



On average sequence complexity[☆]

Svante Janson^a, Stefano Lonardi^{b,*}, Wojciech Szpankowski^c

^a*Department of Mathematics, Uppsala University, Uppsala, Sweden*

^b*Department of Computer Science and Engineering, University of California, Riverside, CA 92521-0135, USA*

^c*Department of Computer Sciences, Purdue University, West Lafayette, IN, USA*

Received 18 November 2003; accepted 23 June 2004

Communicated by H. Prodinger

Abstract

In this paper we study the average behavior of the number of distinct substrings in a text of size n over an alphabet of cardinality k . This quantity is called the *complexity index* and it captures the “richness of the language” used in a sequence. For example, sequences with low complexity index contain a large number of repeated substrings and they eventually become periodic (e.g., tandem repeats in a DNA sequence). In order to identify unusually low- or high-complexity strings one needs to determine how far are the complexities of the strings under study from the average or maximum string complexity. While the maximum string complexity was studied quite extensively in the past, to the best of our knowledge there are no results concerning the average complexity. We first prove that for a sequence generated by a mixing model (which includes Markov sources) the average complexity is asymptotically equal to $n^2/2$ which coincides with the maximum string complexity. However, for a memoryless source we establish a more precise result, namely the average string complexity is $n^2/2 - n \log_k n + (1 + (1 - \gamma)/\ln k + \phi_k(\log_k n) + o(1))n$ where $\gamma \approx 0.577$ and $\phi_k(x)$ is a periodic function with a small amplitude for small alphabet size.

© 2004 Elsevier B.V. All rights reserved.

Keywords: Complexity index; Sequence/subword complexity; Suffix tree

[☆] An preliminary version of this work was presented at the *Symposium on Combinatorial Pattern Matching*, Istanbul, Turkey, and included in its *Proceedings*, Lecture Notes in Computer Science, Vol. 3109, 2004, Springer, Berlin.

* Corresponding author.

E-mail address: stelo@cs.ucr.edu (S. Lonardi).

1. Introduction

In the last decades, several attempts have been made to capture mathematically the concept of “complexity” of a sequence. The notion is related to quite deep mathematical properties, including the rather elusive concept of randomness of a string (see, e.g., [14,19,22]).

In this paper, we are interested in studying a measure of complexity of a sequence called the *complexity index*. The complexity index captures the “richness of the language” used in a sequence. Formally, the *complexity index* $c(x)$ of a string x is equal to the number of distinct substrings in x (see e.g., [20]). The measure is simple but quite intuitive. Sequences with low complexity index contain a large number of repeated substrings and they eventually become periodic. However, in order to classify low complexity sequences one needs to determine average and maximum string complexity. In this paper we concentrate on the average string complexity.

We assume that sequences are generated by some probabilistic source (e.g., Bernoulli, Markov, etc.). As a consequence, the number $c(x)$ of distinct substrings can be modeled by a random variable over a discrete domain. Given a source emitting strings of size n over an alphabet of cardinality k , we call this random variable $C_{n,k}$. The main objective of this study is to give a detailed characterization of the expectation of the random variable $C_{n,k}$.

A related notion is that of the *l-subword complexity* or *l-spectrum* $c^l(x)$ of a string x , which is the number of distinct substrings of length l in x , for $1 \leq l \leq |x|$. We define $C_{n,k}^l$ to be the random variable associated with the number of distinct words of size l in a random string of size n over an alphabet of cardinality k . Clearly, $C_{n,k} = \sum_{l=1}^n C_{n,k}^l$.

The idea of using the complexity index or the *l-spectrum* to characterize sequence statistically has a long history of applications in several fields, such as data compression, computational biology, data mining, computational linguistics, among others.

In dictionary-based data compression, the average length of the pointer is connected with the expected size of the dictionary which in turns depends on the number of distinct subwords (see, e.g., [6]). Low-complexity strings contain more repeated substrings and therefore one can expect them to be more compressible than strings with high complexity index. For example, in [14] bounds between subword complexity and Lempel–Ziv complexity are established.

In the analysis of biosequences, the problem of characterizing the “linguistic complexity” of DNA or proteins is quite old. In the early days of computational biology, it was almost routine to compute the number and/or the frequency of occurrences of substrings and draw conclusions about the string under study based on those counts (see [7,12,15,16,23], just to mention a few).

In these and several other application domains, the typical problem associated with the complexity index is to determine whether a particular sequence x has a *statistically significant* complexity index. An example of a significance score proposed in a recent paper by Troyanskaya et al. [31] in the context of the analysis of prokaryotic genomes, is the following:

$$s(x) = c(x) - \max\{C_{n,k}\} = c(x) - \sum_{i=1}^n \min(k^i, n - i + 1).$$

Here, the authors compare the observed complexity $c(x)$ with the maximum possible complexity for a string of size n over an alphabet of cardinality k . Note however, that the score disregards both the distribution of $C_{n,k}$, and the probabilistic characteristics of the source.

A more statistically sound approach would entail the following steps. First, select an appropriate model for the source that emitted x (Bernoulli, Markov, etc.). Then, measure the statistical significance as a function of the discrepancy between the observed complexity $c(x)$ and the model-based expectation.

This approach of standardizing the value of an observation with respect to the expectation and the standard deviation of the associated random variable is common practice in statistics. The underlying assumption is that the random variable is normally distributed. The standardized z -score for the complexity index would be

$$z(x) = \frac{c(x) - \mathbf{E}(C_{n,k})}{\sqrt{\mathbf{Var}(C_{n,k})}}$$

for a given string x . Although we do not know under which conditions $C_{n,k}$ is distributed normally, such a score is nonetheless more sound than other *ad hoc* scores.

A similar situation takes place when describing the significance of other events in texts, like the number of occurrences, periodicities, etc. Although the normality of the corresponding random variables can be proved under specific conditions, there is a consensus that standardized z -scores should be always preferred over simpler scores (see, e.g., [25] and references therein).

Given an observation x of the source, we would like to compute the statistical significance $z(x)$ of its complexity index. As far as we know, however, the moments $\mathbf{E}(C_{n,k})$ and $\mathbf{Var}(C_{n,k})$ have never been characterized before. The goal of this paper is to study $\mathbf{E}(C_{n,k})$. The asymptotic analysis of the variance remains an open problem.

In order to proceed, we need to introduce some standard concepts and notation about strings. The set Σ denotes a nonempty *alphabet of symbols* and a *string* over Σ is an ordered sequence of symbols from the alphabet. In the rest of the paper, we assume that $|\Sigma| = k$. Given a string x , the number of symbols in x defines the *length* $|x|$ of x . Henceforth, we assume $|x| = n$.

The *concatenation* (or *product*) of two strings x and y is denoted by xy , and corresponds to the operation of appending y to the last symbol of x . Let us decompose a text x in uvw , i.e., $x = uvw$ where u , v and w are strings over Σ . Strings u , v and w are called *substrings*, or *words*, of x .

We write $x_{[i]}$, $1 \leq i \leq |x|$ to indicate the i th symbol in x . We use $x_{[i,j]}$ shorthand for the substring $x_{[i]}x_{[i+1]} \dots x_{[j]}$ where $1 \leq i \leq j \leq |x|$, with the convention that $x_{[i,i]} = x_{[i]}$. Substrings in the form $x_{[1,j]}$ corresponds to the *prefixes* of x , and substrings in the form $x_{[i,n]}$ to the *suffixes* of x .

Finally, we recall that the *subword complexity* function $c^l(x)$ of a string x is defined as the number of distinct substrings of x of length l . The quantity $c(x) = \sum_{l=1}^n c^l(x)$ is the *complexity index* of x . Observe first that $c^l(x)$ is upper bounded by $\min(k^l, n - l + 1)$ since there are precisely $n - l + 1$ words of length l , of which at most k^l can be distinct. Therefore

$$c(x) \leq \sum_{l=1}^n \min(k^l, n - l + 1).$$

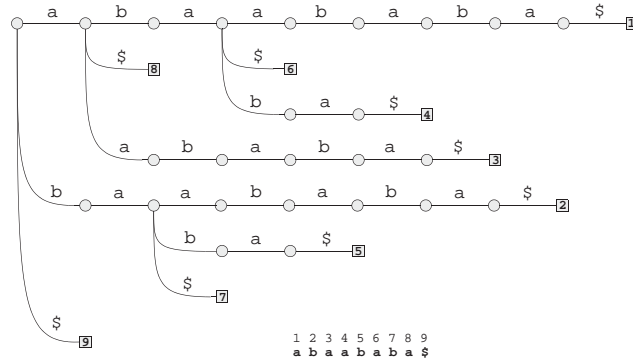


Fig. 1. The non-compact suffix trie T_x for $x = abaababa$. There are 24 internal non-root nodes, therefore $c(x) = 24$.

Note that if we choose $n \leq k$, then $\min(k^l, n - l + 1) = n - l + 1$ for all $1 \leq l \leq n$. Therefore when $n \leq k$, the bound simplifies to $c(x) \leq n(n + 1)/2$.

The value $c(x)$ is strongly connected with the structure of the *non-compact suffix trie* for x . A non-compact suffix trie of a string x is a digital search tree built from all the suffixes of x . The trie for a string of size n has $n + 1$ leaves, numbered 1 to $n + 1$, where leaf $n + 1$ correspond to an extra unique symbol $\$ \notin \Sigma$, and each edge is labeled with a symbol in Σ . No two edges outgoing from a node can have the same label. The trie has the property that for any leaf i , the concatenation of the labels on the path from the root the leaf i spells out exactly the suffix of x that starts at position i , that is $x_{[i,n]}$. The substrings of x can be obtained by spelling out the words from the root to any internal node of the tree. In other words, each internal node (except the root) is in one-to-one correspondence with each distinct substring of x . As a consequence, *the complexity index $c(x)$ can be obtained by counting the non-root internal nodes in the non-compact suffix trie for x* . This would take, however, $O(n^2)$ time and space. The non-compact suffix trie for $abaababa$ is illustrated in Fig. 1.

A faster solution to compute $c(x)$ involves the use of the *suffix tree* \tilde{T}_x of x . The suffix tree can be obtained by “path-compression” of the non-compact suffix trie, that is, by deleting internal nodes with only one child and coalescing consecutive edges in a single edge labeled by the concatenation of the symbols. If one deletes unary nodes only at the *bottom* of the non-compact suffix tries, the resulting tree is called *compact suffix trie*. The compact suffix trie and the suffix tree for $abaababa$ are shown in Fig. 2.

In practice, suffix trees can be built without the need of building the suffix trie first. In fact, several $O(n \log |\Sigma|)$ constructions are available. The algorithm by McCreight [21] and the one by Chen and Seiferas [9] are variations of the Weiner’s algorithm [34]. Note that these algorithms take only linear time for finite alphabets. All these constructions are *off-line* because they process the text right to left. An on-line algorithm by Ukkonen [32] achieves also linear time. Recently, Farach [11] proposed an optimal construction for large alphabets.

The unary nodes that have been removed in the compaction process are called *implicit nodes*. An edge labeled by a string of length $m + 1$ has m implicit nodes. The complexity

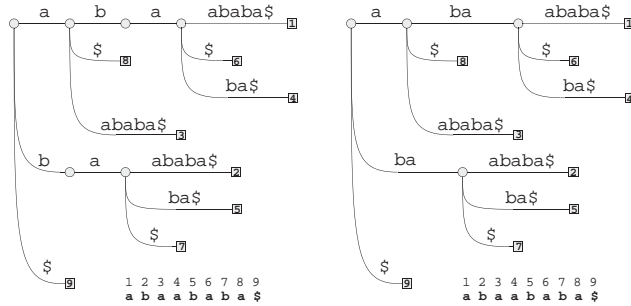


Fig. 2. The compact suffix trie (LEFT) and the suffix tree (RIGHT) for $x = abaababa$.

index $c(x)$ of a text-string x can be computed on the suffix tree by counting the number of implicit and explicit (non-root) nodes in the tree. As a consequence, $c(x)$ can be computed in $O(n)$ time and space. The relation between non-compact suffix tries and suffix trees will be used later in the paper to obtain the leading term of the complexity for a general probabilistic model.

Finally, we briefly describe some recent results regarding the maximum of $c^l(x)$. It is known that $c^l(x)$ is also strongly connected with the structure of the suffix tree \bar{T}_x . De Luca [10] proved that

$$\max\{c^l(x) : 1 \leq l \leq n\} = n - \max\{R, K\} + 1 = n - \max\{L, H\} + 1,$$

where K is the length of the shortest suffix of x that occurs only once, H is the length of the shortest prefix of x that occurs only once, $R - 1$ is the height of the deepest branching node in the suffix tree \bar{T}_x and $L - 1$ is the height of the deepest branching node in the suffix tree \bar{T}_{xR} . De Luca also gave a closed form for $c(x)$

$$c(x) = 1 + \frac{(n + K)(n - K + 1)}{2} - \sum_{j=2}^{|\Sigma|} \sum_{i=0}^R ig(j, i),$$

where $g(j, i)$ is the count of the words of length i which are branching nodes of the suffix tree with at least j children [10].

Shallit [26] derived a simpler bound for $c(x)$ for binary alphabets ($k = 2$)

$$c(x) \leq \frac{(n - d + 1)(n - d)}{2} + 2^{d+1} - 1 \sim \frac{n^2}{2},$$

where d is the unique integer such that $2^d + d - 1 \leq n < 2^{d+1} + d$. More importantly, he proved that the upper bound is attained for all n by using a property of the *de Bruijn* graphs. An extension of this result to larger alphabets was recently described in [13].

Kása [17] studied the probability distribution of random variable associated with the complexity index for the family of words of length equal to the size of the alphabet ($n = k$). He proved several facts about the random variable $C_{k,k}$, and he also conjectured a property of the smallest value of the complexity after which all the frequencies are

non-zero. The conjecture, proved later by Leve and Séébold [18], states that if one chooses $k = l(l+1)/2 + 2 + i$ where $l \geq 2$ and $0 \leq i \leq l$ then

$$\mathbf{P}(C_{k,k} = t) > 0 \quad \text{for all } t \text{ such that } \frac{l(l^2-1)}{2} + 3l + 2 + i(l+1) \leq t \leq \frac{k(k+1)}{2}.$$

In this paper we mostly deal with average string complexity. First, for a general probabilistic model (e.g., Markov source) we prove that the average complexity is asymptotically equal to $n^2/2$ which coincides with the maximum string complexity. We shall strengthen this result for unbiased memoryless sources. In particular, we prove that

$$\mathbf{E}(C_{n,k}) = \binom{n+1}{2} - n \log_k n + \left(\frac{1}{2} + \frac{1-\gamma}{\ln k} + \phi_k(\log_k n) \right) n + O(\sqrt{n \log n}),$$

where $\gamma \approx 0.577$ and $\phi_k(x)$ is a continuous function with period 1 and small amplitude for small alphabet size (e.g., $|\phi_2(x)| < 2 \times 10^{-7}$). To prove this result we use the Stein–Chen method together with the Mellin transform approach.

2. Main results

As a warm-up exercise, we studied the closed forms for $\mathbf{E}(C_{n,k})$ and $\mathbf{Var}(C_{n,k})$ for short strings (e.g., $n \leq 5$) for a symmetric memoryless source. Some facts about $\mathbf{P}(C_{n,k} = t)$ are immediate. For example

$$\begin{aligned} \mathbf{P}(C_{n,k} < n) &= 0, \\ \mathbf{P}(C_{n,k} = n) &= k^{1-n}, \\ \mathbf{P}\left(C_{n,k} > \frac{n(n+1)}{2}\right) &= 0 \quad \text{when } n \leq k, \\ \mathbf{P}\left(C_{n,k} > \sum_{i=1}^n \min(k^i, n-i+1)\right) &= 0 \quad \text{when } n > k. \end{aligned}$$

Following the lines by Kása [17], we were able to obtain closed form for the cases shown in Table 1 assuming a symmetric Bernoulli model for the source. Given the discrete distribution of $C_{n,k}$ one can easily compute the expectation and the variance of $C_{n,k}$.

Corollary 1. *The expectation and the variance of the random variable $C_{n,k}$ for $2 \leq n \leq 5$ over any alphabet of size k , under a symmetric Bernoulli source is*

$$\begin{aligned} \mathbf{E}(C_{2,k}) &= 3 - (1/k), \\ \mathbf{E}(C_{3,k}) &= 6 - (3/k), \\ \mathbf{E}(C_{4,k}) &= 10 - (6/k) + (1/k^2) - (1/k^3), \\ \mathbf{E}(C_{5,k}) &= 15 - (10/k) + (4/k^2) - (6/k^3) + (2/k^4), \end{aligned}$$

and

$$\begin{aligned} \mathbf{Var}(C_{2,k}) &= (k-1)/k^2, \\ \mathbf{Var}(C_{3,k}) &= 3(k-1)/k^2, \end{aligned}$$

Table 1
The probability distribution of $C_{n,k}$ for $2 \leq n \leq 5$ over any alphabet of size k , under a symmetric Bernoulli source

$n \rightarrow$	2	3	4	5
$\mathbf{P}(C_{n,k} = 2)$	$1/k$	0	0	0
$\mathbf{P}(C_{n,k} = 3)$	$1 - 1/k$	$1/k^2$	0	0
$\mathbf{P}(C_{n,k} = 4)$	0	0	$1/k^3$	0
$\mathbf{P}(C_{n,k} = 5)$	0	$3(k-1)/k^2$	0	$1/k^4$
$\mathbf{P}(C_{n,k} = 6)$	0	$(k-1)(k-2)/k^2$	0	0
$\mathbf{P}(C_{n,k} = 7)$	0	0	$3(k-1)/k^3$	0
$\mathbf{P}(C_{n,k} = 8)$	0	0	$4(k-1)/k^3$	0
$\mathbf{P}(C_{n,k} = 9)$	0	0	$6(k-1)(k-2)/k^3$	$3(k-1)/k^4$
$\mathbf{P}(C_{n,k} = 10)$	0	0	$(k-1)(k-2)(k-3)/k^3$	0
$\mathbf{P}(C_{n,k} = 11)$	0	0	0	$10(k-1)/k^4$
$\mathbf{P}(C_{n,k} = 12)$	0	0	0	$2(k-1)(3k-5)/k^4$
$\mathbf{P}(C_{n,k} = 13)$	0	0	0	$19(k-1)(k-2)/k^4$
$\mathbf{P}(C_{n,k} = 14)$	0	0	0	$10(k-1)(k-2)(k-3)/k^4$
$\mathbf{P}(C_{n,k} = 15)$	0	0	0	$(k-1)(k-2)(k-3)(k-4)/k^4$
$\mathbf{P}(C_{n,k} \geq 16)$	0	0	0	0

$$\begin{aligned} \mathbf{Var}(C_{4,k}) &= (k-1)(6k^4 - 5k^3 + 12k^2 - k + 1)/k^6, \\ \mathbf{Var}(C_{5,k}) &= 2(k-1)(5k^6 - 10k^5 + 33k^4 - 28k^3 + 16k^2 - 10k + 2)/k^8. \end{aligned}$$

As it turns out, obtaining a closed form for $\mathbf{E}(C_{n,k})$ and $\mathbf{Var}(C_{n,k})$ for any n, k is a very challenging problem. In practical applications, moreover, having the closed form only for small values of n would be of limited interest. It is certainly more valuable to study the behavior of the moments of $C_{n,k}$ asymptotically, that is, when n is very large.

The main result of our study is a characterization of $\mathbf{E}(C_{n,k})$ for large n . In our first result we show that for quite general sources the average complexity asymptotically coincides with the maximal string complexity. We consider mixing sources in which the probability of two events A and B defined on two substrings separated by g symbols is bounded as follows: $(1 - \psi(g))\mathbf{P}(A)\mathbf{P}(B) \leq \mathbf{P}(AB) \leq (1 + \psi(g))\mathbf{P}(A)\mathbf{P}(B)$ where the mixing coefficient $\psi(g) \rightarrow 0$ as $g \rightarrow \infty$ (cf. [28] for a detailed definition).

Theorem 1. *Let $C_{n,k}$ be the complexity index of a string of length n generated by a strongly mixing stationary source. Then, for large n ,*

$$\mathbf{E}(C_{n,k}) = \binom{n+1}{2} - O(n \log n).$$

Hence $C_{n,k} = n^2/2 + O_p(n \log n)$, i.e., $(n^2/2 - C_{n,k})/n \log n$ is bounded in probability.

Proof. We start with a simple observation. For a given sequence x of size n , build a non-compact suffix trie and a compact suffix trie. Recall that in a compact trie we collapse all unary links at the *bottom* of the suffix trie. In other words, in a compact suffix trie a path from the root to an external node (containing a suffix) is the minimal prefix of any two suffixes

that distinguishes them. Also recall that the string complexity of x is the number of non-root internal nodes in the non-compact trie. We shall argue below that the most contribution to the string complexity comes from the nodes that are in the non-compact trie but not in the compact trie. The upper bound follows immediately from

$$C_{n,k} \leq n(n+1)/2.$$

To find a matching lower bound, we consider the compact and non-compact suffix tries. We know that a typical depth and height in a compact suffix trie is $O(\log n)$. More precisely let H_n be the height of a compact suffix tree. It was shown in [27] that (at least for $\psi(g)$ satisfying $\sum_{g \geq 0} \psi^2(g) < \infty$) $H_n / \ln n \rightarrow 2/h_1$ a.s., where h_1 is Rényi's entropy (cf. [28, p. 157]). More precisely, the proof shows (for any $\psi(g) \rightarrow 0$) that for any $\varepsilon > 0$

$$\mathbf{P}\left(H_n \leq \frac{2}{h_1}(1+\varepsilon)\log n\right) = 1 - O(\log n/n^\varepsilon). \quad (1)$$

We claim that the main contribution to $C_{n,k}$ comes from strings that are in the non-compact trie but not in the compact trie. In fact, the i th suffix string has $n-i$ internal nodes of which at least $n-i-H_n$ are not in the compact trie. These nodes all correspond to unique substrings of x , and thus

$$C_{n,k} \geq \sum_{i=1}^n (n-i-H_n) = \frac{1}{2}n(n+1) - nH_n.$$

By (1), for a suitable constant B and large n , $\mathbf{P}(H_n > B \log n) < n^{-1}$ and thus

$$\mathbf{E}\left(\frac{1}{2}n(n+1) - C_{n,k}\right) \leq n\mathbf{E}H_n \leq n(B \log n + n\mathbf{P}(H_n > B \log n)) = O(n \log n),$$

which completes the proof. \square

However, from a theoretical point of view the most interesting case is when the string is generated by an unbiased source (such a source should have the largest complexity). In this case, we are able to characterize very precisely the average complexity.

Theorem 2. *Let $C_{n,k}$ be the complexity index of a string generated by an unbiased memoryless source. Then the average l -subword complexity is*

$$\mathbf{E}(C_{n,k}^l) = k^l(1 - e^{-nk^{-l}}) + O(l) + O(nlk^{-l}). \quad (2)$$

Furthermore, for large n the average complexity index becomes

$$\mathbf{E}(C_{n,k}) = \binom{n+1}{2} - n \log_k n + \left(\frac{1}{2} + \frac{1-\gamma}{\ln k} + \phi_k(\log_k n)\right)n + O(\sqrt{n \log n})$$

where $\gamma \approx 0.577$ is Euler's constant and

$$\phi_k(x) = -\frac{1}{\ln k} \sum_{j \neq 0} \Gamma\left(-1 - \frac{2\pi i j}{\ln k}\right) e^{2\pi i j x}$$

is a continuous function with period 1. $|\phi_k(x)|$ is very small for small k : $|\phi_2(x)| < 2 \times 10^{-7}$, $|\phi_3(x)| < 5 \times 10^{-5}$, $|\phi_4(x)| < 3 \times 10^{-4}$.

Interestingly enough, the term $O(n)$ of the average complexity contains a fluctuating term $\phi_k(x)$. (Note that $\phi_2(\log_2 n)$ equals $-P_0(n)$ in [28, p. 359].) The formula

$$|\Gamma(-1 - iy)| = |\Gamma(-iy)/(-1 - iy)| = \left(y(1 + y^2) \sinh(\pi y)/\pi\right)^{-1/2} \tag{3}$$

[28, p. 42] shows that the coefficients in ϕ_k are small and decrease geometrically. Numerical evaluations give the bounds for $k = 2, 3, 4$ stated in the theorem and also, for example, $|\phi_k(x)| < 0.01$ for $k \leq 12$ and $|\phi_k(x)| < 0.1$ for $k \leq 200$. Even for very large k , this term is not very large; we have, still using (3) (we omit the details), $|\phi_k(x)| < 0.5$ for $k \leq 10^9$, and $|\phi_k(x)| < \ln \ln(k)/\pi$ for all k . The fluctuating phenomenon is quite common in asymptotics of discrete problems.

3. Proof of Theorem 2

In this section we prove Theorem 2. In order to simplify the notation we restrict ourselves to the binary case $k = 2$. Extension to $k > 2$ is straightforward.

Recall that $C_{n,2}^l$ is the number of distinct substrings of length l in our random binary string of size n , and let

$$A_l = \mathbf{E}(C_{n,2}^l).$$

Thus $C_{n,2} = \sum_{l=1}^n C_{n,2}^l$ and $\mathbf{E}(C_{n,2}) = \sum_{l=1}^n A_l$. Define

$$\delta_l = A_l + l - 1 - 2^l \left(1 - e^{-n2^{-l}}\right) \tag{4}$$

$$= A_l - (n + 1 - l) + 2^l \left(e^{-n2^{-l}} - 1 + n2^{-l}\right). \tag{5}$$

Then

$$\begin{aligned} \mathbf{E}(C_{n,2}) &= \sum_{l=1}^n A_l = \sum_{l=1}^n \left((n + 1 - l) + \delta_l - 2^l \left(e^{-n2^{-l}} - 1 + n2^{-l} \right) \right) \\ &= \binom{n + 1}{2} - \sum_{l=1}^n 2^l \left(e^{-n2^{-l}} - 1 + n2^{-l} \right) + \sum_{l=1}^n \delta_l. \end{aligned} \tag{6}$$

Below we will use the estimates $0 < 1 - e^{-x} < \min(1, x)$ and $0 < e^{-x} - 1 + x < \min(x, x^2)$ (for $x > 0$) several times. In particular,

$$0 < 2^l \left(1 - e^{-n2^{-l}}\right) < \min(2^l, n), \tag{7}$$

$$0 < 2^l \left(e^{-n2^{-l}} - 1 + n2^{-l}\right) < n^2 2^{-l}. \tag{8}$$

We begin by estimating δ_l . First (for short strings, i.e., for small l), we use $1 \leq C_{n,2}^l \leq 2^l$, and thus $0 \leq A_l \leq 2^l$ and, using (4) and (7),

$$\delta_l = O(2^l). \tag{9}$$

For long substrings, i.e., for large l 's, there is another simple estimate. Clearly, $0 \leq C_{n,2}^l \leq n - l + 1$. Observe that

$$\begin{aligned} \mathbf{E}(n - l + 1 - C_{n,2}^l) &\leq \mathbf{E}(\left| \{(i, j) : i < j \text{ and } x_{[i, i+l-1]} = x_{[j, j+l-1]}\} \right|) \\ &= \binom{n-l+1}{2} 2^{-l} \leq n^2 2^{-l} \end{aligned}$$

i.e.,

$$0 \leq n - l + 1 - A_l \leq n^2 2^{-l}.$$

Hence, using (5) and (8),

$$\delta_l = O(n^2 2^{-l}). \quad (10)$$

Note that (9) and (10) easily yield $\sum_{l=1}^n \delta_l = O(n)$, which by (6) yields $\mathbf{E}(C_{n,2})$ up to $O(n)$. In order to obtain our more precise result, we need a better estimate of δ_l when $2^l \approx n$.

Let x be a sequence generated by an i.i.d. source (i.e., $p(0) = p(1) = 1/2$) and let us define

$$\mathbf{P}(n, l) = \mathbf{P}(x_{[1,l]} \neq x_{[j, j+l-1]} \text{ for } j = 2, \dots, n - l + 1).$$

By counting each repeated substring in x only the last time it occurs, it is easily seen (by shift invariance) that

$$A_l = \sum_{m=l}^n \mathbf{P}(m, l). \quad (11)$$

Now fix l and m , and let us define $I_j = \mathbf{1}[x_{[1,l]} = x_{[j, j+l-1]}]$ for $j = 2, \dots, m - l + 1$. Then

$$\mathbf{E}(I_j) = \mathbf{P}(I_j = 1) = \mathbf{P}[x_{[1,l]} = x_{[j, j+l-1]}] = 2^{-l} \quad \text{for every } j \geq 2.$$

In the next lemma, we establish that I_i and I_j are uncorrelated when $i, j > l$.

Lemma 1. *If $i, j \geq l + 1$ and $i \neq j$ then $\mathbf{E}(I_i I_j) = \mathbf{P}(x_{[1,l]} = x_{[i, i+l-1]} = x_{[j, j+l-1]}) = 2^{-2l}$.*

Proof. Assume $i < j$. Scan the three substrings left to right. Each bit in $x_{[i, i+l-1]}$ is either a fresh random bit or coincides with an earlier bit in $x_{[j, j+l-1]}$. In any case, it fixes one new bit each in $x_{[1,l]}$ and $x_{[j, j+l-1]}$, so the probability of success in this step is 2^{-2} . \square

Observe that, if $j \geq l + 1$, to condition on I_j is the same as to change every bit in $x_{[j, j+l-1]}$ to the corresponding bit in $x_{[1,l]}$, leaving all other bits untouched. Let $x^{(j)}$ be the resulting string, and let $J_{ij} = \mathbf{1}[x_{[1,l]}^{(j)} = x_{[i, i+l-1]}^{(j)}]$. Clearly, $J_{ij} = I_i$ if $|i - j| \geq l$. Note also that Lemma 1 yields, when $i \neq j$ and $i, j > l$,

$$\mathbf{E}(J_{ij}) = \mathbf{E}(I_i | I_j = 1) = \frac{\mathbf{E}(I_i I_j)}{\mathbf{E}(I_j)} = 2^{-l}.$$

Let us now set, for fixed integers l and $m \geq 2l$,

$$W = \sum_{j=l+1}^{m-l+1} I_j.$$

Let $d_{TV}(X, Y)$ be the total variation distance between random variables X and Y . Then, with $\text{Po}(\lambda)$ being the Poisson distribution with mean λ , by the Stein–Chen method (cf. [5, p. 24]) we find, with $\lambda = \mathbf{E}W = (m - 2l + 1)2^{-l}$,

$$\begin{aligned} d_{TV} \left(W, \text{Po}((m - 2l + 1)2^{-l}) \right) &\leq \frac{\min(1, \lambda)}{\lambda} \sum_j \mathbf{E}(I_j) \mathbf{E} \left| I_j + \sum_{i \neq j} (I_i - J_{ij}) \right| \\ &\leq \frac{1}{\lambda} \sum_j \mathbf{E}(I_j) \left(\mathbf{E}(I_j) + \sum_{0 < |i-j| < l} (\mathbf{E}(I_i) + \mathbf{E}(J_{ij})) \right) \\ &\leq \frac{1}{\lambda} \sum_{j=l+1}^{m-l+1} 2l \cdot 2 \cdot 2^{-2l} \\ &= 4l2^{-l}. \end{aligned}$$

In particular,

$$\left| \mathbf{P}(W = 0) - e^{-(m-2l+1)2^{-l}} \right| \leq d_{TV} \left(W, \text{Po}((m - 2l + 1)2^{-l}) \right) = O(l2^{-l}). \quad (12)$$

Moreover, by the first moment method

$$\mathbf{P} \left(\sum_{j=2}^l I_j \neq 0 \right) \leq (l - 1)\mathbf{E}(I_j) = (l - 1)2^{-l}. \quad (13)$$

Observe that

$$\mathbf{P}(m, l) = \mathbf{P} \left(\sum_{j=2}^l I_j + W = 0 \right).$$

Then by (13) and (12)

$$\begin{aligned} \mathbf{P}(m, l) &= \mathbf{P}(W = 0) + O(l2^{-l}) = e^{-(m-2l+1)2^{-l}} + O(l2^{-l}) \\ &= e^{-(m-l)2^{-l}} + O(l2^{-l}). \end{aligned} \quad (14)$$

We have assumed $m \geq 2l$. However, by the first moment method directly, the same estimate holds for $l \leq m < 2l$ too.

We thus have, by (11) and (14) and summing a geometric series,

$$A_l = \sum_{m=l}^n \mathbf{P}(m, l) = \frac{1 - e^{-(n+1-l)2^{-l}}}{1 - e^{-2^{-l}}} + O(nl2^{-l}).$$

Since

$$\begin{aligned} \frac{1 - e^{-(n+1-l)2^{-l}}}{1 - e^{-2^{-l}}} &= \frac{1 - e^{-(n+1-l)2^{-l}}}{2^{-l}(1 + O(2^{-l}))} \\ &= (2^l + O(1))(1 - e^{-(n+1-l)2^{-l}}) \\ &= 2^l(1 - e^{-(n+1-l)2^{-l}}) + O(n2^{-l}) \\ &= 2^l(1 - e^{-n2^{-l}}) + O(l) + O(n2^{-l}), \end{aligned}$$

we find

$$A_l = 2^l(1 - e^{-n2^{-l}}) + O(l) + O(nl2^{-l})$$

which proves (2). Thus by (4)

$$\delta_l = O(l) + O(nl2^{-l}). \quad (15)$$

Using (9) for $1 \leq l \leq \log_2 \sqrt{n \ln n}$, (15) for $\log_2 \sqrt{n \ln n} < l \leq 2 \log_2 n$, and (10) for $2 \log_2 n < l \leq n$, we obtain

$$\sum_{l=1}^n \delta_l = O(n \log n)^{1/2}. \quad (16)$$

We turn to the first sum in (6). Let, for $x > 0$,

$$f(x) = \frac{e^{-x} - 1 + x \mathbf{1}[x < 1]}{x}.$$

Then $|f(x)| < x$ for $0 < x < 1$ and $|f(x)| < 1/x$ for $x \geq 1$. Hence,

$$\begin{aligned} \sum_{l=1}^n 2^l (e^{-n2^{-l}} - 1 + n2^{-l}) &= n \sum_{l=1}^n (f(n2^{-l}) + \mathbf{1}[2^l \leq n]) \\ &= n \sum_{l=1}^n f(n2^{-l}) + n \lfloor \log_2 n \rfloor \\ &= n \sum_{l=-\infty}^{\infty} f(n2^{-l}) + O(1) + n \lfloor \log_2 n \rfloor \\ &= n\psi(n) + n \log_2 n + O(1), \end{aligned} \quad (17)$$

where, with $\langle x \rangle = x - \lfloor x \rfloor$, the fractional part of x ,

$$\psi(x) = \sum_{l=-\infty}^{\infty} f(x2^{-l}) - \langle \log_2 x \rangle, \quad x > 0.$$

(The series converges by the estimates above.) It is easily verified that ψ is bounded, continuous (also at powers of 2), and periodic in $\log_2 x$, i.e., $\psi(2x) = \psi(x)$. Hence $\psi(2^y)$ has period 1 and may be expanded in a Fourier series

$$\psi(2^y) = \sum_{v=-\infty}^{\infty} c_v e^{2\pi i v y} \quad (18)$$

where

$$\begin{aligned}
 c_v &= \int_0^1 e^{-2\pi i v y} \psi(2^y) dy \\
 &= \int_0^1 e^{-2\pi i v y} \left(\sum_{l=-\infty}^{\infty} f(2^{y-l}) - y \right) dy \\
 &= \sum_{l=-\infty}^{\infty} \int_0^1 e^{-2\pi i v y} f(2^{y-l}) dy - \int_0^1 e^{-2\pi i v y} y dy \\
 &= \int_{-\infty}^{\infty} e^{-2\pi i v y} f(2^y) dy - \int_0^1 e^{-2\pi i v y} y dy.
 \end{aligned} \tag{19}$$

Further, changing variables back to $(0, \infty)$,

$$\begin{aligned}
 \int_{-\infty}^{\infty} e^{-2\pi i v y} f(2^y) dy &= \frac{1}{\ln 2} \int_0^{\infty} x^{-2\pi i v / \ln 2} f(x) \frac{dx}{x} \\
 &= \frac{1}{\ln 2} \mathcal{M}[f(x); -2\pi i v / \ln 2],
 \end{aligned} \tag{20}$$

where $\mathcal{M}[f(x); s] = \int_0^{\infty} x^{s-1} f(x) dx$ is the Mellin transform of f at the point s (in our case $s = -2\pi i v / \ln 2$). (See [28, Chapter 9] for definition and basic properties.) Since $|f(x)| < \min(x, x^{-1})$, the Mellin transform $\mathcal{M}[f(x); z]$ is analytic in the strip $-1 < \Re z < 1$. In the smaller strip $0 < \Re z < 1$ we have, by [28, Table 9.1],

$$\begin{aligned}
 \mathcal{M}[f(x); z] &= \mathcal{M}\left(\frac{e^{-x} - 1}{x}; z\right) + \mathcal{M}(\mathbf{1}[x < 1]; z) \\
 &= \mathcal{M}(e^{-x} - 1; z - 1) + \mathcal{M}(\mathbf{1}[x < 1]; z) \\
 &= \Gamma(z - 1) + \frac{1}{z}.
 \end{aligned} \tag{21}$$

By analyticity, this extends to the strip $-1 < \Re z < 1$. In particular, taking the limit as $z \rightarrow 0$,

$$\begin{aligned}
 \mathcal{M}[f(x); 0] &= \lim_{z \rightarrow 0} \left(\frac{\Gamma(1+z)}{z(z-1)} + \frac{1}{z} \right) = \lim_{z \rightarrow 0} \frac{\Gamma(1+z) - 1 + z}{z(z-1)} \\
 &= -(\Gamma'(1) + 1) = \gamma - 1.
 \end{aligned} \tag{22}$$

Moreover, elementary integration yields

$$\int_0^1 e^{-2\pi i v y} y dy = \begin{cases} -1/2\pi i v, & v \neq 0, \\ \frac{1}{2}, & v = 0. \end{cases} \tag{23}$$

By (19)–(23),

$$\begin{aligned}
 c_v &= \frac{1}{\ln 2} \Gamma\left(-1 - \frac{2\pi i v}{\ln 2}\right), \quad v \neq 0, \\
 c_0 &= \frac{\gamma - 1}{\ln 2} - \frac{1}{2}.
 \end{aligned}$$

The theorem now follows, with $\phi(x) = -(\psi(2^x) - c_0)$, from (6), (16), (17) and (18).

The numerical bounds for $k \leq 4$ are obtained from (3); for small k , $\sum_1^\infty |c_v|$ is dominated by $|c_1|$ which is very small.

Acknowledgements

This work was supported in part by NSF Grants CCR-0208709, DBI-0321756 and NIH grant R01 GM068959-01. Thanks to Jeffrey Shallit for comments on an earlier version of this paper.

References

- [2] A. Apostolico, M.E. Bock, S. Lonardi, Monotony of surprise and large-scale quest for unusual words, *J. Comput. Biol.* 10 (3–4) (2003) 283–311.
- [3] A. Apostolico, M.E. Bock, S. Lonardi, X. Xu, Efficient detection of unusual words, *J. Comput. Biol.* 7 (1/2) (2000) 71–94.
- [4] A. Apostolico, W. Szpankowski, Self-alignment in words and their applications, *J. Algorithms* 13 (3) (1992) 446–467.
- [5] A.D. Barbour, L. Holst, S. Janson, *Poisson Approximation*, Oxford Press, Oxford, 1992.
- [6] T.C. Bell, J.G. Cleary, I.H. Witten, *Text Compression*, Prentice-Hall, Englewood Cliffs, NJ, 1990.
- [7] C. Burge, A. Campbell, S. Karlin, Over- and under-representation of short oligonucleotides in DNA sequences, *Proc. Natl. Acad. Sci. USA.* 89 (1992) 1358–1362.
- [8] W.I. Chang, E.L. Lawler, Sublinear approximate string matching and biological applications, *Algorithmica* 12 (4/5) (1994) 327–344.
- [9] M.T. Chen, J. Seiferas, Efficient and elegant subword tree construction, in: A. Apostolico, Z. Galil (Eds.), *Combinatorial Algorithms on Words*, NATO Advanced Science Institutes, Series F, Vol. 12, Springer, Berlin, 1985, pp. 97–107.
- [10] A. De Luca, On the combinatorics of finite words, *Theoret. Comput. Sci.* 218 (1) (1999) 13–39.
- [11] M. Farach, Optimal suffix tree construction with large alphabets, in: *Proc. 38th Annual Symp. on Foundations of Computer Science*, 1997, pp. 137–143.
- [12] J.W. Fickett, D.C. Torney, D.R. Wolf, Base compositional structure of genomes, *Genomics* 13 (1992) 1056–1064.
- [13] A. Flaxman, A.W. Harrow, G.B. Sorkin, Strings with maximally many distinct subsequences and substrings, *Electron. J. Combin.* 11 1 (2004).
- [14] L. Ilie, S. Yu, K. Zhang, Repetition complexity of words, in: *Proc. Internat. Computing and Combinatorics Conference 2002*, pp. 320–329.
- [15] S. Karlin, C. Burge, A.M. Campbell, Statistical analyses of counts and distributions of restriction sites in DNA sequences, *Nucleic Acids Res.* 20 (1992) 1363–1370.
- [16] S. Karlin, J. Mrazek, A.M. Campbell, Frequent oligonucleotides and peptides of the *Haemophilus influenzae* genome, *Nucleic Acids Res.* 24 (21) (1996) 4273–4281.
- [17] Z. Kása, On the d -complexity of strings, *Pure Math. Appl.* 9 (1–2) (1998) 119–128.
- [18] F. Leve, P. Séébold, Proof of a conjecture on word complexity, *Bull. Belg. Math. Soc.* 8(2) (2001).
- [19] M. Li, P. Vitányi, *Introduction to Kolmogorov Complexity and its Applications*, Springer, Berlin, August 1993.
- [20] M. Lothaire (Ed.), *Algebraic Combinatorics on Words*, Cambridge University Press, Cambridge, 2002.
- [21] E.M. McCreight, A space-economical suffix tree construction algorithm, *J. Assoc. Comput. Mach.* 23 (2) (1976) 262–272.
- [22] H. Niederreiter, Some computable complexity measures for binary sequences, in: C. Ding, T. Hellseth, H. Niederreiter (Eds.), *Sequences and Their Applications*, Springer, Berlin, 1999, pp. 67–78.
- [23] P.A. Pevzner, M.Y. Borodovsky, A.A. Mironov, Linguistics of nucleotides sequences II: stationary words in genetic texts and the zonal structure of DNA, *J. Biomol. Struct. Dynamics* 6 (1989) 1027–1038.

- [25] G. Reinert, S. Schbath, M.S. Waterman, Probabilistic and statistical properties of words: an overview, *J. Comput. Biol.* 7 (2000) 1–46.
- [26] J. Shallit, On the maximum number of distinct factors in a binary string, *Graphs Combin.* 9 (1993) 197–200.
- [27] W. Szpankowski, A generalized suffix tree and its (un)expected asymptotic behaviors, *SIAM J. Comput.* 22 (1993) 1176–1198.
- [28] W. Szpankowski, *Average Case Analysis of Algorithms on Sequences*, Wiley, New York, 2001.
- [31] O.G. Troyanskaya, O. Arbell, Y. Koren, G.M. Landau, A. Bolshoy, Sequence complexity profiles of prokaryotic genomic sequences: a fast algorithm for calculating linguistic complexity, *Bioinformatics* 18 (5) (2002) 679–688.
- [32] E. Ukkonen, On-line construction of suffix trees, *Algorithmica* 14 (3) (1995) 249–260.
- [34] P. Weiner, Linear pattern matching algorithm, in: *Proc. 14th Annual IEEE Symp. on Switching and Automata Theory*, Washington, DC, 1973, pp. 1–11.