Sketching and Streaming for Dictionary Compression

Ruben Becker[*], Matteo Canton[†], Davide Cenzato[*],
Sung-Hwan Kim[*], Bojana Kodric[*], and Nicola Prezza[*]

[*]Ca' Foscari University of Venice	[†]University of Udine
Via Torino 155	Via delle Scienze 206
30172 Venezia, Italy	33100 Udine, Italy
`firstname.lastname@unive.it`	`canton.matteo@spes.uniud.it`

Abstract

We initiate the study of sub-linear sketching and streaming techniques for estimating the output size of common dictionary compressors such as Lempel-Ziv '77, the run-length Burrows-Wheeler transform, and grammar compression. To this end, we focus on a measure that has recently gained much attention in the information-theoretic community and which approximates up to a polylogarithmic multiplicative factor the output sizes of those compressors: the normalized substring complexity function δ. As a matter of fact, δ itself is a very accurate measure of compressibility: it is monotone under concatenation, invariant under reversals and alphabet permutations, sub-additive, and asymptotically tight (in terms of worst-case entropy) for representing strings, up to polylogarithmic factors.

We present a data sketch of $O(\varepsilon^{-3} \log n + \varepsilon^{-1} \log^2 n)$ words that allows computing a multiplicative $(1 \pm \varepsilon)$-approximation of δ with high probability, where n is the string length. The sketches of two strings S_1, S_2 can be merged in $O(\varepsilon^{-1} \log^2 n)$ time to yield the sketch of $\{S_1, S_2\}$, speeding up the computation of *Normalized Compression Distances* (NCD). If random access is available on the input, our sketch can be updated in $O(\varepsilon^{-1} \log^2 n)$ time for each character right-extension of the string. This yields a polylogarithmic-space algorithm for approximating δ, improving exponentially over the working space of the state-of-the-art algorithms running in nearly-linear time. Motivated by the fact that random access is not always available on the input data, we then present a streaming algorithm computing our sketch in $O(\sqrt{n} \cdot \log n)$ working space and $O(\varepsilon^{-1} \log^2 n)$ worst-case delay per character. We show that an implementation of our streaming algorithm can estimate δ on a dataset of 189GB with a throughput of 203MB per minute while using only 5MB of RAM, and that our sketch speeds up the computation of all-pairs NCD distances by one order of magnitude, with applications to phylogenetic tree reconstruction.

1 Introduction

Sketching techniques allow to summarize in sub-linear space information on big datasets, enabling the approximation of useful statistics such as high-order moments [1], norms [2], and frequencies [3] (to name a few). Additionally, most data sketches can be computed on data streams in sub-linear space, making them attractive in big data scenarios. In this paper, we consider data sketches summarizing the *information content* of a string as approximated by data compression techniques. Previous research on this problem has focused on *empirical entropy*. Chakrabarti et al. [4] showed that the zero-order empirical entropy H_0 of a data stream can be efficiently approximated up to a multiplicative $(1 + \varepsilon)$-factor in poly-logarithmic space, but

any multiplicative approximation of the k-th order entropy H_k requires nearly-linear space for $k \geq 1$. In addition to this fact, it is well-known that H_k is a weak measure when the dataset is highly repetitive [5]. As extensively shown in the literature (see, for example, the survey by Navarro [6]), dictionary compression measures such as the number z of phrases of the Lempel-Ziv'77 factorization (used by `winzip`, `7-zip`, `gzip`, `xz`), the number r of equal-letter runs in the Burrows-Wheeler transform (used by `bzip-2`), and the size g of a smallest context-free grammar generating (only) the text, are exempt from such a limitation. The information-theoretic quality of these measures is strengthened by the fact that *Normalized Compression Distances* based on dictionary compressors yield very precise notions of string similarity [7]. Sketching and streaming techniques for such measures would thus speed up tasks such as the computation of all-pairs similarities when the underlying metric is based on data compression (useful, for example, in the computation of phylogenetic trees [7]).

Motivated by the above considerations, in this paper we present the first sublinear-space sketching and streaming techniques for estimating the output sizes of dictionary compressors. This result is obtained by describing a data sketch yielding a $(1 \pm \varepsilon)$-approximation of the *normalized substring complexity* $\delta = \max_{k \geq 1}\{d_k/k\}$, where d_k is the number of distinct length-k substrings of the string, a measure introduced by Raskhodnikova et al. in [8]. As shown by Kociumaka et al. [9] and Kempa and Kociumaka [10], any of the above dictionary compression measures is lower-bounded by δ and upper-bounded by $\delta(\log n)^c$, where n is the string's length and c is an opportune constant depending on the compressor. Even better, Bonnie et al. in [11] experimentally showed that δ, z, and r (normalized to the interval $[0, 1]$) are almost indistinguishable on collections of genomic data. As a matter of fact, δ is known to be an even more accurate information measure than z, r, and g: it is monotone under string concatenation, invariant under reversals and alphabet permutations, sub-additive, and asymptotically tight (in terms of worst-case entropy) for representing strings, up to polylogarithmic factors [9]. None of the measures z, r, g possesses *simultaneously* all of these properties.

Overview of the paper. After providing all necessary definitions in Section 2, in Section 3 we prove new properties of the normalized substring complexity δ and of the Normalized Compression Distance [7] NCD_δ based on δ. In particular, we show that δ is perfectly sub-additive, that $\mathrm{NCD}_\delta(x, y)$ always lies in $[0, 1]$ (according to [12], this is an indicator that δ is a compressibility measure of good quality), and that $\mathrm{NCD}_{\tilde{\delta}}(x, y)$ is an additive $\Theta(\varepsilon)$-approximation of $\mathrm{NCD}_\delta(x, y)$ if $\tilde{\delta}$ is a multiplicative $(1 \pm \varepsilon)$-approximation of δ. This motivates designing data sketches for δ, a problem that we solve in Section 4. Our sketch is based on the observation (already noted in [13] for the particular case $\varepsilon = 1$) that $\max_{i \geq 0}\{d_{\lceil(1+\varepsilon)^i\rceil}/\lceil(1 + \varepsilon)^i\rceil\}$ is a $(1 - \Theta(\varepsilon))$-approximation of δ. We approximate d_k, for each sampled length $k = \lceil(1 + \varepsilon)^i\rceil$, by keeping a count-distinct sketch [14] for the subset of distinct (Rabin's fingerprints [15] of the) length-k substrings. Our sketch uses space polynomial in $\varepsilon^{-1} \log n$ and supports updates and queries (returning a $(1 \pm \varepsilon)$ approximation of δ), in $O(\varepsilon^{-1} \log^2 n)$ time. The sketches of two strings S_1 and S_2 can moreover be merged in $O(\varepsilon^{-1} \log^2 n)$ time to obtain the sketch of $\{S_1, S_2\}$, from which one can compute an additive ε-

approximation of $\mathrm{NCD}_\delta(S_1, S_2)$. In Section 5 we show how to compute our sketch in sub-linear space on an input stream of length n. The main difficulty in achieving sub-linear space is that, in order to compute the Rabin's fingerprints of the stream's length-k substrings, we need random access to the k-th most recent stream's character. Since the largest k for which we need to compute d_k is linear in n, storing the most recent k characters would require $\Theta(n)$ working space. Our solution relies on the observation that, if $\hat{k} = \mathrm{argmax}_{k \geq 1}\{d_k/k\}$ is small, then we can afford keeping a sliding window of the last $\hat{k}$ stream's characters. If, on the other hand, $\hat{k}$ is large, then the stream is highly repetitive so we can compress it in small space while supporting bookmarked access to its characters. We conclude in Section 6 with experimental results. Complete proofs can be found in the full version [16].

Related work. Bonnie et al. [11] have already observed that d_k can be efficiently estimated by employing count-distinct sketches, and that this can yield an heuristic algorithm for estimating δ. Their strategy relies on estimating d_k/k for increasing values of k, until a local maximum is found. While this strategy works well in practice because, as they showed, $\hat{k} = \mathrm{argmax}_{k \geq 1}\{d_k/k\}$ tends to be a very small number, on particular strings (for example, Thue-Morse) $\hat{k}$ is of the order of $\Theta(n)$ and, as a result, computing all the sketches for d_k requires linear space and quadratic processing time in the worst case. Moreover, local maxima of d_k/k do not always coincide with the global maximum, so this strategy does not yield any provable approximation of δ. We are not aware of other works in the literature describing data sketches for estimating the output sizes of *dictionary* compressors (the literature on estimating empirical entropy is, on the other end, much richer: see [4] and references therein). Our results can be viewed also as a space-efficient way to approximate measure δ. Christiansen et al. [17] showed how to compute δ for a given string T in linear time and space. Recently, Bernardini et al. [13] provided space-time trade-offs for computing/approximating δ in sub-linear working space on top of the input string. If $O(n \, \mathrm{polylog}\, n)$ time is allowed, their algorithms require $\Theta(n/\, \mathrm{polylog}\, n)$ working space, which they proved to be optimal for computing δ exactly. Our algorithm, on the other hand, computes a multiplicative $(1 \pm \varepsilon)$-approximation of δ using working space polynomial in $\varepsilon^{-1} \log n$.

2 Preliminaries

We denote $[n] := \{1, \ldots, n\}$ for any integer n ($[n] = \emptyset$ for $n \leq 0$). For $a \in \mathbb{R}^+$ and a real number $\varepsilon \in [0, 1]$, we write $[(1 \pm \varepsilon)a]$ for the interval $[(1 - \varepsilon)a, (1 + \varepsilon)a]$. Similarly, we write $[a \pm \varepsilon]$ for the interval $[a - \varepsilon, a + \varepsilon]$.

We assume to be given a string S of length $n > 1$ over an alphabet Σ of cardinality $\sigma > 1$. For $k \geq 1$, we define $D_k(S) := \{S[i..i + k - 1] : i \in [n - k + 1]\}$, i.e., the set of all distinct substrings of length k of S. Notice that $D_k(S) = \emptyset$ if $k > n$. The *k-substring complexity* $d_k(S)$ of S is the cardinality of this set, i.e., $d_k(S) := |D_k(S)|$. The *normalized substring complexity* δ is defined as follows:

$$\delta(S) := \max_{k \geq 1}\{|D_k(S)|/k\} = \max_{k \geq 1}\{d_k(S)/k\}.$$

3

We omit the argument from D_k, d_k, and δ in case it is clear from the context. Here, we also extend this measure to pairs of strings S and T. Rather than using $\delta(ST)$, we propose the following natural definition that does not take into account artificial length-k substrings crossing the border between S and T:

$$\delta(S,T) := \max_{k \geq 1}\{|D_k(S) \cup D_k(T)|/k\}.$$

This version also gives mathematically cleaner results (e.g., perfect sub-additivity) and, in any case, differs from $\delta(ST)$ by at most 1. As a consequence, most of our results (read also below) hold also by replacing $\delta(S,T)$ with $\delta(ST)$.

The *Normalized Compression Distance* has been defined by Cilibrasi and Vitányi [7] as a proxy for the non-computable Normalized Information Distance [18]. For two strings S and T and an arbitrary compressibility measure Z (for example, the output size of compression software such as `gzip` and `xz`), it is defined as

$$\mathrm{NCD}_Z(S,T) := \frac{Z(S,T) - \min\{Z(S), Z(T)\}}{\max\{Z(S), Z(T)\}}.$$

Given a uniform prime $q = n^{\Theta(1)}$, the *Rabin's fingerprint* [15] of S is defined as $\rho(S) = \sum_{i=1}^{n} S[i] \cdot \sigma^{n-i} \mod q$. Collisions between substrings of S through ρ happen with low probability, so the results of our paper hold with high probability. We extensively use the fact that the fingerprint of the concatenation of two strings S_1, S_2 can be computed in constant time from (i) the fingerprints of S_1 and S_2 and (ii) $\sigma^{|S_2|} \mod q$ (see [15]). Given a set $B \subseteq \Sigma^*$ of strings, we define $\rho(B) = \{\rho(s) \; : \; s \in B\}$.

Given a set U, a *count-distinct sketch* $CD(U)$ is a sub-linear-space data structure supporting three main operations: $CD(U).add(x)$, which turns the sketch into $CD(U \cup \{x\})$, $CD(U_1).merge(CD(U_2))$, which turns the sketch into $CD(U_1 \cup U_2)$, and $CD(U).estimate()$, which returns a $(1 \pm \varepsilon)$ approximation of $|U|$. In our work, we use the optimal count-distinct sketch of Kane et al. [14]. Letting $U \subseteq [u]$, this sketch uses $O(\varepsilon^{-2} + \log u)$ words of space and computes a $(1 \pm \varepsilon)$ approximation of $|U|$ with high probability of success. All operations are supported in $O(\log u)$ time[1].

Assume $S[1] = \$$, where $\$$ is lexicographically smaller than all other alphabet's characters and does not appear anywhere else in S. The *Burrows-Wheeler transform* (BWT) of the reverse S^R of S is obtained by sorting lexicographically all suffixes of S^R and then taking, in this order, the character preceding each suffix. For example, if $S = \$babba$, then the sorted suffixes and the BWT of S^R are shown in Table 1.

The *LF property* of the BWT states that the i-th occurrence of $c \in \Sigma$ in the BWT corresponds to the position of the i-th suffix starting with $c \in \Sigma$ in Table 1. The *LF function* is the permutation of $[1, n]$ implementing this observation: for instance, in the above example $BWT.LF(2) = 5$ because character $BWT[2]$ corresponds to the first character (b) of the fifth (in lexicographic order) suffix $bab\$$. We denote with r the number of *equal-letter runs* of the BWT; in the above example, $r = 5$ (runs are highlighted in alternating bold/italic). We moreover use the following result:

---

[1] The authors claim $O(\varepsilon^{-2} + \log u)$ *bits* of space and $2/3$ success probability, which can be amplified by taking the median of $\Theta(\log u)$ sketches (thus yielding the bounds we claim above). In our paper, the universe is composed by Rabin's fingerprints and has therefore size $u = n^{\Theta(1)}$.

suffixes of S^R	BWT
$	**b**
a b $	**b**
a b b a b $	*$*
b $	**a**
b a b $	*b*
b b a b $	**a**

Table 1: Burrows-Wheeler transform $BWT(S^R)$ of the string $S^R = abbab\$$.

Lemma 2.1 ([19], Thm. 2). *Letting S be a string and r be the number of equal-letter runs in $BWT(S^R)$, there exists a data structure of $O(r)$ words storing $BWT(S^R)$ supporting right-extensions of S (i.e. $BWT(S^R) \to BWT((Sa)^R)$, for any $a \in \Sigma$) in $O(\log|S|)$ time. Within the same time, the structure supports computing the LF function and retrieving any character of $BWT(S^R)$.*

3 Properties of δ and NCD_δ

We start by proving some properties of δ. Proofs of some statements are omitted due to space limitations and can be found in the full version [16]. The first main property that we show is that δ is both sub-additive and monotone in the following sense.

Lemma 3.1. *For any strings S and T, $\max\{\delta(S), \delta(T)\} \leq \delta(S,T) \leq \delta(S) + \delta(T)$.*

The proof of the lemma uses the properties of the corresponding maximizers together with the fact that the union is a superset of both its arguments (left inequality) and that the union is of smaller cardinality than the sum of the cardinalities of its arguments (right inequality). We remark that it is a well-known fact that the monotonicity property holds for the case of concatenation of the two strings [9]. Using the sub-additivity of δ, we obtain:

Corollary 3.2. *For any strings S and T it holds that $0 \leq \mathrm{NCD}_\delta(S,T) \leq 1$.*

To see why this holds, assume, w.l.o.g., that $\max\{\delta(S), \delta(T)\} = \delta(S)$. Then, $\mathrm{NCD}_\delta(S,T) = \frac{\delta(S,T)-\delta(T)}{\delta(S)} \geq \frac{\delta(T)-\delta(T)}{\delta(S)} = 0$ and $\mathrm{NCD}_\delta(S,T) \leq \frac{\delta(S)+\delta(T)-\delta(T)}{\delta(S)} = 1$. Ming et al. [12] state that common compressors yield a normalized compression distance between 0 and $1 + \varepsilon$, where the ε is due to "imperfections" of the compression algorithm. Above we proved that in the case of the normalized substring complexity δ, the corresponding ε is equal to 0.

We conclude by showing that a multiplicative approximation of δ can be used to obtain an additive approximation of the Normalized Compression Distance NCD_δ.

Lemma 3.3. *Let $\varepsilon \in (0,1)$, $\varepsilon' := \varepsilon/5$ and let S and T be two strings. Assume that $\tilde{\delta}(S)$, $\tilde{\delta}(T)$, and $\tilde{\delta}(S,T)$ are approximations of δ in the sense that $\tilde{\delta}(S) \in [(1\pm\varepsilon')\delta(S)]$, $\tilde{\delta}(T) \in [(1\pm\varepsilon')\delta(T)]$, as well as $\tilde{\delta}(S,T) \in [(1\pm\varepsilon')\delta(S,T)]$. Then*

$$\mathrm{NCD}_{\tilde{\delta}}(S,T) \in [\mathrm{NCD}_\delta(S,T) \pm \varepsilon].$$

We prove this lemma by using the facts that $\tilde{\delta}$ is a multiplicative approximation of δ, that δ is sub-additive (see Lemma 3.1), and that $\text{NCD}_\delta \in [0,1]$ (see Corollary 3.2).

4 A data sketch for estimating δ

We introduce our data sketch, then prove that it yields a good approximation of δ.

Definition 4.1 (Sketch for δ). *Let S be a string, $A := \{\lceil \alpha^i \rceil : i \in [\lfloor \log_\alpha n \rfloor]\}$ be a set of sampled lengths for some real number (sample rate) $\alpha > 1$, and $CD_k = CD(\rho(D_k(S)))$, where CD is the count-distinct sketch described in Section 2 and ρ is Rabin's hash function. Our data sketch is defined as $\kappa(S) = \langle CD_k : k \in A \rangle$.*

We define $\kappa(S).estimate() = \max\{CD_k.estimate()/k : k \in A\}$. When extending the stream S with a new character a, yielding string Sa, the sketch is updated by calling $CD_k.add(\rho(S[|S| - k + 2, |S|]a))$ for all $k \in A$. We denote this operation by $\kappa(S).extend(a)$. Note that, if constant-time random access is available on S and if $\sigma^{k-1} \mod q$ has been pre-computed for all $k \in A$ (in $O(\varepsilon^{-1} \log^2 n)$ time), $\rho(S[|S| - k + 2, |S|]a)$ can be computed in constant time from $\rho(S[|S| - k + 1, |S|])$; see [15]. Finally, $\kappa(S_1).merge(\kappa(S_2))$ returns the sketch $\kappa(\{S_1, S_2\}) = \langle CD'_k : k \in A \rangle$, where $CD'_k = CD^1_k.merge(CD^2_k)$ and CD^i_k is the count-distinct sketch for the (fingerprints of the) length-k substrings of S_i, $i \in \{1, 2\}$. Operation $extend(a)$ is not defined when the sketch represents a set of strings; this is not an issue, since we will call $merge$ only to estimate $\delta(S_1, S_2)$ and $NCD_\delta(S_1, S_2)$.

With the next theorem we show that $\kappa(S).estimate()$ returns a multiplicative $(1 \pm \varepsilon)$-approximation of $\delta(S)$ (analogous for $\kappa(\{S_1, S_2\}).estimate()$).

Lemma 4.2. *Let S be a string of length n. Let $\varepsilon > 0$, $\varepsilon' = \varepsilon/4$, and $\alpha = 1 + \varepsilon'$. Assume that $\tilde{d}_k \in [(1 \pm \varepsilon')d_k(S)]$ for all $k \in A := \{\lceil \alpha^i \rceil : i \in [\lfloor \log_\alpha n \rfloor]\}$, then $\tilde{\delta} := \max\{\tilde{d}_k/k : k \in A\}$, satisfies $\tilde{\delta} \in [(1 \pm \varepsilon)\delta]$.*

We show this theorem by quantifying the impact of two types of errors on the quantity δ. These two types are (1) the error obtained when approximating the values d_k by $\tilde{d}_k$ and (2) the error due to the restriction of the string's offsets $[n]$ to the set A. The error of type (1) directly implies an error of the same magnitude $(1 \pm \varepsilon')$ on δ. We note that this error actually itself has two sources, namely (1.1) errors due to collisions when computing Rabin's fingerprints and (1.2) errors due to the count-distinct sketch when applied to the fingerprints. Both of these errors are accounted for in the assumption $\tilde{d}_k \in [(1 \pm \varepsilon')d_k(S)]$. The error of type (2) instead is more subtle to analyse – the main observation here is that $d_{j+1} \geq d_j - 1$ for every j, as every distinct length-j substring other than possibly $S[n - j + 1..n]$ gives at least one distinct length-$(j + 1)$ substring. Now assume that $i \in [n] \setminus A$ and that $a \in A$ is the minimum element of A larger than i. Then applying the previous observation iteratively yields $d_a \geq d_i - \beta$, where $\beta = a - i$. Hence, we can quantify how much δ gets "perturbed" by restricting to the subset A of the string's offsets $[n]$.

From Lemmas 3.3 and 4.2, the sketch of Definition 4.1 yields a multiplicative $(1 \pm \varepsilon)$-approximation of δ and an additive ε-approximation of NCD_δ if CD is the count-distinct sketch of [14] with error rate $\varepsilon/20$, and the set A is built with sample rate $\alpha =$

$1 + \varepsilon/20$. From [14] and since $|A| \in \Theta(\varepsilon^{-1} \log n)$, our data sketch uses $\Theta(\varepsilon^{-3} \log n + \varepsilon^{-1} \log^2 n)$ words of space and supports all operations in time $O(\varepsilon^{-1} \log^2 n)$.

Using repeatedly operation *extend* on our data sketch we immediately obtain:

Theorem 4.3. *For any string S of length n supporting random access in time at most $O(\log n)$, and any approximation rate $\varepsilon > 0$, we can compute a multiplicative $(1 \pm \varepsilon)$-approximation of $\delta(S)$ in $O(\varepsilon^{-1} n \log^2 n)$ time using $\Theta(\varepsilon^{-3} \log n + \varepsilon^{-1} \log^2 n)$ words of working space on top of the string. The result is correct with high probability.*

5 Streaming algorithm

We now show how to compute the sketch of Definition 4.1 in $O(\sqrt{n} \log n)$ words of working space (on top of the sketch) with one pass over the streamed input string.

Let S denote the current stream, and S^R be the reversed stream. We assume that an upper-bound n to the maximum stream length is known before the algorithm starts. Let r be the number of equal-letter runs in the Burrows-Wheeler transform of S^R. By [10] and by the fact that δ is invariant under string reversals, it holds $r \leq 8\delta \log^2 n$. Our streaming algorithm works as follows. We keep a sliding window $S[|S| - K + 1, |S|]$ of the last K stream characters, for some parameter K to be determined later, and at the same time we keep a dynamic run-length BWT (RLBWT) of S^R that we update by appending the stream's characters using Lemma 2.1. Before the algorithm starts, in $O(\varepsilon^{-1} \log^2 n)$ time we compute $\sigma^{k-1} \mod q$ for all the $|A| \in O(\varepsilon^{-1} \log n)$ sampled substring lengths $k \in A$ in our sketch, using fast exponentiation.

Let $k \in A$ be one of the sampled string lengths in our sketch, and let a be a new character arriving on the stream (so that the new stream is Sa). In order to update our sketch, we need to compute the fingerprint of the last k stream's characters: $\rho(S[|S| - k + 2, |S|]a)$. At any stage of the algorithm, we keep the Rabin's fingerprint $\rho(S[2, |S|])$ of the whole stream, excluding character $S[1] = \$$. If $|Sa| = k + 1$, then $\rho(S[|S| - k + 2, |S|]a)$ is equal to the Rabin's fingerprint of the whole stream. Otherwise, if $|Sa| > k + 1$ then $\rho(S[|S| - k + 1, |S|])$ has already been computed in the previous steps and we can use the formula $\rho(S[|S| - k + 2, |S|]a) = (\rho(S[|S| - k + 1, |S|]) - S[|S| - k + 1] \cdot \sigma^{k-1}) \cdot \sigma + a \mod q$. As a result, updating the fingerprint reduces to extracting character $S[|S| - k + 1]$. We use the window $S[|S| - K + 1, |S|]$ to extract $S[|S| - k + 1]$ for any $k \leq K$, and the RLBWT to extract $S[|S| - k + 1]$ for any $k > K$ using a bookmarking technique that we sketch in Figure 1 and we describe in full detail in the full version [16]. This allows us to update the Rabin's fingerprints for all sampled substring lengths k and thus to implement operation *extend(a)*.

We now describe the policy we employ to keep space usage under control. Let r' be the number of equal-letter runs in the BWT obtained by ignoring (removing) character $\$$. It is easy to see that (i) $r - 2 \leq r' \leq r$ and (ii) r' is non-decreasing upon appending characters at the end of the stream. As soon as $r' \geq 8n(\log^2 n)/K$, we discard the RLBWT and keep only the sliding window for the rest of the stream. As a consequence, from this point on we are only able to extract (fingerprints of) length-k substrings with $k \leq K$. However, we show that this is enough: if we discard the RLBWT, then it means that $\delta \geq \frac{r}{8 \log^2 n} \geq \frac{r'}{8 \log^2 n} \geq n/K$. Let $\hat{k} = \text{argmax}_{k \geq 1}\{d_k/k\}$.

suffixes of S^R	BWT
$	$

suffixes of S^R	BWT
$	b
b$	$

suffixes of S^R	BWT	
$	b	$j = 1$
ab$	$	
b$	a	

suffixes of S^R	BWT	
$	b	
aab$	$	
ab$	a	
b$	a	$j = 4$

Figure 1: Example showing how the BWT (of the reversed stream) is updated upon character right-extensions of the stream, and how the bookmark j corresponding to window length $k = 2$ is initialized and updated. **Top left**: empty stream ($S = \$$). **Top right**: a new character b arrives on the stream ($S = \$b$): in the BWT, $ is replaced by b and a new $ is inserted in the position corresponding to the lexicographic rank $i = 2$ of the new suffix $b\$$. Position i is computed in $O(\log |S|)$ time using the algorithm described in [19, Thm. 2]. **Bottom left**: a new character a arrives on the stream ($S = \$ba$): in the BWT, $ is replaced by a and a new $ is inserted in the position corresponding to the lexicographic rank $i = 2$ of the new suffix $ab\$$. Since the stream length is equal to $k + 1 = 3$, we initialize the bookmark $j \leftarrow BWT.LF(i) = BWT.LF(2) = 1$. Note that $BWT[j] = b$ indeed contains character $S[|S| - k + 1] = b$. **Bottom right**: a new character a arrives on the stream ($S = \$baa$): in the BWT, $ is replaced by a and a new $ is inserted in the position corresponding to the lexicographic rank $i = 2$ of the new suffix $aab\$$. Since $1 = j < i = 2$ ($ is inserted after position j), $j = 1$ is not modified (otherwise, it would have been incremented by 1). Finally, we update j by advancing it by one position in the text: $j \leftarrow BWT.LF(j) = BWT.LF(1) = 4$. Note that $BWT[j] = a$ indeed contains character $S[|S| - k + 1] = a$. Importantly, the data structure of [19, Thm. 2] uses always a space proportional to the number r of equal-letter runs of the BWT.

Then, $\hat{k} = d_{\hat{k}}/\delta \leq n/\delta \leq K$ so to compute δ on the rest of the stream we can focus only on the length-k substrings with $k \leq K$.

The sliding window $S[|S| - K + 1, |S|]$ uses K words of space. We discard the RLBWT when $r' \geq 8n(\log^2 n)/K$, so (since $r \leq r' + 2$) this structure always uses at most $O(r) \subseteq O(n(\log^2 n)/K)$ words. As a consequence, in total we use $O(K + n(\log^2 n)/K)$ words of space, which is optimized asymptotically when $K = \sqrt{n} \log n$; then, our algorithm uses at most $O(\sqrt{n} \log n)$ words of space.

We keep one bookmark (a position in the BWT) for every sampled length $k \in A$, so our bookmarking technique does not affect the asymptotic working space if $\varepsilon \geq n^{-1/2}$ (i.e. $|A| \leq \sqrt{n} \log n$). Updating each bookmark and extracting $S[|S| - k + 1]$ from the RLBWT take $O(\log n)$ time by Lemma 2.1. This running time is absorbed by operation $merge()$ on the count-distinct sketches, see Section 4. We obtain:

Theorem 5.1. *Given an upper-bound n to the stream's length, we can compute the sketch of Definition 4.1 in $O(\sqrt{n} \log n)$ words of working space and $O(\varepsilon^{-1} \log^2 n)$ worst-case delay per stream character, for any approximation factor $\varepsilon \geq n^{-1/2}$.*

6 Implementation and experiments

We implemented a parallel version of our streaming algorithm in C++.[2] We ran experiments on a server with Intel(R) Xeon(R) W-2245 CPU @ 3.90GHz with 16 threads and 128GB of RAM running Ubuntu 18.04 LTS 64-bit. Our complete experimental

---

[2] https://github.com/regindex/substring-complexity

results are reported in [16]. We used the repetitive real Pizza&Chilli dataset (P&C)[3], large Canterbury corpus[4], and datasets from AF Project[5]. We computed the relative error of our approximation $\tilde{\delta}$ with respect to δ for different sampling densities (i.e. parameter α of Definition 4.1). With the sparsest (less precise) sampling scheme (option `-p 1`), $\tilde{\delta}$ always differed from δ by up to 5% and the average throughput was of 174 MB per minute using up to 16 threads (option `-t 0`). For efficiency reasons, the RLBWT is disabled by default: in practice this does not affect precision, since $\hat{k} = \mathrm{argmax}_k \, d_k/k$ was always extremely small ($\hat{k} \leq 100$ in all datasets), meaning that the RLBWT is never required. We also computed $\tilde{\delta}$ on a big dataset of 189GB long reads of Rana Muscosa[6]. Our software finished the computation in 15:31 hours with a throughput of 203MB per minute using only about 5MB of internal memory.

Experiments on repetitiveness measures. We studied the effectiveness of $\tilde{\delta}$ as a repetitive measure. We compared it to exact δ, to the number of runs of the BWT r, to the number of phrases of the LZ77 parse z, and to the output of two popular compressors, `xz` and `7z`. For each dataset in the repetitive P&C corpus, we computed these five measures for prefixes of increasing length. We observe that $\tilde{\delta}$ not only follows closely the values of δ, but it also mirrors the trend of the other four measures. This suggests experimentally that $\tilde{\delta}$ computed by our streaming algorithm is a good indicator of repetitiveness and compressibility.
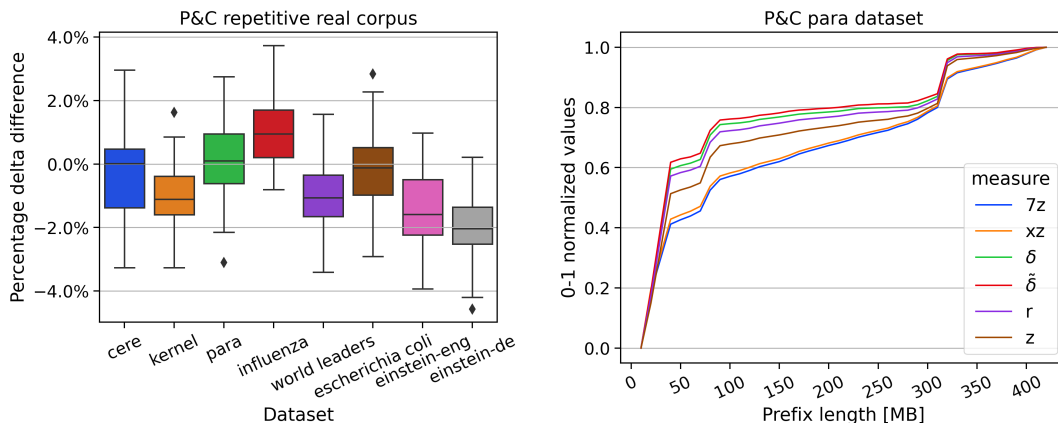

Figure 2: From left to right: $\tilde{\delta}$ error distribution on P&C repetitive corpus, lineplot showing five repetitiveness measures (normalized to $[0, 1]$) computed on increasing prefixes of *para*.

Experiments on phylogeny reconstruction. We verified that NCD based on the compression software `xz`, on δ, and on $\tilde{\delta}$ yield similar phylogenetic trees with the `Gene-trees` dataset from AF Project; the average normalized Robinson-Foulds distances ranged from 0.1 to 0.3, indicating that the reconstructed trees were very similar. We also measured the running time to compute all-pair NCDs on 29 sequences of average length $\sim$81k. This process took only 3 minutes for $\tilde{\delta}$ and 24 minutes for the exact δ, while for `xz` it required 42 minutes.

---

[3]`https://pizzachili.dcc.uchile.cl/repcorpus/real/`
[4]`http://corpus.canterbury.ac.nz/resources/large.tar.gz`
[5]`https://afproject.org`
[6]`https://trace.ncbi.nlm.nih.gov/Traces/?view=run_browser&acc=SRR11606868`

References

[1] Noga Alon, Yossi Matias, and Mario Szegedy, "The space complexity of approximating the frequency moments," in *Proc. of 28th ACM STOC*, 1996, pp. 20–29.

[2] William B. Johnson and Joram Lindenstrauss, "Extensions of Lipschitz mappings into Hilbert space," *Contemporary mathematics*, vol. 26, pp. 189–206, 1984.

[3] Jayadev Misra and David Gries, "Finding Repeated Elements," *Sci. Comput. Program.*, vol. 2, no. 2, pp. 143–152, 1982.

[4] Amit Chakrabarti, Graham Cormode, and Andrew Mcgregor, "A Near-Optimal Algorithm for Estimating the Entropy of a Stream," *ACM Trans. Algorithms*, vol. 6, no. 3, pp. 51:1–21, 2010.

[5] Sebastian Kreft and Gonzalo Navarro, "On compressing and indexing repetitive sequences," *Theoretical Computer Science*, vol. 483, pp. 115–133, 2013.

[6] Gonzalo Navarro, "Indexing Highly Repetitive String Collections, Part I: Repetitiveness Measures," *ACM Comput. Surv.*, vol. 54, no. 2, pp. 29:1–29:31, 2022.

[7] Rudi Cilibrasi and Paul MB Vitányi, "Clustering by compression," *IEEE Transactions on Information theory*, vol. 51, no. 4, pp. 1523–1545, 2005.

[8] Sofya Raskhodnikova, Dana Ron, Ronitt Rubinfeld, and Adam D. Smith, "Sublinear Algorithms for Approximating String Compressibility," *Algorithmica*, vol. 65, no. 3, pp. 685–709, 2013.

[9] Tomasz Kociumaka, Gonzalo Navarro, and Nicola Prezza, "Toward a Definitive Compressibility Measure for Repetitive Sequences," *IEEE Trans. Inf. Theory*, vol. 69, no. 4, pp. 2074–2092, 2023.

[10] Dominik Kempa and Tomasz Kociumaka, "Resolution of the Burrows-Wheeler Transform Conjecture," in *Proc of 61st IEEE FOCS*, 2020, pp. 1002–1013.

[11] Jessica K Bonnie, Omar Ahmed, and Ben Langmead, "DandD: efficient measurement of sequence growth and similarity," *bioRxiv*, pp. 2023–02, 2023.

[12] Ming Li, Xin Chen, Xin Li, Bin Ma, and P.M.B. Vitanyi, "The similarity metric," *IEEE Transactions on Information Theory*, vol. 50, no. 12, pp. 3250–3264, 2004.

[13] Giulia Bernardini, Gabriele Fici, Pawel Gawrychowski, and Solon P. Pissis, "Substring Complexity in Sublinear Space," *CoRR*, vol. abs/2007.08357, 2020.

[14] Daniel M. Kane, Jelani Nelson, and David P. Woodruff, "An Optimal Algorithm for the Distinct Elements Problem," in *Proc. of 29th ACM PODS*, 2010, p. 41–52.

[15] Michael O Rabin, "Fingerprinting by random polynomials," *Technical report*, 1981.

[16] Ruben Becker, Matteo Canton, Davide Cenzato, Sung-Hwan Kim, Bojana Kodric, and Nicola Prezza, "Sketching and Streaming for Dictionary Compression," 2023, https://arxiv.org/abs/2310.17980.

[17] Anders Roy Christiansen, Mikko Berggren Ettienne, Tomasz Kociumaka, Gonzalo Navarro, and Nicola Prezza, "Optimal-Time Dictionary-Compressed Indexes," *ACM Trans. Algorithms*, vol. 17, no. 1, pp. 8:1–8:39, 2021.

[18] Sebastiaan A. Terwijn, Leen Torenvliet, and Paul M. B. Vitányi, "Nonapproximability of the normalized information distance," *J. Comput. Syst. Sci.*, vol. 77, no. 4, pp. 738–742, 2011.

[19] Alberto Policriti and Nicola Prezza, "LZ77 computation based on the run-length encoded BWT," *Algorithmica*, vol. 80, no. 7, pp. 1986–2011, 2018.

A   Missing Proofs

We start with the following two easy observations that we use in our proofs later on.

Observation A.1. *If $S \neq a^n$ for $a \in \Sigma$, then $\delta(S) \geq d_1/1 \geq 2/1 = 2$.*

Note that if at least two distinct letters appear, it follows that $d_1 > 1$ and thus $\delta \geq d_1/1 \geq 2$. Since it is easy to recognize the case $S = a^n$ for some $a \in \Sigma$ in constant space and constant delay per character, from now on we assume w.l.o.g. that $\delta \geq 2$.

We continue with the following simple observation that we will use in the proof of Lemma 4.2.

Observation A.2. *It holds that $\hat{k} = \operatorname{argmax}_{k \geq 1}\{d_k(S)/k\} \leq n/2$.*

Proof. Assume that $\hat{k} > n/2$. Then it is immediate that $d_{\hat{k}} \leq n - \hat{k} + 1 < n/2$ (as the right most character in the substring can be at index at most n). We now obtain that $d_{\hat{k}}/\hat{k} < 1$, contradicting Observation A.1. $\qquad\square$

Lemma 3.1. *For any strings S and T, $\max\{\delta(S), \delta(T)\} \leq \delta(S,T) \leq \delta(S) + \delta(T)$.*

Proof. Let $k_{S,T}$, k_S, and k_T be such that $\delta(S,T) = |D_{k_{S,T}}(S) \cup D_{k_{S,T}}(T)|/k_{S,T}$, $\delta(S) = |D_{k_S}(S)|/k_S$, and $\delta(T) = |D_{k_T}(T)|/k_T$. Let, w.l.o.g., $\delta(S) = \max\{\delta(S), \delta(T)\}$. Then,

$$\delta(S,T) = \frac{|D_{k_{S,T}}(S) \cup D_{k_{S,T}}(T)|}{k_{S,T}} \geq \frac{|D_{k_S}(S) \cup D_{k_S}(T)|}{k_S} \geq \frac{|D_{k_S}(S)|}{k_S} = \delta(S),$$

where the second inequality uses the fact that k_S is the maximizer for S. For the second claim,

$$\begin{aligned} \delta(S,T) = \frac{|D_{k_{S,T}}(S) \cup D_{k_{S,T}}(T)|}{k_{S,T}} &\leq \frac{|D_{k_{S,T}}(S)|}{k_{S,T}} + \frac{|D_{k_{S,T}}(T)|}{k_{S,T}} \\ &\leq \frac{|D_{k_S}(S)|}{k_S} + \frac{|D_{k_T}(T)|}{k_T} = \delta(S) + \delta(T), \end{aligned}$$

where the second inequality uses the fact that k_S and k_T are the respective maximizers for S and T. $\qquad\square$

Lemma 3.3. *Let $\varepsilon \in (0,1)$, $\varepsilon' := \varepsilon/5$ and let S and T be two strings. Assume that $\tilde{\delta}(S)$, $\tilde{\delta}(T)$, and $\tilde{\delta}(S,T)$ are approximations of δ in the sense that $\tilde{\delta}(S) \in [(1 \pm \varepsilon')\delta(S)]$, $\tilde{\delta}(T) \in [(1 \pm \varepsilon')\delta(T)]$, as well as $\tilde{\delta}(S,T) \in [(1 \pm \varepsilon')\delta(S,T)]$. Then*

$$\mathrm{NCD}_{\tilde{\delta}}(S,T) \in [\mathrm{NCD}_\delta(S,T) \pm \varepsilon].$$

11

Proof. We start with the lower bound. Using the definition of $\mathrm{NCD}_{\tilde{\delta}}(S, T)$, we obtain

$$\mathrm{NCD}_{\tilde{\delta}}(S, T) \geq \frac{(1 - \varepsilon') \cdot \delta(S, T) - (1 + \varepsilon') \min\{\delta(S), \delta(T)\}}{(1 + \varepsilon') \max\{\delta(S), \delta(T)\}}$$

$$= \frac{1}{1 + \varepsilon'} \cdot \mathrm{NCD}_{\delta}(S, T) - \frac{\varepsilon'}{1 + \varepsilon'} \cdot \frac{\delta(S, T) + \min\{\delta(S), \delta(T)\}}{\max\{\delta(S), \delta(T)\}}$$

$$= \mathrm{NCD}_{\delta}(S, T) - \frac{\varepsilon'}{1 + \varepsilon'} \cdot \left(\frac{\delta(S, T) + \min\{\delta(S), \delta(T)\}}{\max\{\delta(S), \delta(T)\}} + \mathrm{NCD}_{\delta}(S, T) \right)$$

$$\geq \mathrm{NCD}_{\delta}(S, T) - \frac{4\varepsilon'}{1 + \varepsilon'}$$

$$\geq \mathrm{NCD}_{\delta}(S, T) - \varepsilon$$

using the sub-additivity of δ from Lemma 3.1, the fact that $\mathrm{NCD}_{\delta}(S, T) \leq 1$ according to Corollary 3.2, and the definition of ε'. Similarly, now for the upper bound, we obtain

$$\mathrm{NCD}_{\tilde{\delta}}(S, T) \leq \frac{(1 + \varepsilon') \cdot \delta(S, T) - (1 - \varepsilon') \min\{\delta(S), \delta(T)\}}{(1 - \varepsilon') \max\{\delta(S), \delta(T)\}}$$

$$= \frac{1}{1 - \varepsilon'} \cdot \mathrm{NCD}_{\delta}(S, T) + \frac{\varepsilon'}{1 - \varepsilon'} \cdot \frac{\delta(S, T) + \min\{\delta(S), \delta(T)\}}{\max\{\delta(S), \delta(T)\}}$$

$$= \mathrm{NCD}_{\delta}(S, T) + \frac{\varepsilon'}{1 - \varepsilon'} \cdot \left(\frac{\delta(S, T) + \min\{\delta(S), \delta(T)\}}{\max\{\delta(S), \delta(T)\}} + \mathrm{NCD}_{\delta}(S, T) \right)$$

$$\leq \mathrm{NCD}_{\delta}(S, T) + \frac{4\varepsilon'}{1 - \varepsilon'}$$

$$\leq \mathrm{NCD}_{\delta}(S, T) + \varepsilon$$

again using the sub-additivity of δ from Lemma 3.1, the fact that $\mathrm{NCD}_{\delta}(S, T) \leq 1$ according to Corollary 3.2, the definition of ε' and the assumption that $\varepsilon < 1$. $\square$

Lemma 4.2. *Let S be a string of length n. Let $\varepsilon > 0$, $\varepsilon' = \varepsilon/4$, and $\alpha = 1 + \varepsilon'$. Assume that $\tilde{d}_k \in [(1 \pm \varepsilon')d_k(S)]$ for all $k \in A := \{\lceil \alpha^i \rceil : i \in [\lfloor \log_\alpha n \rfloor]\}$, then $\tilde{\delta} := \max\{\tilde{d}_k/k : k \in A\}$, satisfies $\tilde{\delta} \in [(1 \pm \varepsilon)\delta]$.*

Proof. We first observe that $\lceil \alpha^{\lfloor \log_\alpha n \rfloor} \rceil \leq n$. To see this, assume otherwise, i.e., that $\alpha^{\lfloor \log_\alpha n \rfloor} = n + x$ for some $x > 0$. Then $x = \alpha^{\lfloor \log_\alpha n \rfloor} - n \leq \alpha^{\log_\alpha n} - n = 0$, contradicting the assumption that $x > 0$. It follows that $A \subseteq [n]$. Now, for the upper bound notice that $\tilde{\delta} \leq \max\{(1 + \varepsilon')d_k/k : k \in A\} \leq (1 + \varepsilon') \cdot \delta \leq (1 + \varepsilon) \cdot \delta$.

For the lower bound, let $\hat{k} \in [n]$ be such that $\delta = d_{\hat{k}}/\hat{k}$ and let i be such that $\lceil \alpha^{i-1} \rceil \leq \hat{k} \leq \lceil \alpha^i \rceil$. Notice that obviously $\lceil \alpha^{i-1} \rceil \in A$, but also $\lceil \alpha^i \rceil \in A$ as $\alpha^i \leq \alpha \hat{k} \leq \alpha n/2 \leq n$ by Observation A.2. We now distinguish two cases: (1) $\lceil \alpha^i \rceil = \lceil \alpha^{i-1} \rceil + 1$ and (2) $\lceil \alpha^i \rceil \geq \lceil \alpha^{i-1} \rceil + 2$. In case (1), we get that $\hat{k} \in \{\lceil \alpha^{i-1} \rceil, \lceil \alpha^i \rceil\} \subseteq A$ and consequently $\tilde{\delta} \geq \max\{(1 - \varepsilon')d_k/k : k \in A\} = (1 - \varepsilon') \cdot \delta \geq (1 - \varepsilon) \cdot \delta$. In case (2), it holds that

$$\alpha^{i-1} \cdot \varepsilon' = \alpha^i - \alpha^{i-1} \geq \lceil \alpha^i \rceil - 1 - \alpha^{i-1} \geq \lceil \alpha^{i-1} \rceil + 1 - \alpha^{i-1} \geq 1. \qquad (1)$$

12

Now let $\beta := \lceil \alpha^i \rceil - \hat{k}$. We note that $d_{j+1} \geq d_j - 1$ for every j, as every distinct length-j substring other than possibly $S[n - j + 1..n]$ gives at least one distinct length-$j + 1$ substring. Applying the same observation iteratively yields $d_{\lceil \alpha^i \rceil} \geq d_{\hat{k}} - \beta$. Hence

$$\tilde{\delta} \geq (1 - \varepsilon') \cdot \frac{d_{\lceil \alpha^i \rceil}}{\lceil \alpha^i \rceil} \geq (1 - \varepsilon') \cdot \frac{d_{\hat{k}} - \beta}{\hat{k} + \beta} = (1 - \varepsilon')\delta \cdot \frac{1 - \frac{\beta}{d_{\hat{k}}}}{1 + \frac{\beta}{\hat{k}}} \geq (1 - \varepsilon')\delta \cdot \frac{1 - \frac{\beta}{2\hat{k}}}{1 + \frac{\beta}{\hat{k}}},$$

where we used that $\delta(S) = d_{\hat{k}}/\hat{k} \geq 2$ in the last step. We can now upper bound β by $\lceil \alpha^i \rceil - \lceil \alpha^{i-1} \rceil \leq \alpha^i + 1 - \alpha^{i-1} = \alpha^{i-1} \cdot \varepsilon' + 1 \leq \hat{k}\varepsilon' + 1$. This yields

$$\hat{\delta} \geq (1 - \varepsilon')\delta \cdot \frac{1 - \frac{\varepsilon'}{2} - \frac{1}{2\hat{k}}}{1 + \varepsilon' + \frac{1}{\hat{k}}} \geq (1 - \varepsilon')\delta \cdot \frac{1 - \varepsilon'}{1 + 2\varepsilon'} \geq (1 - \varepsilon) \cdot \delta,$$

where the second inequality uses that $\hat{k} \geq 1/\varepsilon'$ following from (1) and the last inequality uses the definition of ε'. $\qquad\square$

B Details on Bookmarking the RLBWT

We show how to extract $S[|S| - k + 1]$ from the RLBWT, for any of the sampled lengths k. See also the example in Figure 1. We show how to initialize and update (upon character extensions of the stream) an index (bookmark) j such that $BWT[j] = S[|S| - k + 1]$. This allows us retrieving $S[|S| - k + 1]$ in $O(\log |S|) \subseteq O(\log n)$ time with a random access operation $BWT[j]$ on the RLBWT data structure.

We first discuss how to initialize the bookmark j as soon as the stream's length becomes $S = k + 1$ (before that, the window of the last k characters is not completely filled). The initialization works by setting $j = BWT.LF(i)$, where i is the position such that $BWT[i] = \$$. Since the LF mapping on the BWT of the reversed stream corresponds to advancing one position in the stream, it is easy to see that, after this operation, it holds $BWT[j] = S[|S| - k + 1]$. See Figure 1 for an example.

Suppose we are storing the bookmark j such that $BWT[j] = S[|S| - k + 1]$. We now show how to update j when a new character a arrives; let $S' = Sa$ be the updated stream. Our goal is to modify j so that $BWT[j] = S'[|S'| - k + 1]$ holds. The observation is that, upon the extension of the stream by one character a, the algorithm of [19] modifies the BWT as follows: letting i being the index such that $BWT[i] = \$$, the algorithm (1) replaces $BWT[i] \leftarrow a$, and (2) inserts $\$$ in the position i' corresponding to the lexicographic rank of the new reversed stream $(Sa)^R$ (position i' is computed in $O(\log |S|)$ time using basic operations on the RLBWT, see [19] and Example 1): the new BWT becomes $BWT \leftarrow BWT[1, i' - 1] \cdot \$ \cdot BWT[i' + 1, |S|]$. If $j < i'$ (i.e. $\$$ is inserted after position j), then after these modification we have that $BWT[j] = S'[|S'| - k]$; if, on the other hand, $j \geq i'$ (i.e. $\$$ is inserted before position j), then we increment j as $j \leftarrow j + 1$, and $BWT[j] = S'[|S'| - k]$ holds also in this case. Finally, we need to "advance" j by one position on the stream; this operation corresponds to one LF mapping step on the BWT: $j \leftarrow BWT.LF(j)$ ($O(\log |S|)$ time). After these operations, we finally have that $BWT[j] = S'[|S'| - k + 1]$.

C Detailed Experimental Results

C.1 *Estimation of d_k.*

As mentioned above, there are two types of errors in the computation of the approximation $\tilde{\delta}$ of δ: (1) the error obtained when approximating the values d_k by $\tilde{d_k}$ and (2) the error due to the restriction of the string's offsets $[n]$ to the "sampled set" A. The error of type (1) itself has two sources, namely (1.1) errors due to collisions when computing fingerprints with Rabin's hash function and (1.2) errors due to the count-distinct sketch when applied to the fingerprints. We experimentally evaluated the error of type (1.1) and (1.2) as follows. For the Pizza&Chili repetitive corpus, we compute the exact values of d_k and their estimated values $\tilde{d_k}$ for $k \in \{2^i : 0 \le i \le 7\}$. We observe that the error (1.1) caused by collisions in Rabin's fingerprint were negligible; the error in the ratio of the distinct number of fingerprints and the actual number of distinct substrings was less than 0.01%. The error due to the count-distinct sketch (1.2) was dependent on its parameter: the number of registers used for estimation. It is worth noting that the number of registers does not affect the time complexity when updating sketches, but only affects the space usage by a constant factor (and the time to compute the actual estimation at the end, which is negligible). When more than 2^{14} registers were used for count-distinct sketches, the maximum relative error was observed to be below 2%, and the average error on d_k was below 0.5%; see Table 2.

The number of registers	2^{10}	2^{12}	2^{14}	2^{16}
Maximum	0.0641	0.0313	0.0186	0.0077
Average	0.0204	0.0089	0.0041	0.0017

Table 2: The relative error measured in estimating d_k

C.2 *Experiments on phylogenetic tree reconstruction.*

To show similar behavior of $\text{NCD}_{\mathbf{xz}}$, NCD_δ, and $\text{NCD}_{\tilde{\delta}}$, we conducted experiments on phylogenetic tree reconstruction using `Gene-Trees` dataset from AF project[7]. It contains 11 groups of sequences (651 sequences in total) where each group yields a tree. We constructed 11 phylogenetic trees (i.e. one tree for each group) for each of the NCD measures, and compare the constructed trees by measuring the normalized Robinson-Foulds (nRF) distance, a widely-used distance measure for this purpose. The distance tends to 0 as trees become similar, and tends to 1 when comparing with a random tree. The average nRF between $\text{NCD}_{\mathbf{xz}}$ and NCD_δ was measured as 0.2, indicating similar trees were reconstructed. The average nRF between NCD_δ and $\text{NCD}_{\tilde{\delta}}$ ranges from 0.115 to 0.250 depending on the parameters. For ease of interpretation of these
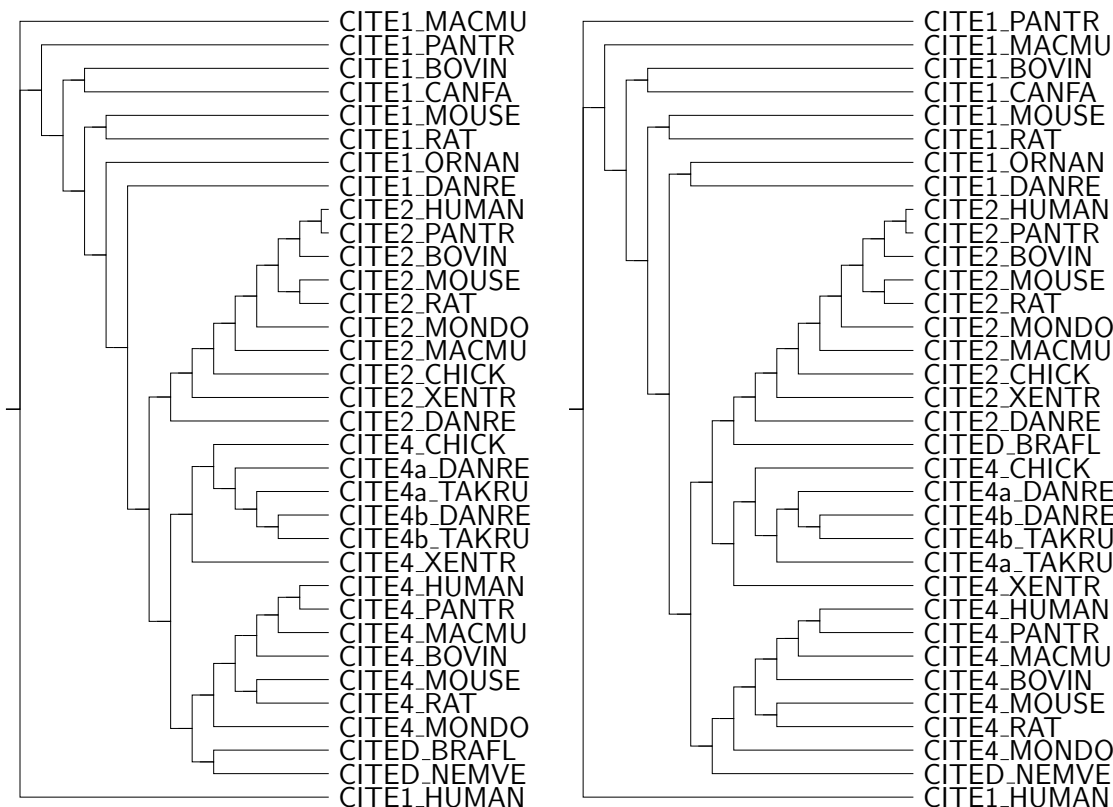
---

[7] https://afproject.org/

Figure 3: Two similar phylogenetic trees constructed using normalized compression distance (NCD) with estimated normalized substring complexity $\tilde{\delta}$ (left) and a popular compression software `xz` (right). Normalized Robinson-Foulds distance is 0.194.

values, we depict two similar phylogenetic trees with nRF=0.194 in Figure 3, which is an actual example of reconstructed trees using NCD with $\tilde{\delta}$ and `xz`.

Running Time. To construct a phylogenetic tree from a sequence set, we usually need to compute all-pair distances. When sequences are long, computing NCDs can be quite costly because we need to compress concatenated sequences for all pairs of sequence in the input set. On the other hand, our sketching can be more efficient because we only need to compute sketches for each sequence, then merging sketches can be done very quickly compared to processing the entire sequences all over again. To demonstrate this, we measure the running time for computing all-pair NCDs on `E.coli` dataset from AF Project that consists of 29 sequences of average length 81,588. Computing all-pair NCDs with `xz` and the exact δ took about 42 and 24 minutes. On the other hand, our sketching method took only 3 minutes.

15

Neither the European Union nor the granting authority can be held responsible for them.