_s[n]$ is the smallest suffix starting with b , and consequently $p_s = n - k$.

For the periodicity, with $T_s[j] = T_s[j + k \bmod n] = T_s[j - (n - k) \bmod n]$ for $j \in [1..n] \setminus \{p_1, p_s\}$ we need to check two conditions:

- If $p_n - (n - k) > 0$, then $T_s[p_n - (n - k)] = T_s[p_1] \neq T_s[p_n]$ breaks the periodicity.
- If $p_s - (n - k) > 0$, then $T_s[p_s - (n - k)] = a \neq b = T_s[p_{s+1}]$ breaks the periodicity.

For Case 1, $p_n = n - k$ and $p_s = n - k - 1$ (hence $p_n - (n - k) = 0$ and $p_s - (n - k) = -1$), thus Case 1 is periodic.

Case 2 is analogous to Case 1.

For Case 3, T_s does not have period $n - k$ as $p_n = n - k + 1$, and hence $p_n - (n - k) > 0$. It cannot have any other period since Case 3 yields a Lyndon word (because the lexicographically smallest suffix $T_s[p_1..n] = T_s[1..n]$ starts at the first text position). Note that Case 3 can be obtained from Case 2 by setting $T_s[1] \leftarrow a$ (the smallest suffix $T_s[k + 1..n]$ thus becomes the second smallest suffix).

Finally, we need to show that no other value for p_1 admits a binary string S having an arithmetically progressed permutation $P := [p_1, \dots, p_n]$ with ratio k as its suffix array. So suppose that $p_1 \notin \{1, k + 1, n\}$, then this would imply the following:

- $S[p_1] = a$ because the smallest suffix starts at text position p_1 , and
- $S[p_1 - 1] = b$ because of the following: First, the text position $S[p_1 - 1]$ exists due to $p_1 > 1$. Second, since $p_1 < n$, there is a text position $j \in [p_1 + 1..n]$ such that $S[p_1] = \dots = S[j - 1] = a$ and $S[j] = b$ (otherwise $S[n]$ would be the smallest suffix). If $S[p_1 - 1] = a$, then the suffix $S[p_1 - 1..n]$ starting with $a^{j-p_1+1}b$ is lexicographically smaller than the suffix $S[p_1..n]$ starting with $a^{j-p_1}b$. Hence, $S[p_1 - 1] = b$ must hold.
- $p_n - 1 \geq 1$ (since $p_1 \neq k + 1$) and $S[p_n - 1] = a$. If $S[p_n - 1] = b$, then the suffix $S[p_n - 1..n]$ has a longer prefix of b 's than the suffix $S[p_n..n]$, and is therefore lexicographically larger.

Since $S[p_n - 1] = a$ and $S[p_1 - 1] = b$ with $p_n - 1 + k \bmod n = p_1 - 1$, the smallest suffix starting with b is located at index $p_1 - 1$. This is a contradiction as $p_1 \neq n$ implies $S[n] = b$ (if $S[n] = a$, then $SA[1] = n$ instead of $SA[1] = p_1$) and thus the smallest suffix starting with b is located at index n (this is a contradiction since we assumed that this suffix starts at $p_1 - 1 \in [1..n - 1]$). This establishes the claim for p_1 . \square

For a given arithmetically progressed permutation with ratio k , and first entry $p_1 \in \{1, k + 1, n\}$, the string T_s of [Theorem 3.9](#) coincides with T of [Theorem 3.8](#).

3.4. Inverse permutations

Since the inverse P^{-1} of a permutation P with $P^{-1}[P[i]] = i$ is also a permutation, one may wonder whether the inverse P^{-1} of an arithmetically progressed permutation is also arithmetically progressed. We affirm this question in the following. For that, we use the notion of the *multiplicative inverse* k^{-1} of an integer k (to the congruence class $[1..n] = \mathbb{Z}/n\mathbb{Z}$), which is given by $k^{-1} \cdot k \bmod n = 1 \bmod n$. The multiplicative inverse k^{-1} is uniquely defined if k and n are coprime.

In what follows, we show that the inverse of an arithmetically progressed permutation P with ratio k is also arithmetically progressed with ratio k^{-1} . As an example consider [Fig. 2](#), where $n = 8$ and $k = 5$. The multiplicative inverse of k is given by $k^{-1} = 5 = k$. Hence, the set of arithmetically progressed permutations with the same parameters $n = 8$ and $k = 5$ contains also all its inverse permutations. Speaking of [Fig. 2](#), the suffix arrays of (1) and (5) are each self-inverse. The suffix arrays of (2) and (4), (3) and (7), as well as (6) and (8) are inverse pairs, meaning that one is the inverse of the other.

Theorem 3.10. *The inverse P^{-1} of an arithmetically progressed permutation P with ratio k is an arithmetically progressed permutation with ratio k^{-1} and $P^{-1}[1] = (1 - P[n]) \cdot k^{-1} \bmod n$.*

Proof. Let $x := P[i]$ for an index $i \in [1..n]$. Then $P[i + k^{-1} \bmod n] = P[i] + k \cdot k^{-1} \bmod n = x + k \cdot k^{-1} \bmod n = x + 1 \bmod n$. For the inverse permutation P^{-1} this means that $P^{-1}[x] = i$ and $P^{-1}[x + 1 \bmod n] = i + k^{-1} \bmod n$. Thus the difference between two consecutive entries of P^{-1} , i.e., $P^{-1}[x + 1 \bmod n] - P^{-1}[x]$, is k^{-1} .

Since $P[i] = j \iff P[n] + ik \bmod n = j$ holds for all indices $i \in [1..n]$, we have (using $i \leftarrow P^{-1}[1]$ and $j \leftarrow 1$ in the above equivalence)

$$\begin{aligned} P[P^{-1}[1]] &= 1 \bmod n \iff P[n] + P^{-1}[1] \cdot k = 1 \bmod n \\ &\iff P^{-1}[1] \cdot k = 1 - P[n] \bmod n \\ &\iff P^{-1}[1] = (1 - P[n]) \cdot k^{-1} \bmod n. \quad \square \end{aligned}$$

Consequently, using the split index s of p_s for SA and

$$\begin{aligned} \text{ISA}[i] &= \text{ISA}[1] + (i - 1)k^{-1} \pmod n \\ &= (1 - \text{SA}[n]) \cdot k^{-1} + (i - 1)k^{-1} \pmod n \\ &= (i - \text{SA}[n]) \cdot k^{-1} \pmod n, \end{aligned}$$

we can rewrite T_s defined in Eq. (2) as

$$T_s[i] = \begin{cases} a & \text{if } \text{ISA}[i] \leq s, \text{ or} \\ b & \text{otherwise} \end{cases} \tag{3}$$

where SA and ISA denote the suffix array and the inverse suffix array of T_s , respectively. Another result is that $\text{ISA}[p_s] = s$ is the number of a's in T_s , for which we split the study into the cases of Theorem 3.9:

1. If $\text{SA}[1] = n$ and $p_s = n - k - 1$, then $\text{SA}[n] = n - k$ and $\text{ISA}[i] = (i - n + k)k^{-1} \pmod n$. Consequently, $s = \text{ISA}[p_s] = (-1)k^{-1} \pmod n = n - k^{-1} \pmod n$.
2. If $\text{SA}[1] = k + 1$ and $p_s = n - k$, then $\text{SA}[n] = 1$ and $\text{ISA}[i] = (i - 1)k^{-1} \pmod n$. Consequently, $s = \text{ISA}[p_s] = (n - k - 1)k^{-1} \pmod n = nk^{-1} - 1 - k^{-1} \pmod n = n - 1 - k^{-1} \pmod n$.
3. If $\text{SA}[1] = 1$ and $p_s = n - k$, then $\text{SA}[n] = n - k + 1$ and $\text{ISA}[i] = (i - n + k - 1)k^{-1} \pmod n$. Consequently, $s = \text{ISA}[p_s] = (-1)k^{-1} \pmod n = n - k^{-1} \pmod n$ as in Case (1).

For Fig. 6 with $k = 5$ and $n = 8$, we know that the number of a's is $\text{ISA}[p_s] = 3$ in Cases (1) and (3), and $\text{ISA}[p_s] = 2$ in Case (2) because $k^{-1} = 5 \Leftrightarrow k \cdot k^{-1} \pmod n = 1 \pmod n$. This also determines the constant m used in Lemma 3.7. Finally, we can fix the parameter t in Lemma 3.4 defined such that $p_t = p_1 - 1 - k \pmod n$: For that, write $\text{ISA}[i] = (i - p_n)k^{-1} \pmod n = (i + k - p_1)k^{-1} \pmod n$ and compute $\text{ISA}[p_t] = \text{ISA}[p_1 - 1 - k] = (-1)k^{-1} \pmod n = n - k^{-1} \pmod n$.

4. Applications

We conclude our main results of Section 3 by drawing connections between strings having arithmetically progressed suffix arrays and Christoffel words (Section 4.1), balanced words (Section 4.2), and Fibonacci words (Section 4.3).

4.1. Christoffel words

Christoffel words are binary strings whose origins are considered to date from Bernoulli's 1771 work [28]. Christoffel words can be described geometrically in terms of a line segment and associated polygon traversal [29]: let $(p, q) \in \mathbb{N}^2$ where (p, q) are coprime and let S be the line segment with endpoints $(0, 0)$ and (p, q) . The induced path of a binary string T is a list of points $v_0, \dots, v_n \in \mathbb{N}^2$ such that $v_0 = (0, 0)$, $v_n = (p, q)$, and for each $i \in [1..n]$, $v_i - v_{i-1} = (1, 0)$ if $T[i] = a$ and $v_i - v_{i-1} = (0, 1)$ if $T[i] = b$. The string $T \in \{a, b\}^*$ is a lower Christoffel word if the path induced by T from the origin to (p, q) is below the line segment S and the path and S determines a simple polygon¹² which contains no other point in \mathbb{N}^2 . An upper Christoffel word is defined analogously by taking the path above S . Hence, a Christoffel word is defined by a direction (above or below) and the slope $\frac{q}{p}$, which determines p and q since p and q are coprime. We illustrate the upper and lower Christoffel words for the slope $\frac{4}{7}$ in Fig. 8. Furthermore, Case (3) of Fig. 6 determines a Christoffel word with slope $\frac{5}{3}$ and Fig. 11 gives an example of the geometric description where $p = 7$ and $q = 5$.

It follows from the coprimality of p and q that Christoffel words are necessarily primitive. In what follows, we focus on lower Christoffel words, and may drop the adjective lower when it is clear we are speaking of (lower) Christoffel words.

To show that every lower Christoffel word has an arithmetically progressed suffix array, we use an alternative characterization of Christoffel words based on Cayley graphs [30, Def. 1.4]. The Cayley graph (V, E) of a group G with generating set Σ is defined by $V = G$ and $E = \{(g, ga) \mid g \in G, a \in \Sigma\}$. Let again $p, q \in \mathbb{N}$ be coprime. Fig. 9 is the Cayley graph of $\mathbb{Z}/(p + q)\mathbb{Z}$ generated by q . Cayley graphs are always simple cycles since q and $p + q$ are coprime. In what follows, we use mod_0 with $n \text{ mod}_0 n = 0 \text{ mod}_0 n = 0$ to match the definition in [30], as opposed to mod with $0 \text{ mod } n = n \text{ mod } n = n$ elsewhere.

An edge $s \rightarrow t$ in the Cayley graph has the label a if $s < t$, otherwise it has the label b . Reading the edge labels, starting from node 0 , following the edges of the Cayley graph and stopping when reaching node 0 again, yields the lower Christoffel word T_c parameterized by p and q . The i th node in the Cayley graph (0 being the first node) is $(i - 1)q \text{ mod}_0 (p + q)$. Hence the i th character of T_c is

$$T_c[i] = \begin{cases} a & \text{if } (i - 1)q \text{ mod}_0 (p + q) < iq \text{ mod}_0 (p + q), \\ b & \text{if } (i - 1)q \text{ mod}_0 (p + q) > iq \text{ mod}_0 (p + q). \end{cases} \tag{4}$$

Given the suffix array SA_{T_c} of a lower Christoffel word T_c , the split index s (defined in Theorem 3.9) is given by p , the total number of units along the x -axis in the polygonal path. All lower Christoffel words are Lyndon words [29] and so necessarily border-free. Hence $\text{SA}_{T_c}[p_1] = 1$, and $\text{SA}_{T_c}[s + 1] = n$ since the string T_c must end b . The following theorem now gives the connection between lower Christoffel words and the strings of Section 3:

¹² Note that this polygon is unique and the line segment S contains no point in \mathbb{N}^2 as p and q are coprime.




Fig. 8. An upper Christoffel word $T_U = babaabaabaa$ and its reversal, the lower Christoffel word $T_L = aabaabaabab$, shown geometrically with respect to the line segment $S = ((0, 0), (7, 4))$.




Fig. 9. Cayley graph of $\mathbb{Z}/(p+q)\mathbb{Z}$ generated by q . See [30, Fig. 1.4] for a concrete example.

Theorem 4.1. Let $p, q \in \mathbb{N}$ be coprime. Then the lower Christoffel word T_c characterized by p and q has an arithmetically progressed suffix array. The suffix array is given by the arithmetic progression P with $p_1 = 1$ and $k = q^{-1}$, where q^{-1} is the multiplicative inverse of q in $\mathbb{Z}/n\mathbb{Z}$. The string T_c is identical to Case 3 of Theorem 3.9 characterizing the binary case.

Proof. We prove the theorem by showing that the Christoffel word T_c is equal to the string described in Theorem 3.9 as Case 3 whose suffix array is the arithmetically progressed array P .

Let $n = p + q$. By Eq. (4) the i th character of T_c is an a if and only if $(i - 1)q \bmod_0 n < iq \bmod_0 n$. We can rewrite that to $(i - 1)q \bmod_0 n < (i - 1)q \bmod_0 n + q \bmod_0 n$. This condition is fulfilled if and only if $((i - 1)q \bmod_0 n) + q < n$. We can rewrite that to $(i - 1)q \bmod_0 n < n - q$. Replacing \bmod_0 with \bmod , we obtain

$$T_c[i] = \begin{cases} a & \text{if } 1 + (i - 1)q \bmod n < n + 1 - q, \\ b & \text{otherwise.} \end{cases}$$

Let $k = q^{-1}$. Let T_s denote the string of Case 3 in Theorem 3.9, which is characterized by $p_1 = 1$ and $p_s = n - k = n - q^{-1}$. Using the results from Section 3.4, Eq. (3), T_s can be written as

$$T_s[i] = \begin{cases} a & \text{if } (i + k - 1)k^{-1} \bmod n \leq s, \\ b & \text{otherwise.} \end{cases} \tag{5}$$

Substituting $k = q^{-1}$, $k^{-1} = q$ and $s = \text{ISA}[p_s] = \text{ISA}[n - q^{-1}] = n - q$, we can rewrite the condition for a in Eq. (5) as $1 + (i - 1)q \bmod n \leq n - q$. Replacing \leq with $<$ we obtain the same definition for T_s as for T_c , concluding the proof.

$$T_s[i] = \begin{cases} a & \text{if } 1 + (i - 1)q \bmod n < n + 1 - q, \\ b & \text{otherwise.} \end{cases} \quad \square$$

In the geometric setting, traversing the path of the polygon from $(0, 0)$ to (p, q) is equivalent to scanning the characters in the defining Christoffel word T_c ; hence, we can associate each character with its polygon coordinates. The BWT can be obtained from iterating across the suffix array and for each index selecting the preceding character in the text – the BWT for the first string in Fig. 10 has maximal runs in the form b^5a^7 ; more generally, the BWT of a Christoffel word over $\Sigma = \{a, b\}$ with slope $\frac{q}{p}$ takes the form $b^q a^p$ [30].

The progression ratio q^{-1} of Theorem 4.1 may be useful for accessing geometries of interest in the polygon or aiding discovery of geometric repetitive structures. For instance, to access the ‘steps’ in the polygon in decreasing width, start at

	T	SA	p_s	s
	1 2 3 4 5 6 7 8 9 10 11 12	1 2 3 4 5 6 7 8 9 10 11 12		
(1)	a a b a b a a b a b a b	[1, 6, 11, 4, 9, 2, 7, 12, 5, 10, 3, 8]	7	7
(2)	b a b a b a a b a b a a	[1, 8, 3, 10, 5, 12, 7, 2, 9, 4, 11, 6]	5	5

Fig. 10. Two Christoffel words with $S = ((0, 0), (7, 5))$, slope $\frac{5}{7}$, and length 12. The suffix array of the lower Christoffel word (1) has an arithmetically progressed permutation with ratio $k = 5$, $p_s = 7 = n - k$, and hence is an instance of Theorem 3.9, Case 3. Then, assuming $b < a$, the suffix array of the upper Christoffel word (2) has an arithmetic progression ratio of $n - k = 7$ with split index $n - s$ and position of largest suffix starting b of $n - p_s$.




Fig. 11. An upper Christoffel word $T_U = bababaababaa$ and its reversal, the lower Christoffel word $T_L = aababaababab$, shown geometrically with respect to the line segment $S = ((0, 0), (7, 5))$. The dotted lines indicate the balanced₂ factorization of T_L into the factors $aabab$ and $aabab$.

the origin and apply increments of k : for the first string in Fig. 10 the widest steps are at coordinates $(0, 0)$ and $(3, 2)$ – see Fig. 11. Note that the geometric width of this polygon at a certain height (that is, a step) reflects a run of the character a in its representing Christoffel word T_c .

We now extend a known result for lower Christoffel words [31, Corollary 5.1], [30, Sect. 6.1], which distinguishes consecutive positions in consecutive rows in an associated BWT matrix:

Lemma 4.2. *Let T_c be a Christoffel word of length $n \geq 2$ over $\Sigma = \{a, b\}$. Suppose T_c has an arithmetically progressed suffix array with progression ratio k . Then in the BWT matrix \mathbf{M} of T_c with $\mathbf{M}[i, j] = T_c[(SA_{T_c}[i] + j - 1) \bmod n]$, rows i and $i + 1$ for $1 \leq i < n$ differ in exactly two consecutive positions. For each row i , these positions are at $\mathbf{M}[i, i(n - k) \bmod n]$ and $\mathbf{M}[i, i(n - k) + 1 \bmod n]$.*

Proof. The Christoffel word T_c is characterized by Theorem 3.9 (3). Further, since T_c is also a Lyndon word then it starts with a and ends with b . The second row in the BWT matrix \mathbf{M} of T_c is $T_c[SA_{T_c}[2]..n]T_c[1..SA_{T_c}[2] - 1]$ where the prefix $T_c[SA_{T_c}[2]..n]$ has length $n - k$ and ends with b .

From Case 3 of Theorem 3.9, $p_s = n - k$, and so $T_c[n - k] = a$, and since $T_c[n - k..n]$ is the largest suffix starting with a and $n \geq 2$, then it has the prefix ab . So $\mathbf{M}[1, n - k] = a$ and $\mathbf{M}[1, n - k + 1] = b$, while from $T_c[1] = a$ and $T_c[n] = b$, $\mathbf{M}[2, n - k] = b$ and $\mathbf{M}[2, n - k + 1] = a$.

This gives an arrangement, denoted by \mathcal{A} , across two adjacent rows i and $i + 1$ and depicted as $\mathcal{A} = \frac{ab}{ba}$, where at some index h , $\mathbf{M}(i, h) = a$, $\mathbf{M}(i, h + 1) = b$, $\mathbf{M}(i + 1, h) = b$, $\mathbf{M}(i + 1, h + 1) = a$. We will map the locations of \mathcal{A} in the matrix \mathbf{M} which for the first two rows start at $\mathbf{M}(1, n - k)$ and $\mathbf{M}(2, 2(n - k) \bmod n)$. Similarly for each subsequent row i in \mathbf{M} , $i < n$, there is an additional factor of $(n - k) \bmod n$ for the position, where \mathcal{A} occurs.

We proceed to show that there are no other differences between adjacent rows other than those occurring in \mathcal{A} . So suppose instead that $\mathbf{M}[i, j] = c$ and $\mathbf{M}[i + 1, j] = d$ for two distinct characters $c, d \in \Sigma$, $j \neq i(n - k) \bmod n$ and $j \neq i(n - k) + 1 \bmod n$. Then the character d in row $i + 1$ is the same character in T_c as the one at position $j - (n - k) \bmod n$ in row i . Consider the character c at position $j - (n - k) - k \bmod n$ in row i . If c is a b then d cannot be an a . If c is an a and d is a b then this must occur at \mathcal{A} , that is $i = SA^{-1}[p_s]$ and $j = 1$, thereby contradicting the conditions on j . So if c is an a then d cannot be a b .

Hence the only differences between adjacent rows in \mathbf{M} occur at \mathcal{A} which coincides with $T_c[p_s]$ and $T_c[n]$ and their adjacent letters in the rows. In particular, \mathcal{A} determines the lexicographic ordering of the rows of the matrix \mathbf{M} . \square

The arithmetic progression of arithmetically progressed suffix arrays of Christoffel words allows us to determine the following factorizations of such words in constant time (without looking at the explicit characters of the word):

- **Right factorization** [30]; originally called the **standard factorization** [32,33]. If $w = uv$ is a Lyndon word with v its lexicographically least proper suffix, then u and v are also Lyndon words and $u < v$. Equivalently, the right factor of the standard factorization can be defined to be the longest proper suffix of w that is a Lyndon word [34].
- **Left factorization** [35,36]. If $w = uv$ is a Lyndon word with u a proper Lyndon prefix of maximal length, then v is also a Lyndon word and $u < v$.
- **balanced₂ factorization** [37]. A Lyndon word w is balanced₂ if w is a character or there is a factorization $w = (u, v)$ that is simultaneously the left and the right factorization of w , and u and v are also balanced₂.

For any of the three above factorizations $w = (u, v)$, we will say that the **factorization index** is the index of the last character of u in w . The factorization index determines the split position and hence the factorization.

Lemma 4.3. *Let T_c be a Christoffel word of length $n \geq 3$ over $\Sigma = \{a, b\}$. Suppose T_c has an arithmetically progressed suffix array SA_{T_c} with progression ratio $k \geq 1$ and split index s . Then T_c is a balanced₂ Lyndon word with a factorization $T_c = T_u T_v$ such that $|T_u| = SA_{T_c}[(s + 2) \bmod n]$. For $n = 1, 2$ we have that $T_c \in \{a, b, ab\}$, $k = 1$, and that the split index s is 1.*

Proof. We apply the result that a word is a lower Christoffel word if and only if it is a balanced₂ Lyndon word [37] or [30, Theorem 6.7]. So given T_c , it remains to prove that $|T_u| = SA_{T_c}[(s+2) \bmod n]$. Let $j \in [1..n]$ be such that $|T_u| = SA_{T_c}[j]$. Then $T_u = T_c[1..SA_{T_c}[j]]$ and $T_v = T_c[SA_{T_c}[j] + 1..n]$. Since T_u and T_v are Lyndon words, we have that

- $j > s$, otherwise $T_c[SA_{T_c}[j]] = a$ and T_u would have a border.
- $SA_{T_c}[(s + 1) \bmod n] = n$ with $T_c[n] = b$ since T_c is a Lyndon word larger than one, and the smallest suffix starting with b is $T_c[n..n]$.
- The second smallest suffix starting with b is $T_c[SA_{T_c}[2] - 1..n]$. To see that observe that, because T_c is a Lyndon word, $SA_{T_c}[1] = 1$, and hence $SA_{T_c}[2] \geq 2$. Next, $T_c[SA_{T_c}[2] - 1] = b$, otherwise $T_c[SA_{T_c}[2] - 1..n] < T_c[SA_{T_c}[2]..n]$ would give a smaller suffix. Hence, $SA_{T_c}[(s + 2) \bmod n] = SA_{T_c}[2] - 1$.

Finally, $T_v = T_c[SA_{T_c}[2]..n]$ since otherwise we yield a contradiction to the right factorization that T_v is the lexicographically least proper suffix. So $j = SA_{T_c}[(s + 2) \bmod n]$. \square

For example, the Christoffel word in Fig. 10 Case (1) has factorization index $SA_{T_c}[9] = 5$ with the balanced₂ factorization $(aabab)(aababab)$; the first level of recursion gives the balanced₂ factorizations $(aab)(ab)$ and $(aabab)(ab)$.

Conditions for the factorization of a Lyndon word into exactly two non-overlapping Lyndon factors are given in [38], where *overlapping factors* have non-empty suffixes and prefixes sharing the same characters. We can geometrically consider the balanced₂ factorization of Christoffel words as follows:

Lemma 4.4. *Let $p, q \in \mathbb{N}$ be coprime and T_c be the lower Christoffel word characterized by p and q with factorization index i . Further, let the line segment $S = ((0, 0), (p, q))$ be associated with T_c . Then for all points on the path determined by T_c (apart from the end points), $T_c[i]$ has the shortest Euclidean distance to S .*

Proof. Since T_c is a Christoffel word it is also a balanced₂ Lyndon word. Let the balanced₂ factorization be $T_c = T_u T_v$, where the factors have defining line segments S_u, S_v respectively, i.e., S_u is the line from $(0, 0)$ to the point s being the geometric representation of $T_c[1..|T_u|]$, and S_v is the line from s to (p, q) , cf. Fig. 11 with $S_u = aabab$ and $s = (3, 2)$. Now assume that there exists a point r that geometrically represents a prefix $T_c[1..i]$ with the property that r does not have the shortest Euclidean distance to S . Then at least one of the paths for T_u and T_v must cross their associated line S_u or S_v , contradicting the geometric definition of a Christoffel word. \square

We can also interpret the increasing sequences in the suffix array of a Christoffel word geometrically:

Lemma 4.5. *Let $p, q \in \mathbb{N}$ be coprime and T_c be the lower Christoffel word over $\Sigma = \{a, b\}$ characterized by p and q and with the associated line segment $S = ((0, 0), (p, q))$. Suppose T_c of length $n \geq 2$ has the arithmetically progressed suffix array SA_{T_c} with progression ratio $k \geq 1$. Let $SA_{T_c} = A_1, A_2, \dots, A_j$, $1 \leq j < n$, be the factorization of SA_{T_c} into j maximally increasing sequences. Then,*

- (i) for $A_i = a_1 < a_2 < \dots < a_f$, $1 \leq i \leq j$, the geometry given by the coordinates of the characters indexed in T_c by A_i is a line (possibly degenerate with $f = 1$), and
- (ii) the j lines in (i) are parallel, and
- (iii) the line for A_i has shorter Euclidean distance to S than that for A_{i+1} , $1 \leq i < j$.

(For $n = 1$, $T_c = a$ and there is one degenerate line, a point. For $n = 2$, $T_c = ab$ and the unique line is $((0, 0), (1, 0))$.)




Fig. 12. The factorization of the suffix array [1, 6, 11, 4, 9, 2, 7, 12, 5, 10, 3, 8] of the lower Christoffel word $T_L = \text{aababaababab}$ in Figs. 10 & 11 shown geometrically. The parallel lines associated with the increasing sequences in the suffix array, namely [1, 6, 11], [4, 9], [2, 7, 12], [5, 10], [3, 8], are depicted with decreasing thickness as they get further from the line S .

Proof. Part (i). Consider the sequence $A_i = a_1, a_2, \dots, a_f$ with $f > 1$. For $1 < d < f$, if $T_c[a_1] = a$ then $T_c[a_d] = a$, otherwise $T_c[a_1] = T_c[a_d] = b$. Suppose there is a minimal $p > 0$ such that $T_c[a_1 + p] \neq T_c[a_d + p]$. If $T_c[a_1 + p] = a$ and $T_c[a_d + p] = b$ then the geometric path of $T_c[a_d..a_d + p]$ crosses S ; otherwise, $T_c[a_1 + p] = b$ and $T_c[a_d + p] = a$ and the region between $T_c[a_d..a_d + p]$ and S contains a point in \mathbb{N}^2 – either case contradicts the definition of a Christoffel word. Therefore the paths $T_c[a_1..a_1 + k - 1]$ and $T_c[a_d..a_d + k - 1]$ have the same shape and the same slope. For $d = f$, if $a_d \neq n$ then $T_c[a_d..n - 1]$ is a proper prefix of $T_c[a_1..a_1 + k - 1]$ hence with $T_c[a_d]$ on the slope, while if $a_d = n$ then $T_c[a_d]$ is connected to $T_c[a_d - 1]$, the suffix of the shape, and likewise is on the same slope.

Part (ii). Consider two maximal sequences of SA_{T_c} , $A_i = a_1, a_2, \dots, a_f$ and $A_{i+1} = a'_1, a'_2, \dots, a'_g$ with $f, g > 1$. From part (i), A_i and A_{i+1} both describe lines which have the same slope, namely the relative number of b's to a's in $T_c[1..k - 1]$. Since the lines do not have any point in common they must be parallel.

Part (iii). Here, either of the above sequences $A_i = a_1, a_2, \dots, a_f$ and $A_{i+1} = a'_1, a'_2, \dots, a'_g$ may index a single character, that is $f, g \geq 1$. Since $a'_1 < a_f$, $T_c[a_f..n]$ is a proper suffix of $T_c[a'_1..n]$. Suppose that $a_f \neq n$. For $f \geq 1$, if $T_c[a'_1..a_f - 1]$ does not contain a b, then this would imply $\text{ISA}_{T_c}[a'_1] < \text{ISA}_{T_c}[a_f]$, thus showing that the point for $T_c[a'_1]$ is vertically further from S than that for $T_c[a_1]$. If $a_f = n$, then $f > 1$, and since a_f is the index of the smallest suffix starting b, $T_c[a_1..a_f] = a..b$ while $T_c[a'_1..n] = b..b$. Further, $a'_1 = k$ and so $a_1 < a'_1$. Since $T_c[a'_1] = b$, the point given by $T_c[a'_1]$ is vertically further from S than the line given by the points $T_c[a_1]$ and $T_c[a_f]$. The rest follows from parts (i) & (ii). \square

In a practical geometric application, Lemma 4.5 implies that for a maximal increasing sequence $A_i = a_1, a_2, \dots, a_f$, the associated geometric path needs only be stored for $T_c[a_1..a_2 - 1]$ due to repetitions of the path structure. The lemma is illustrated in Fig. 12:

4.2. Balanced words

A binary string T is called a *balanced word* if for each character $c \in \{a, b\}$, the number of occurrences of c in U and in V differ by at most one, for all pairs of substrings U and V with $|U| = |V|$ of the infinite concatenation $T \cdot T \cdots$ of T .

Lemma 4.6. Let T be a string over the binary alphabet $\{a, b\}$ with an arithmetically progressed suffix array $P = [p_1, \dots, p_n]$ with ratio k . Given that $p_1 \neq k - 1$, T is balanced.

Proof. According to Theorem 3.3 for T the BWT defined on the suffix array and the BWT defined on the BWT matrix are identical. From Corollary 3.5 we know that the BWT of T has the shape $b^x a^y$. By [39, Thm. 2] for binary words the following conditions are equivalent:

1. There are two coprime numbers p and q such that $\text{BWT}_T = b^p a^q$;
2. T is balanced.

We conclude the proof by showing that x and y are co-prime. Since $p_1 \neq k - 1$ we have $p_1 \in \{1, n\}$. Using the results from Section 3.4 the number of a's in T is $n - k^{-1}$. Thus the number of b's is k^{-1} . As k^{-1} is co-prime to n , $n - k^{-1}$ must be co-prime to k^{-1} . \square

4.3. Relation to the Fibonacci word sequence

The Fibonacci sequence is a sequence of binary strings $\{F_m\}_{m \geq 1}$ with $F_1 := b$, $F_2 := a$, and $F_m := F_{m-1}F_{m-2}$. Then F_m and $f_m := |F_m|$ are called the m th *Fibonacci word* and the m th *Fibonacci number*, respectively.

Köppl and I [11, Thm. 1] observed that the suffix array of F_m for even m is the arithmetically progressed permutation SA_{F_m} with ratio $f_{m-2} \bmod f_m$ and $SA_{F_m}[1] = f_m$. Theorem 3.9 generalizes this observation by characterizing all binary strings whose suffix arrays are arithmetically progressed. Hence, F_m must coincide with Case 1 of Theorem 3.9 since it ends with character a.

Lemma 4.7. *The Fibonacci word F_m for even m is given by*

$$F_m[i] = \begin{cases} a & \text{if } 1 + i \cdot f_{m-2} \bmod f_m \leq f_{m-1}, \text{ or} \\ b & \text{otherwise.} \end{cases}$$

Proof. We use the following facts:

- The greatest common divisor of f_i and f_j is the Fibonacci number whose index is the greatest common divisor of i and j [40, Fibonacci numbers]. Hence, f_{m-1} and f_m are coprime for every $m \geq 2$.
- $f_{m-2}^2 \bmod f_m = 1$ holds for every even $m \geq 3$ [41]. Hence, $k^{-1} = k = f_{m-2}$.
- By definition, $F_m[f_m] = a$ if m is even, and therefore $SA_{F_m}[1] = f_m$.

The split position p_s is $p_s = f_m - k = f_{m-1}$. So $SA_{F_m}[f_m] = f_m - k = f_m - f_{m-2}$. By Theorem 3.10, $ISA_{F_m}[i] = if_{m-2} - SA_{F_m}[f_m]f_{m-2} \bmod f_m = if_{m-2} + 1 \bmod f_m$, where $-SA_{F_m}[f_m]f_{m-2} \bmod f_m = (f_{m-2} - f_m)f_{m-2} \bmod f_m = 1 - f_m f_{m-2} \bmod f_m = 1$. The rest follows from Eq. (3). \square

Let \bar{F}_m denote the m th Fibonacci word whose a's and b's are exchanged, i.e., $\bar{F}_m = a \Leftrightarrow F_m = b$.

Lemma 4.8. *$SA_{\bar{F}_m}$ is arithmetically progressed with ratio f_{m-2} for odd m .*

Proof. Since $\bar{F}_m[|\bar{F}_m|] = a$ for odd m , $\bar{F}_m[f_m..]$ is the lexicographically smallest suffix. Hence, $SA := SA_{\bar{F}_m} = |\bar{F}_m|$. If SA is arithmetically progressed with ratio k , then its split position must be $p_s = n - k - 1$ according to Theorem 3.9. We show that $k = f_{m-2}$ by proving

$$\begin{aligned} \bar{F}_m[f_m..] < \bar{F}_m[f_m + f_{m-2} \bmod f_m..] < \bar{F}_m[f_m + 2f_{m-2} \bmod f_m..] < \dots \\ < \bar{F}_m[f_m + (f_m - 1)f_{m-2} \bmod f_m..] \end{aligned}$$

in a way similar to [11, Lemma 8]. For that, let \bar{S} of a binary string $S \in \{a, b\}^*$ denote S after exchanging a's and b's (i.e., $\bar{S} = a \Leftrightarrow S = b$). Further, let $<$ be the relation on strings such that $S < T$ if and only if $S < T$ and S is *not* a prefix of T . We need this relation since $S < T \iff \bar{S} > \bar{T}$ while $S < T$ and $\bar{S} < \bar{T}$ holds if S is a prefix of T .

- For $i \in [1..f_{m-1})$, we have $F_m[i..] > F_m[i + f_{m-2}..]$ due to [11, Lemma 7], thus $\bar{F}_m[i..] < \bar{F}_m[i + f_{m-2}..]$.
- For $i \in (f_{m-1}..f_m]$, since $F_m = F_{m-1}F_{m-2} = F_{m-2}F_{m-3}F_{m-2}$, $F_m[i..]$ is a prefix of $F_m[i - f_{m-1}..] = F_m[i + f_{m-2} \bmod f_m..]$. Therefore, $\bar{F}_m[i..]$ is a prefix $\bar{F}_m[i + f_{m-2} \bmod f_m..]$ and $\bar{F}_m[i..] < \bar{F}_m[i + f_{m-2} \bmod f_m..]$.

Since f_m and f_{m-2} are coprime, $\{i + f_{m-2} \bmod f_m \mid i > 0\} = [1..n]$. Starting with the smallest suffix $\bar{F}_m[f_m..]$, we end up at the largest suffix $\bar{F}_m[f_m + (f_m - 1)f_{m-2} \bmod f_m..]$ after $m - 1$ arithmetic progression steps of the form $\bar{F}_m[f_m + if_{m-2} \bmod f_m..]$ for $i \in [0..f_m - 1]$. By using one of the two above items we can show that these arithmetic progression steps yield a list of suffixes sorted in lexicographically ascending order. \square

5. Extension to larger alphabets

In this section we extend our results of Section 3 to alphabets of arbitrary size. Let $P = [p_1, \dots, p_n]$ be an arithmetically progressed permutation with ratio k . Let $\Sigma = [1..\sigma]$ be an alphabet of size σ where the order is given by $1 < 2 < \dots < \sigma$. To construct a string T over the alphabet Σ having P as its suffix array we proceed similarly to the construction presented in the previous sections: First we split P into subarrays S_1, \dots, S_σ , then for each subarray S_i we assign the character i to each position $p \in S_i$. When splitting P into subarrays there are some fixed positions where we are required to split P , while the remaining splitting positions can be chosen freely. Let σ_{\min} be the size of the minimal alphabet over which there is a string having P as its suffix array. Then there are $\sigma_{\min} - 1$ positions where we are required to split P . Those required splitting positions are after the following entries, modulo n :

$$\begin{aligned} \{p_1 - k - 1, n - k\} & \quad \text{if } p_1 \notin \{1, k + 1, n\}, \\ \{n - k\} & \quad \text{if } p_1 \in \{1, k + 1\}, \\ \{p_1 - k - 1\} & \quad \text{if } p_1 = n \text{ and } k \neq n - 1. \end{aligned}$$

Example 5.1. Consider the alphabet $\Sigma = \{a, b, c, d, e\}$ with $\sigma := |\Sigma| = 5$ and order $a < b < c < d < e$, and the permutation $P = [5, 2, 7, 4, 1, 6, 3, 8]$ which has length $n = 8$, is arithmetically progressed with ratio $k = 5$ and starts with the entry $p_1 = 5$. The minimum alphabet size to construct a string having P as suffix array is $\sigma_{\min} = 3$. Thus there are two required splitting positions. Those are after 3 and 7. Using vertical bars $|$ to denote the splitting positions we obtain $P = [5, 2, 7|4, 1, 6, 3|8]$. There are $\sigma - \sigma_{\min}$ splitting positions left which can be chosen freely. Assume we choose to split after entries 1 and 2. This leads to $P = [5, 2|7|4, 1|6, 3|8]$. Assigning characters we obtain the string $T = \text{cadcadbe}$.

At the beginning of this paper we looked at the strings over an alphabet of size σ having the suffix array $[n, n-1, \dots, 1]$. The number of those strings is given by [Theorem 2.2](#) and is bounded by a polynomial in n and σ . We extend upon that result: For an arbitrary permutation, which is not necessarily arithmetically progressed we give a bound on the number of strings with that permutation as suffix array. For a fixed arithmetically progressed permutation P we give an exact formula for the number of strings having P as their suffix array. We conclude that in total the number of strings having an arithmetically progressed suffix array is bounded by a polynomial in n and σ .

Lemma 5.2. *Let Σ be an alphabet of size $\sigma = |\Sigma|$. Given a permutation P of length n , there are at most $\binom{n+\sigma-1}{n}$ strings of length n over the alphabet Σ having P as their suffix array.*

Proof. Let the alphabet Σ be given by $\Sigma = [1..\sigma]$ with the order $1 < 2 < \dots < \sigma$. Let T be a string over Σ of length n whose suffix array is P . Permuting T by its suffix array P we obtain the string T' with $T'[i] = T[P[i]]$. As P is the suffix array of T we have $T'[i..n] < T'[j..n]$ and thus $T'[i] \leq T'[j]$ for all $i < j$. Therefore T' is given by $T' = 1^{t_1} \dots \sigma^{t_\sigma}$, where t_i describes the number of occurrences of i in T' . The problem of finding the number of strings with this particular shape can be reduced to the classic stars and bars problem¹³ [[27](#), Chp. II, Sect. 5] with n stars and $\sigma - 1$ bars, yielding $\binom{n+\sigma-1}{n}$ possible strings. As the permutation P is a bijection on Σ^n , there is only one string T that, permuted by P , gives the string T' . Thus there are at most $\binom{n+\sigma-1}{n}$ strings having P as their suffix array. There may be less because it is possible that a string V , which is a pre-image of some $T' = 1^{t_1} \dots \sigma^{t_\sigma}$ under the permutation P , has a suffix array different to P . \square

Lemma 5.3. *Let Σ be an alphabet of size $\sigma = |\Sigma|$. Given a permutation P of length n , let σ_{\min} be the size of the smallest alphabet over which there exists a string having P as its suffix array. Then there are at most $\binom{n+\sigma-1}{\sigma-\sigma_{\min}}$ strings of length n over the alphabet Σ having P as their suffix array.*

Proof. As described at the beginning of this section, all strings over the alphabet Σ having the suffix array P can be constructed by splitting P into $\sigma - 1$ subarrays. There are $\sigma_{\min} - 1$ positions where we are required to split P and $\sigma - \sigma_{\min}$ splitting positions that can be chosen freely. We model the selection of the positions that can be chosen freely using the stars and bars problem¹⁴ [[27](#), Chp. II, Sect. 5]. The stars represent the entries of the permutation, the bars the splitting positions. Assume we start with $n + \sigma - 1$ stars. Then $\sigma - \sigma_{\min}$ bars (splitting positions) can be freely chosen. There are $\binom{n+\sigma-1}{\sigma-\sigma_{\min}}$ ways to do this. Then we replace the stars at the $\sigma_{\min} - 1$ required splitting positions with bars. This gives us n stars and $\sigma - 1$ bars, which we map to a partition of P into σ subarrays. \square

Lemma 5.4. *Let Σ be an alphabet of size $\sigma = |\Sigma|$. Then the number of strings of length n with an arithmetically progressed suffix array is at most $p(n, \sigma) = n(n - 1)\binom{n+\sigma-1}{\sigma-\sigma_{\min}}$.*

Proof. Each arithmetically progressed permutation P of length n can be described by two parameters: Its first entry $P[1]$ and its ratio k . For $P[1]$ there are n different options, for k there are at most $n - 1$ different options. Thus the number of arithmetically progressed permutations of length n is at most $n(n - 1)$. By [Lemma 5.3](#) the number of strings of length n having a specific permutation as suffix array is bounded by $\binom{n+\sigma-1}{\sigma-\sigma_{\min}}$. Putting those two facts together we obtain that the number of strings of length n with an arithmetically progressed suffix array is bounded by $p(n) = n(n - 1)\binom{n+\sigma-1}{\sigma-\sigma_{\min}}$. \square

Given a fixed alphabet Σ , by the above lemma, the number of strings of length n with an arithmetically progressed suffix array is bounded by a polynomial in n and σ .

6. Applications on meta strings

In this final section we overview some connections of suffix arrays with generalized forms of words and meta strings. A generalization of strings was proposed in the 1974 groundbreaking paper of Fischer and Paterson [[42](#)], where string matching was considered in a more general setting than with the usual solid letters, whereby either string could have *don't care*¹⁵ symbols, and was achieved in time nearly as fast as linear. Uncertain sequences, including indeterminate

¹³ See Footnote 9.

¹⁴ See Footnote 9.

¹⁵ Don't care symbols match with all characters.

strings, have application in inexact matching tasks, for instance, allowing for errors such as with Web interface data entry and Internet search. They are also useful for expressing DNA polymorphisms, that is biological sequence positions that can have multiple possibilities and encoded with IUPAC¹⁶ meta characters, for example N denotes any of the DNA nucleotides. A *codon* is a form of meta character whereby a sequence of three nucleotides encodes a specific amino acid and are used for protein expression; so a genetic code can be composed of concatenated codon units. The *truncated generalized suffix automaton* was introduced for indexing length-bounded *k*-factors of degenerate DNA and RNA sequences [43]. An *elastic-degenerate* string is a sequence of sets of strings used for succinctly representing a multiple alignment of a collection of closely related sequences (e.g. a pan-genome, that is all genes and genetic variation within a species), and also supports approximate pattern matching [44]. Sequence alignment is useful for inferring evolutionary relationships between biological sequences.

Daykin and Watson proposed a simple degenerate BWT, the *D-BWT* [45], constructed by applying lex-extension order (cf. Example 6.1) to relabel the sets and order the degenerate strings in the *D-BWT* matrix. Subsequently in [46], the *D-BWT* was applied to pattern matching using the backward search technique. This formalized and extended work implemented in [47] presenting a bioinformatics software tool, BWBBLE, for pattern matching on a pan-genome that they called a reference multi-genome.

In what follows, we first formally define indeterminate strings, and subsequently illustrate various approaches for defining an *indeterminate suffix array*: An *indeterminate string*¹⁷ $T^I = T^I[1..n]$ on an alphabet Σ is a sequence of nonempty subsets of Σ . Specifically, an indeterminate string T^I has the form $T^I = t_1 \cdots t_n$ where each t_i is a set of characters over Σ ; a singleton is known as a *solid letter*. For example,

$$T^I = \{c, a, b, e\}\{d\}\{c, a, b, e\}\{d\}\{d\}\{c, a, b, e, f, g\}$$

is a ternary border-free indeterminate string of length 6, with the indeterminate letters $\{c,a,b,e\}$, $\{d\}$, $\{c,a,b,e,f,g\}$ where $\{d\}$ is a singleton.

Example 6.1. Let $T^I = \{b\}\{b\}\{c,a\}\{b,d,a\}$ be an indeterminate string, then SA_{T^I} can be alternatively defined by the following approaches:

- We can apply the lex-extension order by treating each character set as a string whose characters are sorted, and use the lexicographic order of these strings to sort the sets, resulting in ranked meta characters: $\{a, b, d\} = A$; $\{a, c\} = B$; $\{b\} = C$. Then $T^I = CCBA$, and $SA_{T^I} = [4, 3, 2, 1]$, an arithmetically progressed permutation. This method enables linear-time construction of an indeterminate suffix array.
- Given we do not want the input to be rearranged, then we can form the meta characters as follows: $\{b\} = A$; $\{b, d, a\} = B$; $\{c, a\} = C$. Here, $T^I = AACB$ with suffix array $SA_{T^I} = [1, 2, 4, 3]$, which is not arithmetically progressed.
- Finally, we can incorporate the suffix arrays of the individual sets treated as strings, so the above suffix array $[1,2,4,3]$, with both indeterminate and solid letter positions, becomes

$$[(1, 1), (2, 1), ((4, 3), (4, 1), (4, 2)), ((3, 2), (3, 1))].$$

Clearly, there is a natural generalization of many types of well-known patterned strings over solid letters to the indeterminate form, such as *indeterminate Fibonacci words*, where the solid form can give a meta representation of the generalized form.

7. Conclusion and problems

Given an arithmetically progressed permutation P with ratio k , we studied the minimum alphabet size and the shape of those strings having P as their suffix array. Only in the case $P = [n, n - 1, \dots, 1]$, a unary alphabet suffices. For general $P = [p_1, \dots, p_n] \neq [n, n - 1, \dots, 1]$, there is exactly one such string on the binary alphabet if and only if $p_1 \in \{1, k + 1, n\}$. In all other cases, there is exactly one such string on the ternary alphabet. Altogether, for an arbitrary but fixed alphabet size σ , only a polynomial portion of the σ^n strings of length n have an arithmetically progressed suffix array.

We conclude by proposing some research directions.

- A natural question arising from this research is to characterize strings having arithmetic progression properties for the run length exponents of their BWTs, particularly for the bijective [48] or extended BWT [49], which are always invertible.

For example, given the arithmetically progressed permutation 3214, then the run-length compressed string $a^3c^2b^4$ (a) matches the permutation 3214 and (b) is a BWT image because its inverse is $b^2cb^2ca^3$, which can be computed by the Last-First mapping. However, for the same permutation, $a^3b^2b^4$ does not work since it is not a BWT image. Further examples of arithmetically progressed BWT exponents are: $a^3c^4b^2$, $a^4c^3e^2d^6b^5$, and $a^4c^3e^2f^7d^6b^5$.

¹⁶ International Union of Pure and Applied Chemistry.

¹⁷ Also known as *degenerate string* in the literature, particularly in the field of microbiology.

- Arithmetic properties can likewise be considered for the following stringology integer arrays:
 - Firstly the *longest common prefix* (LCP) array LCP, whose entry $LCP[i]$ is the length of the longest common prefix of the *lexicographically* i th smallest suffix with its lexicographic predecessor for $i \in [2..n]$.
 - Given a string $T \in \Sigma^+$ of length n , the *prefix table* P_T of T is given by $P_T[i] = LCP(T, T[i..n])$ for $i \in [1..n]$; equivalently, the *border table* B_T of T is defined by

$$B_T[i] = \max\{|S| \mid S \text{ is a border of } T[1..i]\} \text{ for } i \in [1..n].$$
 - Integer *prefix lists* are more concise than prefix tables and give the lengths of overlapping LCPs of T and suffixes of T (cf. [50]).
 - The i th entry of the *Lyndon array* $\lambda = \lambda_T[1..n]$ of a given string $T = T[1..n]$ is the length of the longest Lyndon word that is a prefix of $T[i..]$ – reverse engineering in [51] includes a linear-time test for whether an integer array is a Lyndon array. Likewise, the *Lyndon factorization array* $F = F_T[1..n]$ of T stores in its i th entry the size of the Lyndon factorization (i.e., the number of factors) of the suffix $T[i..n]$. The problems are to characterize those arithmetic progressions that define a valid Lyndon array, respectively Lyndon factorization array. For example, consider the string $T = azyx$, then its Lyndon array is $\lambda_T = [4, 1, 1, 1]$, while the Lyndon factorization array is $F_T = [1, 3, 2, 1]$. Trivially, for $T = abc\dots z$ the Lyndon array is an arithmetic progression and likewise for the Lyndon factorization array of $T = z^t y^t x^t \dots a^t$ for $z > y > x > \dots > a$.
- A challenging research direction is to consider arithmetic progressions for multi-dimensional suffix arrays and Fibonacci word sequences.

Data availability

No data was used for the research described in the article.

Acknowledgments

We thank the reviewer for the careful proofreading and perceptive comments, which helped us polish this article. We also thank Gabriele Fici for the initial help in guiding this research started at the 2019 Pisa StringMasters.

In line with the community norm, we used alphabetic author order. We confirm that all authors contributed to the mathematical concepts and solutions along with writing and editing the manuscript, where Florian Stober is distinguished as the lead.

Funding

This research was part-funded by JSPS KAKENHI, Japan with grant numbers JP21K17701 and JP23H04378, and by the European Regional Development Fund through the Welsh Government [Grant Number 80761-AU-137 (West)]:



References

- [1] U. Manber, E.W. Myers, Suffix arrays: A new method for on-line string searches, *SIAM J. Comput.* 22 (5) (1993) 935–948.
- [2] M. Burrows, D.J. Wheeler, A Block Sorting Lossless Data Compression Algorithm, Technical Report 124, Digital Equipment Corporation, Palo Alto, California, 1994.
- [3] D. Adjeroh, T. Bell, A. Mukherjee, *The Burrows–Wheeler Transform: Data Compression, Suffix Arrays, and Pattern Matching*, Springer, 2008.
- [4] T. Bingmann, J. Fischer, V. Osipov, Inducing suffix and LCP arrays in external memory, *ACM J. Exp. Algorithmics* 21 (1) (2016) 2.3:1–2.3:27.
- [5] J. Fischer, F. Kurpicz, Lightweight distributed suffix array construction, in: *Proc. ALENEX*, 2019, pp. 27–38.
- [6] S. Mantaci, A. Restivo, M. Sciortino, Burrows–Wheeler transform and Sturmian words, *Inform. Process. Lett.* 86 (5) (2003) 241–246.
- [7] M. Christodoulakis, C.S. Iliopoulos, Y.J.P. Ardila, Simple algorithm for sorting the Fibonacci string rotations, in: *Proc. SOFSEM*, in: LNCS, vol. 3831, 2006, pp. 218–225.
- [8] J. Simpson, S.J. Puglisi, Words with simple Burrows–Wheeler transforms, *Electron. J. Combin.* 15 (1) (2008).
- [9] W. Rytter, The structure of subword graphs and suffix trees of Fibonacci words, *Theoret. Comput. Sci.* 363 (2) (2006) 211–223.
- [10] J. Berstel, A. Savelli, Crochemore factorization of Sturmian and other infinite words, in: *Proc. MFCS*, in: LNCS, vol. 4162, 2006, pp. 157–166.
- [11] D. Köppl, T. I. Arithmetics on suffix arrays of Fibonacci words, in: *Proc. WORDS*, in: LNCS, vol. 9304, 2015, pp. 135–146.
- [12] S. Giuliani, Z. Lipták, R. Rizzi, When a dollar makes a BWT, in: *Proc. ICTCS*, in: *CEUR Workshop Proceedings*, vol. 2504, 2019, pp. 20–33.
- [13] G. Nong, S. Zhang, W.H. Chan, Linear suffix array construction by almost pure induced-sorting, in: *Proc. DCC*, 2009, pp. 193–202.
- [14] M. Crochemore, R. Grossi, J. Kärkkäinen, G.M. Landau, Computing the Burrows–Wheeler transform in place and in small space, *J. Discrete Algorithms* 32 (2015) 44–52.
- [15] D. Köppl, D. Hashimoto, D. Hendrian, A. Shinohara, In-place bijective Burrows–Wheeler transformations, in: *Proc. CPM*, in: *LIPICs*, 2020, pp. 21:1–21:15.

- [16] A. Alatabbi, J.W. Daykin, N. Mhaskar, M.S. Rahman, W.F. Smyth, Applications of V -order: Suffix arrays, the Burrows–Wheeler transform & the FM-index, in: Proc. WALCOM, in: LNCS, vol. 11355, 2019, pp. 329–338.
- [17] F. Gray, Pulse code communication, 1953, (filed 1947). U.S. Patent 2, 632, 058.
- [18] B. Chapin, S.R. Tate, Higher compression from the Burrows–Wheeler transform by modified sorting, in: Data Compression Conference, DCC 1998, IEEE Computer Society, 1998, p. 532.
- [19] J.W. Daykin, R. Groult, Y. Guesnet, T. Lecroq, A. Lefebvre, M. Léonard, É. Prieur-Gaston, Binary block order Rouen transform, Theoret. Comput. Sci. 656 (2016) 118–134.
- [20] J.W. Daykin, D. Köppl, D. Kübel, F. Stober, On Arithmetically Progressed Suffix Arrays, 2020, pp. 96–110.
- [21] A. Amir, A. Levy, L. Reuveni, The practical efficiency of convolutions in pattern matching algorithms, Fund. Inform. 84 (1) (2008) 1–15.
- [22] A. Amir, Y. Aumann, A. Levy, Y. Roshko, Quasi-distinct parsing and optimal compression methods, Theoret. Comput. Sci. 422 (2012) 1–14.
- [23] A. Amir, A. Apostolico, E. Eisenberg, G.M. Landau, A. Levy, N. Lewenstein, Detecting approximate periodic patterns, Theoret. Comput. Sci. 525 (2014) 60–67.
- [24] M. Lothaire, Combinatorics on Words, second ed., Cambridge Mathematical Library. Cambridge University Press, 1997.
- [25] C. Hohlweg, C. Reutenauer, Lyndon words, permutations and trees, Theoret. Comput. Sci. 307 (1) (2003) 173–178.
- [26] S. Mantaci, A. Restivo, G. Rosone, M. Sciortino, Suffix array and Lyndon factorization of a text, J. Discrete Algorithms 28 (2014) 2–8.
- [27] W. Feller, An Introduction to Probability Theory and Its Applications, Wiley, 1968.
- [28] J. Bernoulli, Sur une nouvelle espèce de calcul, Recueil Pour Astron. (1771) 255–284.
- [29] J.-P. Borel, F. Laubie, Quelques mots sur la droite projective Réelle, J. Théor. Nombres Bordeaux 5 (1) (1993) 23–51.
- [30] J. Berstel, A. Lauve, C. Reutenauer, F. Saliola, Combinatorics on Words: Christoffel Words and Repetition in Words, in: CRM Monograph Series, vol. 27, American Mathematical Society, 2008.
- [31] J. Borel, C. Reutenauer, On christoffel classes, RAIRO Theor. Inform. Appl. 40 (1) (2006) 15–27.
- [32] K.T. Chen, R.H. Fox, R.C. Lyndon, Free differential calculus, IV – the quotient groups of the lower central series, Ann. of Math. 68 (1958) 81–95.
- [33] J.-P. Duval, Factorizing words over an ordered alphabet, J. Algorithms 4 (4) (1983) 363–381.
- [34] F. Bassino, J. Clément, C. Nicaud, The standard factorization of Lyndon words: an average point of view, Discret. Math. 290 (1) (2005) 1–25.
- [35] A.I. Shirshov, On bases for free Lie algebra, Algebra Log. 1 (1) (1962) 14–19, (in Russian).
- [36] X.G. Viennot, Algèbres de Lie Libres et Monoïdes Libres, in: Lecture Notes in Mathematics, vol. 691, Springer, Berlin, 1978, p. 124.
- [37] G. Melançon, Visualisation de Graphes et Combinatoire des Mots (Thèse d'Habilitation à Diriger les Recherches), Université Bordeaux I, 1999.
- [38] D.E. Daykin, J.W. Daykin, W.F. Smyth, Combinatorics of unique maximal factorization families (UMFFs), Fund. Inform. 97 (3) (2009) 295–309.
- [39] A. Restivo, G. Rosone, Balanced words having simple Burrows–Wheeler transform, in: International Conference on Developments in Language Theory, Springer, 2009, pp. 431–442.
- [40] D. Wells, Prime Numbers: The Most Mysterious Figures in Math, Wiley, 2005.
- [41] V. Hoggatt, M. Bicknell-Johnson, Composites and primes among powers of Fibonacci numbers increased or decreased by one, Fibonacci Quart. 15 (1977) 2.
- [42] M.J. Fischer, M.S. Paterson, String matching and other products, in: Complexity of Computation, Vol. 7, 1974, pp. 113–125.
- [43] T. Flouri, C.S. Iliopoulos, M.S. Rahman, L. Vagner, M. Voráček, Indexing factors in DNA/RNA sequences, in: Proc. BIRD, in: Communications in Computer and Information Science, vol. 13, Springer, 2008, pp. 436–445.
- [44] G. Bernardini, N. Pisanti, S.P. Pissis, G. Rosone, Approximate pattern matching on elastic-degenerate text, Theoret. Comput. Sci. 812 (2020) 109–122.
- [45] J.W. Daykin, B.W. Watson, Indeterminate string factorizations and degenerate text transformations, Math. Comput. Sci. 11 (2) (2017) 209–218.
- [46] J.W. Daykin, R. Groult, Y. Guesnet, T. Lecroq, A. Lefebvre, M. Léonard, L. Mouchard, É. Prieur, B.W. Watson, Efficient pattern matching in degenerate strings with the Burrows–Wheeler transform, Inform. Process. Lett. 147 (2019) 82–87.
- [47] L. Huang, V. Popic, S. Batzoglou, Short read alignment with populations of genomes, Bioinformatics 29 (13) (2013) 361–370.
- [48] J.Y. Gil, D.A. Scott, A bijective string sorting transform, 2012, arXiv:1201.3077.
- [49] S. Mantaci, A. Restivo, G. Rosone, M. Sciortino, An extension of the Burrows–Wheeler transform, Theoret. Comput. Sci. 387 (3) (2007) 298–312.
- [50] J. Clément, L. Giambruno, Representing prefix and border tables: results on enumeration, Math. Structures Comput. Sci. 27 (2) (2017) 257–276.
- [51] J.W. Daykin, F. Franek, J. Holub, A.S.M.S. Islam, W.F. Smyth, Reconstructing a string from its Lyndon arrays, Theoret. Comput. Sci. 710 (2018) 44–51.