

Data Deduplication with Random Substitutions

Hao Lou, Farzad Farnoud

Electrical and Computer Engineering, University of Virginia, VA, USA. Email: {haolou,farzad}@virginia.edu

Abstract—Data deduplication saves storage space by identifying and removing repeats in the data stream. Compared with traditional compression methods, data deduplication schemes are more computationally efficient and are thus widely used in large scale storage systems. In this paper, we provide an information-theoretic analysis of the performance of deduplication algorithms on data streams in which repeats are not exact. We introduce a source model in which probabilistic substitutions are considered. More precisely, each symbol in a repeated string is substituted with a given edit probability. Deduplication algorithms in both the fixed-length scheme and the variable-length scheme are studied. The fixed-length deduplication algorithm is shown to be unsuitable for the proposed source model as it does not take into account the edit probability. Two modifications are proposed and shown to have performances within a constant factor of optimal for a specific class of source models with the knowledge of model parameters. We also study the conventional variable-length deduplication algorithm and show that as source entropy becomes smaller, the size of the compressed string vanishes relative to the length of the uncompressed string, leading to high compression ratios.

I. INTRODUCTION

The task of reducing data storage costs is gaining increasing attention due to the explosive growth of the amount of digital data, especially redundant data [3], [10], [18]. Data deduplication is a data reduction approach that eliminates duplicate data at the file or subfile level. Compared with traditional data compression approaches, data deduplication is more efficient when dealing with large-scale data. It has been widely used in mass data storage systems, e.g., LBFS (low-bandwidth network file system) [12] and Venti [14]. In this paper, we aim to study the performance of data deduplication algorithms from an information-theoretic point of view when repeated data segments are not necessarily exact copies.

A typical data deduplication system uses a chunking scheme to parse the data stream into multiple data ‘chunks’. Chunks are entered into the dictionary at the first occurrences, and duplicates are replaced by pointers to the dictionary. The chunks can be of equal length (fixed-length chunking) or of lengths that are content-defined (variable-length chunking) [8]. The fixed-length scheme has low complexity but suffers from the boundary-shift problem: if insertions or deletions occur in a part of the data stream, then all subsequent chunks are

changed because the boundaries are shifted. In the variable-length scheme, chunk breakpoints are determined using pre-defined patterns and therefore edits will not affect subsequent chunks and repeated data segments can still be identified.

An information-theoretic analysis of deduplication algorithms was first performed by Niesen [13]. Niesen’s work introduced a source model, formalized deduplication algorithms in both fixed-length and variable-length schemes, including (conventional) fixed-length deduplication (FLD) and variable-length deduplication (VLD), and analyzed their performance. We adopt a similar strategy in this paper. The source model introduced by Niesen produces data strings that are composed of blocks, with each block being an exact copy of one of the source symbols, where the source symbols are pre-selected strings. It is often the case, however, that the copies of a block of data that is repeated many times are approximate, rather than exact. This may occur, for example, due to edits to the data, or in the case of genomic data¹ due to mutations. Thus, in our source model, we add probabilistic substitutions to each block, resulting in data streams composed of approximate copies of the source symbols.

We then analyze data deduplication algorithms over source models with probabilistic edits. For the fixed-length scheme, three algorithms: a generalization of FLD [13] named *modified fixed-length* deduplication (mFLD), a variant of mFLD named *adaptive fixed-length* deduplication (AFLD), and the *edit-distance* deduplication (EDD), are presented and analyzed. Due to the boundary-shift problem, algorithms in the fixed-length scheme are studied over the source model where all source symbols have the same length. We show that for mFLD, if the chunk length is not properly chosen, the average length of the compressed strings is greater than source entropy by an arbitrarily large multiplicative factor for small enough edit probability. Meanwhile, AFLD and EDD take source model parameters into account and are shown to have performances within a constant factor of optimal. For the variable-length scheme, we consider the general scenario where source symbols are of random lengths. We show that VLD can achieve large compression ratios relative to the length of the uncompressed strings.

A large number of works have studied data deduplication; see [20] for a comprehensive survey. However, the problem is not well-studied from an information-theoretic point of view. This is important because information-theoretic analysis would enable comparing the performance of deduplication algorithms with theoretical limits, under appropriate probabilistic models, and guide the development of more efficient, possibly optimal,

This paper was presented in part at ISIT 2020 [7].

Hao Lou is with the Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, 22903, USA, (email: haolou@virginia.edu).

Farzad Farnoud (Hassanzadeh) is with the Department of Electrical and Computer Engineering and the Department of Computer Science, University of Virginia, Charlottesville, VA, 22903, USA, (email: farzad@virginia.edu).

This work was supported by NSF grant under grant no. CIF-1755773.

¹Repeats are common in genomic data. For example, a majority of the human genome consists of interspersed and tandem repeated sequences [6].

algorithms. In addition to the seminal work by Niesen [13], the work [19] also analyzed deduplication from an information-theoretical point of view but with a source model that is incompatible with the current analysis. The problem of deduplication under edit errors was also considered in [1]. While [1] focuses on performing deduplication on two files, one being an edited version of the other by insertions and deletions, we consider a single data stream with substitution errors.

The rest of the paper is organized as follows. Notation and preliminaries are given in the next section. In Section III, we introduce the information source model and bound its entropy. In Section IV, we formally state the deduplication algorithms. In Section V, we summarize the main results of this paper. Bounds on the performances of algorithms in the fixed- and the variable-length schemes are derived in Section VI and VII, respectively. We close the paper with concluding remarks and open problems in Section VIII.

II. PRELIMINARIES

We consider the binary alphabet $\{0, 1\}$, denoted Σ . The set of all finite strings over Σ (including the unique empty string) is denoted Σ^* . A j -(sub)string is a (sub)string of length j . For a non-negative integer m , let Σ^m be the set of all strings of length m over Σ . For strings $u, v \in \Sigma^*$, the concatenation of u and v is denoted uv , and the concatenation of i copies of u is denoted u^i . We denote the substring of length ℓ starting from the j -th symbol of u by $u_{j,\ell}$, which is also referred to as the j -th ℓ -substring of u . The length of u is denoted $|u|$. The cardinality of a set S is also denoted $|S|$. For a set T of strings, u is said to be a substring of T if u is a substring of one or more strings in T .

In this paper, all logarithms are to the base 2. For $0 \leq p \leq 1$, $H(p)$ denotes the binary entropy function: $p \log\left(\frac{1}{p}\right) + (1-p) \log\left(\frac{1}{1-p}\right)$. For $0 \leq p, q \leq 1$, $H(p, q)$ denotes the cross entropy function: $p \log\left(\frac{1}{q}\right) + (1-p) \log\left(\frac{1}{1-q}\right)$. For an event \mathcal{E} , we use $\bar{\mathcal{E}}$ to denote its complement and use $I_{\mathcal{E}}$ to denote the indicator variable for \mathcal{E} , which takes value 1 when \mathcal{E} is true, and 0 otherwise.

The following inequalities are used frequently: for $x \in (0, 1)$ and a positive integer n ,

$$\frac{1}{2} \min(1, nx) \leq 1 - (1-x)^n \leq \min(1, nx). \quad (1)$$

A binary string is k -runlength-limited (k -RLL) [9] if it does not contain k consecutive zeros, i.e., all runs of zeros in the string are of lengths less than k . We denote the set of binary k -RLL strings by R_k and denote the set of binary k -RLL strings of length n by R_k^n . The following lemma provides bounds on the size of R_k^n .

Lemma 1. *Let k be a positive integer. The number of binary k -RLL strings of length n , $|R_k^n|$, satisfies*

$$(2 - \frac{1}{2^{k-2}})^n \leq |R_k^n| \leq 2(2 - \frac{1}{2^k})^n.$$

Lemma I is proved by induction in Appendix A. By Lemma I, we bound the number of binary k -RLL strings of lengths at most 2^k in the following corollary.

Corollary 1. *The number of binary k -RLL strings of lengths at most 2^k satisfies*

$$\sum_{n=0}^{2^k} |R_k^n| \geq \sum_{n=0}^{2^k} \left(2 - \frac{1}{2^{k-2}}\right)^n \geq 2^{2^k-2}.$$

III. SOURCE MODEL

The source model studied in this paper extends the one described in [13] by allowing probabilistic substitutions. The output data stream s is a concatenation of approximate copies of source symbols. The A source symbols, denoted X_1, X_2, \dots, X_A , are iid binary strings generated in the following way. Fix a length distribution \mathbb{P}_l over positive integers with mean L . For each $1 \leq a \leq A$, we draw L_a from \mathbb{P}_l and draw X_a uniformly from Σ^{L_a} . It is important to note that, as a result of sampling with replacement, the source symbols are distributed uniformly and independently. The probability that $(X_1, \dots, X_A) = (\mathbf{x}_1, \dots, \mathbf{x}_A)$ given the lengths L_a is $\prod_{a=1}^A \frac{1}{2^{L_a}}$, for any set of strings (\mathbf{x}_a) where \mathbf{x}_a has length L_a . So the same sequence can be drawn multiple times as source symbols. The draws are treated as separate symbols, but with the same content. We use \mathcal{X} to denote the source symbol alphabet, i.e., $\mathcal{X} = \{X_1, X_2, \dots, X_A\}$. The alphabet is thus a multiset. To simplify some of the derivations, we adopt the same assumption as [13] that \mathbb{P}_l is concentrated around its mean, specifically, $\mathbb{P}_l(\frac{L}{2} \leq l \leq 2L) = 1$.

After generating the source symbols X_1, X_2, \dots, X_A , we generate an iid sequence of length B , denoted Y_1, \dots, Y_B , where each Y_b is an approximate copy of a randomly chosen source symbol. Specifically, for each $1 \leq b \leq B$, we first pick J_b uniformly at random from $\{1, 2, \dots, A\}$. Next, we generate Y_b by flipping each bit of X_{J_b} independently with probability δ , as a way of simulating edits and other changes to the data in a simple manner. The bit flipping process is referred to as a δ -edit. As an example, if $X_{J_b} = 000000$, then a possible outcome of the δ -edit could be 001001, which has probability $\delta^2(1-\delta)^4$. The data stream s will be a concatenation of Y_1, Y_2, \dots, Y_B , i.e., $s = Y_1 Y_2 \dots Y_B$. The approximate copies Y_1, Y_2, \dots, Y_B are referred to as source blocks. The real number δ is referred to as the *edit probability*. The entropy of this source is denoted $H(s)$. Note that given s , the boundaries between source blocks are not known to us.

In this paper, we study the asymptotic regime in which $B, A, L \rightarrow \infty$ while the edit probability δ remains a constant less than $\frac{1}{2}$. We consider the situation where A, L are functions of B with $A \leq B^{1-k_2}$ for some $0 < k_2 < 1$ and $L = \Theta(B^{k_1})$ for some $k_1 > 0$. We allow A to grow large because it is reasonable to assume that as the dataset gets larger, the number of unique blocks is also higher. This necessitates L to also grow large. The assumption $A \leq B^{1-k_2}$ ensures that, on average, every source symbol has repeats. The polynomial relationship between L and B ensures that B is much smaller than $2^{\Theta(L)}$. So only a small fraction of all possible strings of length $\Theta(L)$ can appear as source symbols, or edited versions of the source symbols, in the datastream. This is compatible with our intuition that only a small number of all possible strings are valid data, e.g., an image, or a piece of text or

code. Furthermore, the polynomial relationship between B and L appears to agree with results from experiments in [18] (also referred to in [13]) suggesting that the reasonable range for L is from a few KB to a few MB ($\approx 10^4$ to 10^7 bits) and for B is on the order of 10^5 to 10^9 . Nevertheless, other asymptotic regimes may also be appropriate but are left to future work for simplicity.

The following lemma provides asymptotic bounds on $H(s)$. Under our assumptions, $H(s)$ is shown to be dominated by the term $H(\delta)BL$, i.e., the main component of the entropy is the uncertainty arising from the random substitutions.

Lemma 2. *As $B \rightarrow \infty$, the entropy of the source model with edit probability δ satisfies*

$$H(\delta)BL \leq H(s) \leq H(\delta)BL + B \log A + A(2L + 1).$$

Proof: For the lower bound,

$$H(s) \geq H(s|\mathbf{X}_{J_1}, \dots, \mathbf{X}_{J_B}) = \sum_{b=1}^B H(Y_b|\mathbf{X}_{J_b}) = H(\delta)BL.$$

For the upper bound,

$$H(s) \leq H(s|\mathbf{X}_{J_1}, \dots, \mathbf{X}_{J_B}) + H(\mathbf{X}_{J_1}, \dots, \mathbf{X}_{J_B}|\mathcal{X}) + H(\mathcal{X})$$

$$\leq H(\delta)BL + B \log A + A(2L + 1),$$

where $H(\mathcal{X}) \leq A(2L + 1)$ follows from the fact that for each \mathbf{X}_a , there are at most 2^{2L+1} different possibilities since we assume $L_a \leq 2L$. ■

A deduplication algorithm is said to *(asymptotically) achieve a constant factor of optimal* if there exists a constant c (independent of δ) such that $\mathbb{E}[\mathcal{L}(s)] \leq cH(s)$, for all $0 < \delta < \frac{1}{2}$ and all sufficiently large B , where $\mathcal{L}(s)$ is the length of the encoding produced by the algorithm. Given our assumptions on A, B, L , and the result from Lemma 2, the entropy $H(s)$ is dominated by the term $H(\delta)\mathbb{E}[|s|]$. If δ is close to $\frac{1}{2}$, $H(s)$ is close to the length of the uncompressed sequence (s is close to an iid Bernoulli(1/2) process), while if δ is close to 0, there is large gap between the two. Hence, to determine whether an algorithm achieves a constant factor of optimal, the case of small δ is especially important, which is also the case where compression is more beneficial.

We also define the compression ratio $R = \frac{\mathbb{E}[|s|]}{\mathbb{E}[\mathcal{L}(s)]}$. Note that if there exists a constant c_1 independent of δ such that $R \leq c_1$, then the algorithm uses more bits than the entropy by an arbitrarily large multiplicative factor as δ goes to 0. While if $R \rightarrow \infty$ as $\delta \rightarrow 0$, then the algorithm can achieve arbitrarily large compression ratios as entropy decreases. Finally, if there exists a constant c_2 such that $R \geq \frac{c_2}{H(\delta)}$ for all valid δ , then the algorithm achieves a constant factor of optimal.

We discuss some strategies that we use in the rest of the paper for computing $\mathbb{E}[\mathcal{L}(s)]$. We say \mathbf{X}_{J_b} is the ancestor of Y_b and Y_b is a descendant of \mathbf{X}_{J_b} . For each a , we use $Y(a)$ to denote the set $\{1 \leq b \leq B : J_b = a\}$ and use $Y_{1/2}(a)$ to denote the set $\{1 \leq b \leq \lceil B/2 \rceil : J_b = a\}$. In other words, $Y(a)$ is the set of source block indexes of the descendants of \mathbf{X}_a and $Y_{1/2}(a)$ is the set of source block indexes of the descendants of \mathbf{X}_a among the first half of source blocks.

Note that $\mathbb{E}[|Y(a)|] = B/A$ and $\mathbb{E}[|Y_{1/2}(a)|] = B/(2A)$. We use \mathcal{E}_u to denote the event that $|Y(a)| \leq \frac{3B}{2A}$ for all $1 \leq a \leq A$, and use \mathcal{E}_l to denote the event that $|Y_{1/2}(a)| \geq \frac{B}{4A}$ for all $1 \leq a \leq A$. Since $|Y(a)| = \sum_{b=1}^B I_{J_b=a}$, where all summands are iid with expected value $\frac{1}{A}$, by the Chernoff bound [11] and the union bound,

$$\Pr(\mathcal{E}_u) \geq 1 - Ae^{-\frac{B}{10A}}, \quad \Pr(\mathcal{E}_l) \geq 1 - Ae^{-\frac{B}{16A}}. \quad (2)$$

Given our assumption that $A \leq B^{1-k_2}$, asymptotically $\frac{B}{16A} - \log A$ goes to infinity. So the probability of \mathcal{E}_u goes to 1 (also true for \mathcal{E}_l). In the performance analysis of deduplication algorithms, we generally only need to consider the case in which \mathcal{E}_l or \mathcal{E}_u holds. Specifically, we use the following inequalities as bounds on $\mathbb{E}[\mathcal{L}(s)]$:

$$\begin{aligned} \mathbb{E}[\mathcal{L}(s)] &\leq \mathbb{E}[\mathcal{L}(s)|\mathcal{E}_u] + \mathbb{E}[\mathcal{L}(s)|\bar{\mathcal{E}}_u] \cdot \Pr(\bar{\mathcal{E}}_u), \\ \mathbb{E}[\mathcal{L}(s)] &\geq \mathbb{E}[\mathcal{L}(s)|\mathcal{E}_l] \cdot \Pr(\mathcal{E}_l) = \mathbb{E}[\mathcal{L}(s)|\mathcal{E}_l] \cdot (1 - \Pr(\bar{\mathcal{E}}_l)). \end{aligned}$$

To find $\mathbb{E}[\mathcal{L}(s)]$, we generally compute the terms $\mathbb{E}[\mathcal{L}(s)|\mathcal{E}_u]$, $\mathbb{E}[\mathcal{L}(s)|\mathcal{E}_l]$ and show that the terms $\mathbb{E}[\mathcal{L}(s)|\bar{\mathcal{E}}_u] \cdot \Pr(\bar{\mathcal{E}}_u)$ and $\mathbb{E}[\mathcal{L}(s)|\bar{\mathcal{E}}_l] \cdot \Pr(\bar{\mathcal{E}}_l)$ are asymptotically negligible, using trivial bounds on $\mathcal{L}(s)$.

IV. DEDUPLICATION SCHEMES

In this section, we formally state the deduplication algorithms, which can be regarded as mathematical abstractions of real-world deduplication systems. All algorithms are dictionary-based and composed of two stages: chunking and encoding. In particular, the conventional fixed-length deduplication (FLD) and variable-length deduplication (VLD) algorithms were formalized in [13] and are restated here.

In FLD, the chunk length ℓ is fixed. Source string s is parsed into segments of length ℓ , i.e., $s = z_1 z_2 \dots z_{C+1}$, where $|z_1| = |z_2| = \dots = |z_C| = \ell$, $C = \lfloor |s|/\ell \rfloor$. The substrings $\{z_c\}_{c=1}^{C+1}$ are collected as deduplication chunks. The encoding process starts with encoding the length of s by a prefix-free code for positive integers, such as the Elias gamma code [2], to ensure that the whole scheme is prefix-free. The chunks are then encoded sequentially. Starting with $c = 1$, if chunk z_c appears for the first time, i.e., $z_c \neq z_i$ for all $i < c$, then it is encoded as the bit 1 followed by z_c itself and is entered into the dictionary. Otherwise, when there already exists an entry in the dictionary storing the same string as z_c , it will be encoded as the bit 0 followed by a pointer to that entry. The pointer is an index of the dictionary entries and thus can be encoded by at most $\log|T^{c-1}| + 1$ bits, where T^{c-1} denotes the dictionary just after z_{c-1} is processed. The number of bits FLD takes to encode s is denoted $\mathcal{L}_F(s)$. It was shown in [13] that FLD is ineffective when source symbols have different lengths. So in this paper, we study FLD (as well as its variations, mFLD and AFLD, described below) only for sources in which all source symbols have the same length. We note that such sources are not realistic except for some scenarios such as deduplication in virtual machine disk images [5]. However, the analysis of FLD and its variants is helpful for the study of VLD, described next, as it reveals

important insights about the effect of chunk lengths on the performance.

Example 1. For $s = 01101101$ and $\ell = 2$, the chunks generated by fixed-length chunking are $z_1 = 01, z_2 = 10, z_3 = 11, z_4 = 01$. The encoding of length $|s| = 8$ by Elias gamma coding is 0001000. Chunks z_1, z_2 and z_3 are new chunks and thus are encoded as 101, 110, 111, respectively. Chunk z_4 is a duplicate of z_1 . When z_4 is processed, the dictionary contains three strings 01, 10 and 11. So z_3 is encoded as 000, where the first 0 indicates that the chunk is repeated and the following 00 represents the first entry of the dictionary. Concatenating all components, the final encoding of s is 0001000101110111000. Note that after encoding terminates, the dictionary is the ordered set $\{01, 10, 11\}$, which appear in the encoded string as the set of chunks preceded by indicator bits with value 1.

In VLD, a string of length M (we assume 0^M) is chosen as the marker string. The source string s is parsed into chunks that end with the marker string. Specifically, the source string s is parsed as $s = z_1 \cdots z_C$, where each z_c (except for perhaps the last one) contains a single appearance of 0^M at the end. We again use z_1, \dots, z_C to represent the chunks. After splitting s into the chunks $\{z_c\}_{c=1}^C$, the same dictionary encoding process as in FLD is conducted. The number of bits variable-length deduplication takes to encode s is denoted $\mathcal{L}_{VL}(s)$.

Example 2. Consider the same string $s = 01101101$ as Example 1. VLD, with marker length $M = 1$, parses s as chunks 0, 110, 110, 1. The length of s is still encoded by 0001000. Chunks 0, 110, 1 are new and are encoded with 10, 1110, 11, respectively. The second occurrence of 110 is encoded by a 0 followed by the pointer 1. The final encoding of s is thus 00010001011100111.

The *modified fixed-length* deduplication (mFLD) has the same encoding process as FLD but with a two-stage chunking process. In mFLD, first, the source string s is parsed into segments of length D , and then, each segment is parsed into chunks of length ℓ , where $\ell \leq D$. Specifically, the source string s is parsed as

$$s = x_1 x_2 \cdots x_{K+1}, \quad |x_1| = |x_2| = \cdots = |x_K| = D,$$

where $K = \lfloor |s|/D \rfloor$ and

$$x_k = z_k^1 z_k^2 \cdots z_k^{N+1}, \quad |z_k^1| = |z_k^2| = \cdots = |z_k^N| = \ell,$$

with $1 \leq k \leq K$, $N = \lfloor D/\ell \rfloor$ (x_{K+1} is parsed in the same way). The number of bits mFLD takes to encode s is denoted $\mathcal{L}_{mF}(s)$.

Note that mFLD is a generalization of FLD since with $D = \ell$, mFLD is equivalent to FLD with the same chunk length ℓ . For FLD to perform well, the source symbols must all have the same length L and the chunk length ℓ must also be chosen equal to L to maintain synchronization between the chunks and symbols. The generalization to mFLD allows us to maintain synchronization by setting $D = L$ and frees us to choose values other than the symbol length for the chunk length ℓ . This flexibility enables us to study the effect of chunk length, which as we will see, will provide important intuitions for

more practical algorithms such as VLD. We will focus on analyzing the performance of mFLD and report that of FLD as a corollary.

The *adaptive fixed-length* deduplication (AFLD) is a specialization of mFLD with source model parameters taken into account. Given A, B, L, δ , AFLD is the version of mFLD with chunk length specified as $\ell = \lceil \frac{\log(B/A)}{H(\gamma, \delta)} \rceil$ ($\ell = D$ if $D < \lceil \frac{\log(B/A)}{H(\gamma, \delta)} \rceil$) for some $\gamma \in (\delta, 1/2)$. AFLD thus contains two parameters D and γ . Note that in practice, source model parameters can be estimated from data. We will show later that AFLD is an optimized version of mFLD. The distinction in names is made to emphasize the optimality and also for the convenience of referring to this version of the algorithm. The number of bits AFLD takes to encode s is denoted $\mathcal{L}_{AF}(s)$.

Edit-distance deduplication (EDD) extends FLD by encoding chunks relative to previously observed *similar* chunks, if any. EDD takes the source model parameters into account and is only defined for source models with edit probability $\delta < 1/4$. EDD has two parameters, chunk length ℓ and mismatch ratio β , where $\delta < \beta \leq 1/4$. The chunking scheme is the same as in FLD, i.e., parsing the source string s into chunks of length ℓ , denoted z_1, z_2, \dots, z_{C+1} . The encoding starts with a prefix-free code representing the length of the source string. Next, each chunk z_c is encoded as the bit 1 followed by itself if no chunk has appeared before whose Hamming distance from z_c is at most $2\beta\ell$. Otherwise, let c' be the smallest index such that the Hamming distance between $z_{c'}$ and z_c is $\leq 2\beta\ell$. Chunk z_c will be encoded as the bit 0 followed by a pointer to the dictionary entry where $z_{c'}$ is stored, along with the bits describing the mismatches between z_c and $z_{c'}$. The mismatches are the indexes of positions in which $z_{c'}$ and z_c differ. Since we restrict the number of mismatches to be no more than $2\beta\ell$, the mismatches can be encoded in at most $\log\left(\sum_{i=0}^{\lfloor 2\beta\ell \rfloor} \binom{\ell}{i}\right) + 1 \leq H(2\beta)\ell + 1$ bits. The number of bits EDD uses to store s is denoted by $\mathcal{L}_{ED}(s)$.

Encoding differences between similar chunks is usually used as a post-deduplication process, which spends extra computation to eliminate redundancy among distinct but similar chunks [16], [17], [21], [22]. In this paper, we study EDD as a simple abstraction of this type of algorithms and only consider the fixed-length chunking scheme. An edit-distance based variable-length algorithm may potentially lead to better performance and be more practically important. We leave it to future consideration due to the technical challenges in the analysis, primarily arising from the facts that chunk boundaries may shift because of edits and that deriving the statistics of the number of detected copies within a certain distance does not appear readily tractable.

V. RESULTS

In this section, we summarize the main results of the paper. Detailed analysis and proofs of these results will be provided in the following corresponding sections.

A. Modified fixed-length deduplication and its variants

We first present results for mFLD and its variants AFLD and FLD. Fixed-length deduplication has been shown in [13] to not

perform well when source symbols have variable lengths. So for algorithms in the fixed-length scheme, we assume \mathbb{P}_l is degenerate and let the first-stage parsing length be equal to the source symbol length.

The mFLD algorithm allows us to set the chunk length ℓ . The effect of this length is investigated in Theorems 6, 7 and 8. For simplicity of presentation, we give detailed analysis about the theorems in Section VI and provide corollaries here as summaries.

Corollary 2. *Consider the source model in which source symbols all have length L . For mFLD with $D = L$, if the chunk length $\ell = o(\log B) \cup \omega(\log B)$, the compression ratio $\frac{\mathbb{E}[|s|]}{\mathbb{E}[\mathcal{L}_{mF}(s)]}$ is upper bounded by a universal constant for any edit probability $\delta > 0$.*

Corollary 2 follows directly from Theorems 6 and 7. It characterizes the performance of mFLD when the chunk length ℓ is chosen too small or too large. With the chunk length improperly chosen, the average length of the compressed strings is always at least a constant factor of the original length, regardless of the edit probability δ . This is not desirable for small δ since, as δ goes to 0, the entropy gets smaller and the ratio $\frac{\mathbb{E}[\mathcal{L}_{mF}(s)]}{H(s)}$ grows unboundedly. It can be seen later from the proofs of Theorems 6 and 7 that when the chunk length is chosen too small, the dictionary becomes so large that the pointers become of similar lengths to the chunks. On the other hand, when the chunk length is chosen too large, repeats can not be identified and deduplication thus fails. It is therefore important to pick a suitable chunk length when implementing deduplication algorithms in practice.

If we pick $\ell = L$, mFLD becomes FLD with chunk length equal to source symbol length, which was shown in [13] to be asymptotically optimal on sources with fixed symbol length and no edits. However, in the case when edit probability δ is nonzero, since we assume $L = \Theta(B^{k_1})$, Corollary 2 implies that the compression ratio of FLD is bounded and the gap between FLD and entropy can be arbitrarily large, as stated in the next corollary.

Corollary 3. *Consider the source model in which source symbols all have length L . For FLD, with chunk length L , the compression ratio $\frac{\mathbb{E}[|s|]}{\mathbb{E}[\mathcal{L}_F(s)]}$ is upper bounded by a universal constant for any edit probability $\delta > 0$.*

AFLD has its chunk length chosen adapted to source parameters and is shown in Theorem 8 to be nearly optimal. The following corollary is a summary of Theorem 8.

Corollary 4. *For any edit probability $\delta \in (0, \frac{1}{2})$ and any $a > 1$, there exists $\delta < \gamma < \frac{1}{2}$ such that*

$$\frac{\mathbb{E}[\mathcal{L}_{AF}(s)]}{H(s)} \leq \frac{a(1 + k_1)}{k_2}(1 + o(1)).$$

With k_1, k_2 being fixed constants, the preceding corollary states that AFLD achieves a constant factor of optimal for any edit probability δ . Thus, to achieve high compression ratio, deduplication algorithm parameters, especially the chunk length, should be chosen based on the data. In practice, it can thus be beneficial to first obtain an estimate of the

parameters of the data and then apply deduplication with algorithm parameters properly chosen. A fixed chunk length is unlikely to be universally effective for all datasets.

B. Edit-distance deduplication

The edit-distance deduplication is studied in Theorem 9 and shown to achieve performance a constant factor of optimal.

Theorem 9. *Consider the source model in which source symbols have the same length L and the edit probability is $\delta < \frac{1}{4}$. The performance of edit-distance deduplication with chunk length $\ell = L$ and mismatch ratio β satisfies*

$$1 \leq \frac{\mathbb{E}[\mathcal{L}_{ED}(s)]}{H(s)} \leq \frac{H(2\beta)}{H(\delta)}(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

for any $\delta < \beta \leq \frac{1}{4}$.

Note that for any $\delta < \frac{1}{4}$, we can always find β larger than but close enough to δ such that $\frac{H(2\beta)}{H(\delta)}$ is upper bounded by a constant value. With such choices of β , the preceding theorem states that $\mathbb{E}[\mathcal{L}_{ED}(s)]$ is at most a constant factor of $H(s)$. As an example, let $\beta = \min(\frac{3\delta}{2}, \frac{1}{4})$. The ratio $\frac{H(2\beta)}{H(\delta)}$ is upper bounded by

$$\frac{H(2\beta)}{H(\delta)} \leq \frac{H(\min(3\delta, 1/2))}{H(\delta)} \leq 3,$$

where the last inequality follows from the fact that $\frac{H(3p)}{H(p)} \leq 3$ for all $p \leq \frac{1}{3}$ and $H(\frac{1}{6}) \leq \frac{1}{2}$. Hence, EDD also achieves a constant factor of optimal, as formalized in the following corollary.

Corollary 5. *Consider the source model in which source symbols have the same length L and edit probability $\delta < \frac{1}{4}$. There exists a mismatch ratio β such that the performance of EDD with chunk length $\ell = L$ satisfies*

$$\frac{\mathbb{E}[\mathcal{L}_{ED}(s)]}{H(s)} \leq \frac{H(3\delta)}{H(\delta)}(1 + o(1)) \leq 3(1 + o(1)).$$

We note however that EDD is more complex than AFLD as it identifies chunks that are within a certain Hamming distance.

C. Variable-length deduplication

Similar to the algorithms in the fixed-length scheme, the performance of VLD depends on the chunk length. In VLD, the chunk length is controlled by the length M of the marker (the expected chunk length is approximately 2^M). The effect of M on the performance is studied in Theorems 14, 15, 17 and 18 in Section VII.

As a summary of Theorems 14, 15, and 17, we first present the following corollary, showing that an inappropriate choice of M leads to poor performance.

Corollary 6. *Consider the source model with edit probability δ and variable-length deduplication with marker length M . If $2^M = o(\log B) \cup \omega(\log B)$, the compression ratio $\frac{\mathbb{E}[|s|]}{\mathbb{E}[\mathcal{L}_{VL}(s)]}$ is upper bounded by a universal constant for any edit probability $\delta > 0$.*

We also show that a well-chosen marker length M can lead to arbitrarily large compression ratios as edit probability δ approaches 0.

Theorem 18. *Consider the source model with edit probability $\delta < \frac{1}{2}$. For any $\gamma \in (\delta, 1/2)$, the performance of variable-length deduplication with marker length M such that $2^M = \Theta(\log(B/A))$ satisfies*

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(s)] &\leq \\ &\left(12e^{-c_M}(c_M + 1) + 4H(\gamma, \delta)\frac{(1+k_1)}{k_2}c_M\right)BL(1 + o(1)), \end{aligned} \quad (21)$$

as $B \rightarrow \infty$, where $c_M = \frac{\log(B/A)}{H(\gamma, \delta)2^{M+1}}$.

We perform the following analysis for minimizing the upper bound given by Theorem 18. For any given $c > 0$, there exists an integer value for M such that $c \leq c_M \leq 2c$. For this M , (21) is upper bounded by

$$\left(12e^{-c}(c + 1) + 8H(\gamma, \delta)\frac{(1+k_1)}{k_2}c\right)BL(1 + o(1)),$$

since $e^{-c}(c+1)$ is decreasing in c when $c > 0$. We can always find γ such that $H(\gamma, \delta) \leq 2H(\delta)$. Such γ gives

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(s)] & \\ &\leq \left(12e^{-c}(c + 1) + 16H(\delta)\frac{(1+k_1)}{k_2}c\right)BL(1 + o(1)). \end{aligned} \quad (3)$$

Let $h = 4H(\delta)\frac{(1+k_1)}{3k_2}$. Upper bounding the above expression is equivalent to upper bounding the function $f(c) = e^{-c}(c + 1) + hc$, $c \in (0, +\infty)$. If $h < e^{-1}$, then $f(c)$ has a local minimum at $c = -W_{-1}(-h)$, where W_{-1} is the lower branch of the Lambert W function. If $h \geq e^{-1}$, then $f(c)$ is monotonically increasing in $(0, +\infty)$. Therefore, $c = -W_{-1}(-\min(e^{-1}, h))$ provides an upper bound on $f(c)$. As an example, for $A = L = B^{1/2}$ (i.e., $k_1 = k_2 = \frac{1}{2}$), Figure 1 shows the upper bound given by (3) with $c = -W_{-1}(-\min(e^{-1}, h))$, as well as $H(\delta)$, as δ ranges from 10^{-5} to 10^{-1} .

Note that $h \leq e^{-1}$ holds for small enough δ . When this holds, the upper bound (3) can be rewritten as

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(s)] & \\ &\leq \left(12e^{-c}(c + 1) + 16H(\delta)\frac{(1+k_1)}{k_2}c\right)BL(1 + o(1)) \\ &\leq 12e^{-c}(c^2 + c + 1)BL(1 + o(1)), \end{aligned}$$

where $c = -W_{-1}(-4H(\delta)(1+k_1)/(3k_2))$. Hence the upper bound on the normalized expected compressed length approaches 0 as δ approaches 0. This means that as the entropy becomes smaller, the compression ratio grows if the length of the marker is chosen appropriately. In particular, it can be seen that the proper length of the marker depends on δ , which represents the degree of variability between the copies.

Large compression ratios when entropy is small is desirable and variable-length deduplication achieves this. However, it can be shown and also observed in Figure 1 that the upper bound of the ratio $\mathbb{E}[\mathcal{L}_{VL}(s)]/H(s)$ given by Theorem 18

increases as δ decreases. Therefore, despite the large compression ratios, the gap to entropy may become large for small δ . Determining whether this is indeed the case or the bound provided here is loose is left to future work.

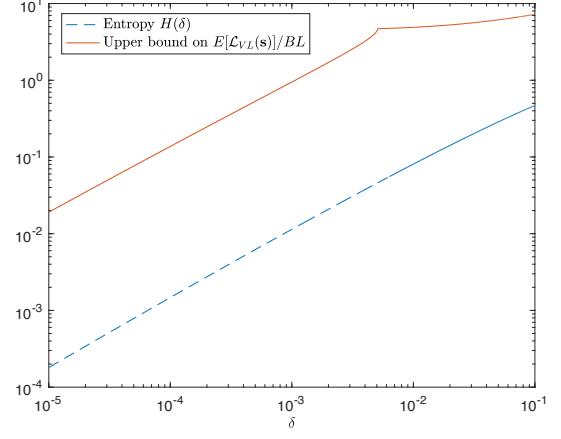


Figure 1. Upper bound on $\frac{\mathbb{E}[\mathcal{L}_{VL}(s)]}{BL}$ and $H(\delta)$ vs the edit probability δ with $A = L = B^{1/2}$, as δ ranges from 10^{-5} to 10^{-1} .

VI. DEDUPLICATION IN THE FIXED-LENGTH SCHEME

In this section, we study the performances of the deduplication algorithms in the fixed-length scheme. It is pointed out by [13] that when all source symbols have the same length and there are no edits, FLD with knowledge of the symbol length can parse data strings in a way that chunk boundaries align with source block boundaries (by setting the chunk length equal to source block length) and achieve asymptotically optimal performance under mild conditions. However, when symbols have different lengths, the loss of synchronization leads to poor performance. For instance, [13] considered the scenario in which there are $A = 2$ source symbols, with the source symbol length distribution \mathbb{P}_l assigning equal probability to L and $L+1$ (here L is an independent parameter rather than the expected value of \mathbb{P}_l) and with $B = 3L$ source blocks. FLD with chunk length $\ell = L$ was shown to satisfy $\frac{\mathbb{E}[\mathcal{L}_F(s)]}{H(s)} \geq \Omega(B)$. In the case where copies are not exact, the question of interest is then whether fixed-length deduplication can still perform well when chunk boundaries align with repeat boundaries. To answer this question, we need to ensure that the two groups of boundaries are aligned. So we consider only source models where source symbols all have the same length L (\mathbb{P}_l is degenerate).

We first study in detail the performance of mFLD and then specialize the results to FLD. The first-stage parsing length of mFLD (including AFLD) and the chunk length of EDD are both assumed to be equal to L .

We present a lemma that will be used frequently. For positive integers m, ℓ and $\delta \in (0, \frac{1}{2})$, define

$$\mathcal{S}_\delta(\ell, m) = \sum_{t=0}^{\ell} \binom{\ell}{t} \min(1, m\delta^t(1-\delta)^{\ell-t}).$$

Lemma 3. Let \mathbf{r} be a string drawn uniformly at random from Σ^ℓ . Let $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m$ be m iid descendants of \mathbf{r} by δ -edit and let $\mathbf{r}_{[m]} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m\}$. For any $\mathbf{w} \in \Sigma^\ell$, let $\mathbf{w} \in \mathbf{r}_{[m]}$ denote the event that $\mathbf{w} = \mathbf{r}_i$ for some i . Then

$$\frac{1}{2} \frac{\mathcal{S}_\delta(\ell, m)}{2^\ell} \leq \Pr(\mathbf{w} \in \mathbf{r}_{[m]}) \leq \frac{\mathcal{S}_\delta(\ell, m)}{2^\ell}, \quad (4)$$

and thus the expected number of unique strings in $\mathbf{r}_{[m]}$ is bounded between $\frac{1}{2}\mathcal{S}_\delta(\ell, m)$ and $\mathcal{S}_\delta(\ell, m)$.

Furthermore, $\mathcal{S}_\delta(\ell, m)$ takes the following values for different values of ℓ and m :

- If $\ell \geq \frac{\log m}{H(\delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \geq \frac{1}{4}m. \quad (5)$$

In particular if $\ell \geq \frac{\log m}{\log(\frac{1}{1-\delta})}$, then

$$\mathcal{S}_\delta(\ell, m) = m. \quad (6)$$

- If $\ell \leq \frac{\log m}{H(\frac{1}{2}, \delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \geq 2^{\ell-1}. \quad (7)$$

In particular if $\ell \leq \frac{\log m}{\log(\frac{1}{\delta})}$, then

$$\mathcal{S}_\delta(\ell, m) = 2^\ell.$$

- For any $\delta < \delta' < \frac{1}{2}$,

$$\mathcal{S}_\delta(\ell, m) \leq 2^{\ell H(\delta')} + m 2^{-\ell D(\delta' || \delta)}. \quad (8)$$

In particular if $\ell = \frac{\log m}{H(\delta', \delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \leq 2^{\ell H(\delta')} + m 2^{-\ell D(\delta' || \delta)} = 2^{\ell H(\delta')+1}. \quad (9)$$

- For any values of ℓ and m ,

$$\mathcal{S}_\delta(\ell, m) \leq \min(2^\ell, m). \quad (10)$$

The proof of Lemma 3 is presented in Appendix B.

A. Modified and adaptive fixed-length deduplication

We show that, even with knowledge of the source symbol length, if the chunk length is not properly chosen, mFLD encodes \mathbf{s} with a constant number of bits per symbol regardless of δ . Therefore, the ratio $\frac{\mathbb{E}[\mathcal{L}_{mF}(\mathbf{s})]}{H(\mathbf{s})}$ can be arbitrarily large for small δ . Meanwhile for AFLD, with the adaptive chunk length $\ell = \lceil \frac{\log(B/A)}{H(\gamma, \delta)} \rceil$, the ratio $\frac{\mathbb{E}[\mathcal{L}_{AF}(\mathbf{s})]}{H(\mathbf{s})}$ is shown to be upper bounded by a constant for all δ and for γ properly chosen.

Consider the two-stage parsing of \mathbf{s} with $D = L$. The length- D segments after the first-stage parsing are exactly the source blocks Y_1, Y_2, \dots, Y_B . Let $C = \lfloor L/\ell \rfloor$ and $r = L - C\ell$. Each Y_b , $1 \leq b \leq B$, is then parsed into chunks $Z_1^b, Z_2^b, \dots, Z_{C+1}^b$ with $|Z_c^b| = \ell$ for all $c \leq C$ and $|Z_{C+1}^b| = r$ (see Figure 2). If we also divide each source symbol X_a into substrings of length ℓ as $X_a = U_1^a U_2^a \dots U_{C+1}^a$, then for all $1 \leq c \leq C+1$, $\{Z_c^b\}_{b \in Y(a)}$ are iid δ -edit descendants of U_c^a .

Before performing a detailed evaluation of the algorithm, let us first provide a rough analysis for a special case, which will provide some insights into the general problem. Suppose the alphabet \mathcal{X} only has a single symbol X of length L , whose

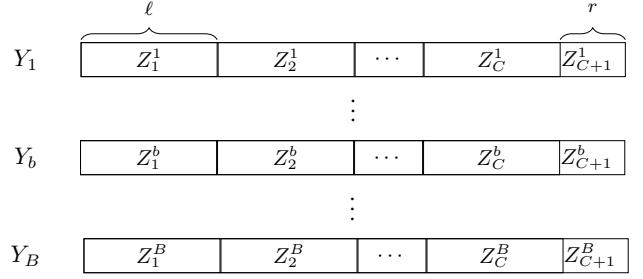


Figure 2. Modified fixed-length chunking with segment length $D = L$ and chunk length ℓ .

ℓ -prefix is denoted by U_1 . We consider encoding only the set $Z_1^1, Z_2^1, \dots, Z_1^B$, where each Z_1^b is a descendant of U_1 by δ -edit. The expected size of the dictionary, i.e., the number of distinct ℓ -strings in $\{Z_1^1, Z_2^1, \dots, Z_1^B\}$, by Lemma 3 is approximately

$$S := \mathcal{S}_\delta(\ell, B) = \sum_{t=0}^{\ell} \min\left(\binom{\ell}{t}, \binom{\ell}{t} B \delta^t (1-\delta)^{\ell-t}\right). \quad (11)$$

We can interpret (11) as follows. At a given distance t from U_1 , there are $\binom{\ell}{t}$ sequences of length ℓ . Further, if we generate B sequences, the expected number of sequences at distance t is $\binom{\ell}{t} B \delta^t (1-\delta)^{\ell-t}$. The number of sequences in the dictionary at distance t is then approximated by the minimum of the two terms. (This analysis of S is helpful whenever $\mathcal{S}_\delta(\cdot, \cdot)$ appears in the sequel as well.)

We would like S to be small enough that $\log S \ll \ell$ (so that pointers to the dictionary have much smaller lengths than the sequences being encoded) and $S \ll B$ (so that each sequence in the dictionary is repeated many times).² As t ranges from 0 to ℓ in the sum in (11), the term $\binom{\ell}{t}$ attain its maximum at $t \simeq \ell/2$ while the second term inside the min attains its maximum at $t \simeq \ell\delta$. We investigate which term determines the behavior of the sum. Let $\ell = \frac{\log B}{H(\gamma, \delta)}$ for a constant $0 \leq \gamma \leq 1$. Note that since $\delta < \frac{1}{2}$, $H(\gamma, \delta)$ and ℓ are increasing and decreasing functions of γ , respectively. With this choice, $B \delta^t (1-\delta)^{\ell-t} \geq 1$ for $t \leq \ell\gamma$ and $B \delta^t (1-\delta)^{\ell-t} \leq 1$ for $t \geq \ell\gamma$.

- If $\gamma < \delta$, then $B \delta^{\ell\delta} (1-\delta)^{(\ell-\delta)\ell} < 1$, and $S \geq \sum_{t=\lceil \gamma \ell \rceil}^{\ell} \binom{\ell}{t} B \delta^t (1-\delta)^{\ell-t} \geq B(1 - 2^{-\ell D(\gamma || \delta)})$. In this case, almost all Z_1^b are distinct and thus not compressible.
- If $\gamma = \delta$, then $\ell = \frac{\log B}{H(\delta)}$, and $S \geq \frac{B}{4}$ by (5). In this case, a constant fraction of Z_1^b are distinct and thus not compressible.
- If $\gamma \geq 1/2$, then $\ell \leq \frac{\log B}{H(\frac{1}{2}, \delta)}$, and $S \geq 2^{\ell-1}$ by (7). In this case, due to the fact that ℓ is chosen too small, the dictionary is so large that pointers to the dictionary are as long as the chunks and there is no compression gain.
- If $\delta < \gamma < 1/2$, then by (9),

$$S \leq 2^{\ell H(\gamma)+1}.$$

²Note that the size of the dictionary, and hence the length of the pointers, vary as the encoding progresses; we ignore this fact for now and approximate pointer lengths based on the final size of the dictionary.

Hence, pointers have an approximate length of $\ell H(\gamma)$ and are smaller than ℓ by a factor of $\frac{1}{H(\gamma)}$. Furthermore, each sequence is repeated approximately $2^{\ell D(\gamma||\delta)}$ times since $B = 2^{\ell H(\gamma,\delta)}$. The number of bits required to encode the dictionary is $2\ell 2^{\ell H(\gamma)}$, which is negligible compared to $B\ell$, the length of the uncoded sequences since $\gamma \neq \delta$. Hence, we can encode $\{Z_1^1, \dots, Z_1^B\}$ using essentially $B\ell H(\gamma)$ bits, achieving a compression ratio of $\frac{1}{H(\gamma)}$.

This analysis highlights that ℓ should be chosen appropriately to avoid a large dictionary or a situation in which there are no repetitions in the sequence. If these conditions are satisfied, then we can successfully deduplicate the data, as shown rigorously in Theorem 8 for AFLD.

Now we return to the general setting. It can be seen from the description of mFLD that the compressed string is composed of two parts: the bits used to encode the chunks at their first occurrences and the bits used to encode repeated chunks by pointers to the dictionary. For both parts, our first step is to compute the expected size of the dictionary, i.e., the number of distinct chunks, for which we present Lemma 4 and Lemma 5.

Lemma 4. *Suppose K strings of length n are chosen independently and uniformly from Σ^n . Assume each string produces at least m_1 and at most m_2 descendants by δ -edits. For any string w with $|w| = n$, let G_w denote the event that w equals one or more descendants. Then*

$$\begin{aligned} \frac{1}{2} \min\left(1, \frac{1}{2} K \frac{\mathcal{S}_\delta(n, m_1)}{2^n}\right) \\ \leq \Pr(G_w) \leq \\ \min\left(1, K \frac{\mathcal{S}_\delta(n, m_2)}{2^n}\right). \end{aligned}$$

The proof of Lemma 4 is presented in Appendix C. This lemma considers the probability of observing a string w when multiple random strings produce δ -edit descendants simultaneously. This setting models exactly our source string generation process where the A source symbols correspond to K random strings, and the source blocks correspond to the δ -edit descendants. In particular, \mathcal{E}_u being true corresponds to $m_2 = \frac{3B}{2A}$ and \mathcal{E}_l being true corresponds to $m_1 = \frac{B}{4A}$.

Let $T_F^1(s)$ denote the dictionary after all chunks of s are processed, i.e., $T_F^1(s)$ contains all distinct strings in $\{Z_c^b\}_{b,c}$. Let $T_F^{1/2}(s)$ denote the dictionary immediately after all chunks in the first half of s , i.e., $Y_1 Y_2 \dots Y_{\lceil B/2 \rceil}$, are processed. We apply Lemma 4 to find bounds on the sizes of $T_F^1(s)$ and $T_F^{1/2}(s)$ in the following lemma.

Lemma 5. *Consider the two-stage fixed-length chunking process with first-stage parsing length $D = L$ and chunk length ℓ . The dictionary sizes $T_F^1(s)$ and $T_F^{1/2}(s)$ satisfy*

$$\begin{aligned} \mathbb{E}[|T_F^1(s)| | \mathcal{E}_u] &\leq \min\left(2^\ell, ACS_\delta\left(\ell, \frac{3B}{2A}\right)\right) + B, \\ \mathbb{E}[|T_F^{1/2}(s)| | \mathcal{E}_l] &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} ACS_\delta\left(\ell, \frac{B}{4A}\right)\right). \end{aligned} \quad (12)$$

The proof of Lemma 5 is presented in Appendix C.

Next, we show using Lemma 5 that if ℓ is chosen too small relative to the scale of the system, then mFLD spends a constant number of bits per symbol. The proof strategy is as follows: with ℓ small enough, the term $\min(2^\ell, \frac{1}{2} ACS_\delta(\ell, \frac{B}{4A}))$ in (12) equals 1, which makes $\mathbb{E}[|T_F^{1/2}(s)| | \mathcal{E}_l]$ greater than $2^{\ell-1}$. Therefore, when encoding duplicated chunks using pointers, each pointer takes approximately ℓ bits and there is no compression gain.

Theorem 6. *Consider the source model in which source symbols have the same length L . For mFLD with first-stage parsing length $D = L$ and chunk length ℓ , if $\ell 2^\ell = O(AL)$ or $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$, then*

$$\mathbb{E}[\mathcal{L}_{mF}(s)] \geq \frac{1}{12} BL(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

where the $o(1)$ term is independent of δ .

Proof: We first claim, to be proved later, that if $\ell 2^\ell = O(AL)$ or $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$, then

$$\mathbb{E}[|T_F^{1/2}(s)| | \mathcal{E}_l] \geq 2^{\ell-1}. \quad (13)$$

It follows from Markov's inequality that

$$\Pr\left(2^\ell - |T_F^{1/2}(s)| \geq \frac{3}{4} \cdot 2^\ell | \mathcal{E}_l\right) \leq \frac{\frac{1}{2} \cdot 2^\ell}{\frac{3}{4} \cdot 2^\ell} = \frac{2}{3},$$

which is equivalent to

$$\Pr\left(|T_F^{1/2}(s)| \geq \frac{2^\ell}{4} | \mathcal{E}_l\right) \geq \frac{1}{3}. \quad (14)$$

Next, we consider the second half of s , $Y_{\lceil B/2 \rceil+1} \dots Y_B$. There are $\lfloor B/2 \rfloor C$ chunks of length ℓ , and encoding each of them takes at least either ℓ or $\log|T_F^{1/2}(s)|$ bits plus an additional bit indicating whether the chunk is stored in full or represented by a pointer. So in total, we need at least

$$\left(\min\left(\ell, \log|T_F^{1/2}(s)|\right) + 1\right) \cdot \left\lfloor \frac{B}{2} \right\rfloor C$$

bits. It follows that for B sufficiently large,

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{mF}(s) | \mathcal{E}_l] &\geq \mathbb{E}\left[\left(\min\left(\ell, \log|T_F^{1/2}(s)|\right) + 1\right) \cdot \left\lfloor \frac{B}{2} \right\rfloor C | \mathcal{E}_l\right] \\ &\geq \frac{1}{3} \left(\min\left(\ell, \log \frac{2^\ell}{4}\right) + 1\right) \cdot \left\lfloor \frac{B}{2} \right\rfloor C \\ &\geq \frac{BL}{12} (1 + o(1)), \end{aligned}$$

where the second inequality follows from (14).

Finally, since (2) gives that $\Pr(\mathcal{E}_l) = 1 + o(1)$, we get

$$\mathbb{E}[\mathcal{L}_{mF}(s)] \geq \mathbb{E}[\mathcal{L}_{mF}(s) | \mathcal{E}_l] \Pr(\mathcal{E}_l) \geq \frac{BL}{12} (1 + o(1)).$$

It remains to prove the claim: $\mathbb{E}[|T_F^{1/2}(s)| | \mathcal{E}_l] \geq 2^{\ell-1}$ when $\ell 2^\ell = O(AL)$ or $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$. Consider the case when

$\ell 2^\ell = O(AL)$. For sufficiently large B (and thus A and L), $\frac{B}{4A} \geq \frac{4\ell 2^\ell}{AL}$. Therefore, by Lemma 5,

$$\begin{aligned}\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} A C S_\delta\left(\ell, \frac{B}{4A}\right)\right) \\ &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} A C S_\delta\left(\ell, \frac{4\ell 2^\ell}{AL}\right)\right),\end{aligned}$$

where the last inequality follows from the fact that $S_\delta(\ell, m)$ is non-decreasing in m . By (6), if $m(1 - \delta)^\ell \leq 1$, then $S_\delta(\ell, m) = m$. Since asymptotically $\frac{4\ell 2^\ell}{AL}(1 - \delta)^\ell \leq 1$,

$$\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] \geq 2^{\ell-1} \min\left(1, \frac{4\ell C}{2L}\right) \geq 2^{\ell-1},$$

where the last step follows from the fact that $C \geq \frac{L}{2\ell}$.

When $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$, again by Lemma 5,

$$\begin{aligned}\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} A C S_\delta\left(\ell, \frac{B}{4A}\right)\right) \\ &\geq 2^{\ell-1} \min\left(1, \frac{AC}{4}\right) \geq 2^{\ell-1},\end{aligned}$$

where the second inequality follows from (7) that when $\ell \leq \frac{\log m}{H(\frac{1}{2}, \delta)}$, $S_\delta(\ell, m) \geq 2^{\ell-1}$. ■

The preceding theorem shows that when ℓ is chosen too small, the size of the dictionary will be of order 2^ℓ . Specifically, if $\ell 2^\ell = O(AL)$, the number of distinct ℓ -substrings in the source alphabet is already of order 2^ℓ . If $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$, then the δ -edits are able to produce almost all ℓ -strings instead of only producing strings that are on the $\delta\ell$ Hamming sphere.

In the next theorem, we show that if ℓ is chosen too large, then mFLD again spends a constant number of bits per symbol. The proof strategy is to show that if ℓ is chosen too large, then almost every chunk is distinct, thus making the source string incompressible.

Theorem 7. Consider the source model in which source symbols have the same length L . For mFLD with first-stage parsing length $D = L$ and chunk ℓ , if $\ell \geq \frac{\log(B/A)-2}{H(\delta)}$, then

$$\mathbb{E}[\mathcal{L}_{mF}(\mathbf{s})] \geq \frac{1}{128} BL(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

where the $o(1)$ term is independent of δ .

Proof: When $\ell \leq \frac{\log(B/A)-2}{H(\frac{1}{2}, \delta)}$,

$$\begin{aligned}\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} A C S_\delta\left(\ell, \frac{B}{4A}\right)\right) \\ &\geq 2^{\ell-1} \min\left(1, \frac{1}{2} A C \cdot \frac{1}{4} \cdot \frac{B}{4A2^\ell}\right) \\ &= 2^{\ell-1} \min\left(1, \frac{BC}{32 \cdot 2^\ell}\right),\end{aligned}$$

where the first inequality follows from Lemma 5 and the second from (5).

In the case where $1 \leq \frac{BC}{32 \cdot 2^\ell}$ and hence $\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] \geq 2^{\ell-1}$, the proof follows from the discussion that follows (13). So it remains to consider the case when $\frac{BC}{32 \cdot 2^\ell} \leq 1$, i.e.,

$$\mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] \geq 2^{\ell-1} \cdot \frac{BC}{32 \cdot 2^\ell} = \frac{BC}{64}.$$

Since it takes $\ell + 1$ bits to store distinct chunks in the dictionary,

$$\begin{aligned}\mathbb{E}[\mathcal{L}_{mF}(\mathbf{s}) \mid \mathcal{E}_l] &\geq (\ell + 1) \mathbb{E}\left[\left|T_F^{1/2}(\mathbf{s})\right| \middle| \mathcal{E}_l\right] = \ell \frac{B \lfloor L/\ell \rfloor}{64} \\ &\geq \frac{1}{64} B \max(\ell, L - \ell) \geq \frac{1}{128} BL.\end{aligned}$$

The desired result thus follows again from $\mathbb{E}[\mathcal{L}_{mF}(\mathbf{s})] \geq \mathbb{E}[\mathcal{L}_{mF}(\mathbf{s}) \mid \mathcal{E}_l] \Pr(\mathcal{E}_l)$ and the fact that $\Pr(\mathcal{E}_l) = 1 + o(1)$. ■

In the Results section, Theorems 6 and 7 imply Corollary 2, which shows that choosing ℓ in $o(\log B)$ or $\omega(\log B)$ results in poor performance, and Corollary 3, which shows that FLD cannot compress the sequences effectively.

Next, we show that with the adapted chunk length, AFLD can achieve performance within a constant factor of optimal.

Theorem 8. Consider the source model in which source symbols have the same length L . The performance of AFLD with $D = L$ and $\ell = \left\lceil \frac{\log(B/A)}{H(\gamma, \delta)} \right\rceil$ satisfies

$$1 \leq \frac{\mathbb{E}[\mathcal{L}_{AF}(\mathbf{s})]}{H(\mathbf{s})} \leq \frac{1 + k_1}{k_2} \cdot \frac{H(\gamma, \delta)}{H(\delta)} \cdot (1 + o(1)),$$

as $B \rightarrow \infty$, for any $\gamma \in (\delta, \frac{1}{2})$.

Proof: We first note that the length of \mathbf{s} can be encoded in at most $2 \log(|\mathbf{s}|) + 3$ bits with Elias gamma coding.

The number of bits used to encode chunks at their first occurrences is upper bounded by $|T_F^1(\mathbf{s})|(\ell + 1)$ since chunks are all of lengths less than or equal to ℓ . Consider the upper bound on $\mathbb{E}[|T_F^1(\mathbf{s})| \mid \mathcal{E}_u]$ in Lemma 5. Note that by (8) and $\frac{B}{A} \leq 2^{\ell H(\gamma, \delta)}$ with our choice of ℓ ,

$$S_\delta\left(\ell, \frac{3B}{2A}\right) \leq 2^{\ell H(\gamma)} + \frac{3B}{2A} 2^{-\ell D(\gamma \mid \delta)} \leq \frac{5}{2} \cdot 2^{\ell H(\gamma)}.$$

It follows that

$$\begin{aligned}\mathbb{E}[|T_F^1(\mathbf{s})| \mid \mathcal{E}_u](\ell + 1) &\leq \left(\min\left(2^\ell, A C S_\delta\left(\ell, \frac{3B}{2A}\right)\right) + B \right) (\ell + 1) \\ &\leq \min\left(2^\ell, \frac{5AC}{2} \cdot 2^{\ell H(\gamma)}\right) (\ell + 1) + B(\ell + 1) \\ &\leq \frac{5AL}{2\ell} \cdot 2^{\ell H(\gamma)} \cdot (\ell + 1) + B(\ell + 1) \\ &= \frac{5}{2} AL \left(\frac{B}{A}\right)^{H(\gamma)/H(\gamma, \delta)} \left(1 + \Theta\left(\frac{1}{\log(B/A)}\right)\right) \\ &\quad + \Theta(B \log B) \\ &= o(BL),\end{aligned}\tag{15}$$

where the last equality follows from $\frac{H(\gamma)}{H(\gamma, \delta)} < 1$ and thus $B^{\frac{H(\gamma)}{H(\gamma, \delta)}} A^{1 - \frac{H(\gamma)}{H(\gamma, \delta)}} = o(B)$.

Next, we derive an upper bound on the number of bits used by pointers for encoding repeated chunks. There are $(C+1)B$ chunks and the number of bits needed for encoding one pointer

is at most $\log(BL) + 1$. So in total, the number of bits we need is at most

$$\begin{aligned} (C + 1)B(\log(BL) + 1) &\leq (L + \ell)B \frac{\log(BL) + 1}{\ell} \\ &\leq \frac{BL}{\ell} \log(BL) \left(1 + O\left(\frac{1}{\log B}\right)\right) \\ &\leq H(\gamma, \delta)BL \cdot \frac{\log(BL)}{\log(B/A)}. \end{aligned} \quad (16)$$

Combining (15), (16), and including the number of bits used for encoding the length of s by Elias coding, we get

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{AF}(s)|\mathcal{E}_u] &\leq H(\gamma, \delta)BL \cdot \frac{\log(BL)}{\log(B/A)} + o(BL) \\ &\leq H(\gamma, \delta)BL \frac{1 + k_1}{k_2} (1 + o(1)), \end{aligned}$$

by noting that $\frac{\log(BL)}{\log(B/A)} \leq \frac{1 + k_1}{k_2} (1 + o(1))$.

On the complement of \mathcal{E}_u , the number of bits needed for storing the dictionary is at most $2BL$ since the lengths of chunks in total is at most BL and there are at most BL chunks. The number of bits for encoding repeated chunks by pointers is at most $BL(\log(BL) + 1)$. It follows that

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{AF}(s)|\bar{\mathcal{E}}_u] \Pr(\bar{\mathcal{E}}_u) \\ \leq (2ABL + 2\log(BL) + 3) \log(BL) e^{-\frac{B}{10A}} = o(1). \end{aligned}$$

The desired result thus follows from

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{AF}(s)] \\ = \mathbb{E}[\mathcal{L}_{AF}(s)|\mathcal{E}_u] \Pr(\mathcal{E}_u) + \mathbb{E}[\mathcal{L}_{AF}(s)|\bar{\mathcal{E}}_u] \Pr(\bar{\mathcal{E}}_u), \end{aligned}$$

and the fact that $\Pr(\mathcal{E}_u) = 1 + o(1)$. \blacksquare

For any $\delta < \frac{1}{2}$ and $a > 1$, we can find γ in the range $(\delta, \frac{1}{2})$ such that $H(\gamma, \delta)/H(\delta) \leq a$. It thus follows from Theorem 8 that adaptive fixed-length deduplication can compress the sequence within a constant factor of the entropy, as stated in Corollary 4 in the Results section.

B. Edit-distance deduplication

Next, we study the edit-distance deduplication algorithm. EDD identifies positions in which the current chunk and previously observed similar chunks differ. We show that with chunk length being equal to source symbol length, EDD can achieve a constant factor of optimal.

Theorem 9. *Consider the source model in which source symbols have the same length L and the edit probability is $\delta < \frac{1}{4}$. The performance of edit-distance deduplication with chunk length $\ell = L$ and mismatch ratio β satisfies*

$$1 \leq \frac{\mathbb{E}[\mathcal{L}_{ED}(s)]}{H(s)} \leq \frac{H(2\beta)}{H(\delta)} (1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

for any $\delta < \beta \leq \frac{1}{4}$.

Proof: With $\ell = L$, the B source blocks, Y_1, \dots, Y_B , are parsed as chunks. We know that each Y_b is a descendant of one of the source symbols. Let \mathcal{E}_d denote the event that every source block Y_b is within Hamming distance βL from

its ancestor. By the Chernoff bound, the probability that more than βL symbols of a source symbol are flipped in a δ -edit is at most $2^{-D(\beta||\delta)L}$. We then apply the union bound and get $\Pr(\mathcal{E}_d) \geq 1 - B2^{-D(\beta||\delta)L}$.

When \mathcal{E}_d holds, the source blocks are covered by A Hamming balls of radius βL . Therefore, with mismatch ratio β , the dictionary is of size at most A , and takes $A(L + 1)$ bits to store. The pointer length is thus upper bounded by $\log A + 1$. The difference with the referenced chunk can be encoded in at most $H(2\beta)L + 1$ bits. Including the $2\log(BL) + 3$ bits for encoding $|s|$ at the beginning, we get

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{ED}(s)|\mathcal{E}_d] &\leq 2\log(BL) + 3 + A(L + 1) \\ &\quad + (1 + \log A + 1 + H(2\beta)L + 1)B \\ &= H(2\beta)BL + o(BL). \end{aligned}$$

When the complement of \mathcal{E}_d holds, we trivially upper bound dictionary size by B . It follows that

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{ED}(s)|\bar{\mathcal{E}}_d] &\leq 2\log(BL) + 3 + B(L + 1) \\ &\quad + (1 + \log B + 1 + H(2\beta)L + 1)B \\ &\leq 2BL. \end{aligned}$$

Thus,

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{ED}(s)] &= \Pr(\mathcal{E}_d)\mathbb{E}[\mathcal{L}_{ED}(s)|\mathcal{E}_d] + \Pr(\bar{\mathcal{E}}_d)\mathbb{E}[\mathcal{L}_{ED}(s)|\bar{\mathcal{E}}_d] \\ &\leq H(2\beta)BL(1 + o(1)) + 2B^2L2^{-D(\beta||\delta)L} \\ &= H(2\beta)BL(1 + o(1)), \end{aligned}$$

where the term $2B^2L2^{-D(\beta||\delta)L}$ is absorbed into the $o(1)$ term since $D(\beta||\delta) > 0$. \blacksquare

The theorem is used in the Results section to establish that EDD performs within a constant factor of entropy in Corollary 5.

VII. DEDUPLICATION IN THE VARIABLE-LENGTH SCHEME

In this section, we study the variable-length deduplication algorithm, which is more widely applicable than the algorithms in the fixed-length scheme and does not require the source symbol lengths to be the same or known. In the previous section, we saw that for AFLD to achieve optimality, the chunk length should be adapted to the source. Similarly for VLD, the performance depends on chunk lengths which in turn depend on the length of the marker M .

Before presenting the detailed analysis, we provide some insights on how the marker length M affects the distribution of chunk contents. In variable-length chunking, the chunks (except perhaps the last one) end with the marker string 0^M . We write $s = U_1 0^M U_2 0^M \dots 0^M U_N$, where each $U_n, n < N$, is either empty or of the form $u1$ for some M -RLL string u . We can approximately treat s as a Bernoulli(1/2) process for now. The lengths of strings U_n are thus equivalent to the stopping time in an infinite-length Bernoulli(1/2) process until the beginning of the first occurrence of 0^M , which is of expected length approximately 2^M . The behavior of VLD with marker length M is thus similar to that of mFLD with chunk length 2^M . When M is chosen so small that the number N of chunks becomes much larger than the total number of strings

of lengths around 2^M , the dictionary becomes exhaustive and pointers have similar lengths to chunks. When M is chosen too large, most of U_1, \dots, U_N are distinct and thus not compressible. In the following, we study in detail how $\mathbb{E}[\mathcal{L}_{VL}(s)]$ varies for different values of M .

Similar to the fixed-length schemes, the dictionary size is an essential first-step in computing $\mathbb{E}[\mathcal{L}_{VL}(s)]$. To determine the expected dictionary size, we again start with the probability of occurrences of chunks. However, now the chunks are of different lengths and the occurrences are not restricted to a fixed set of positions. So we bound the probability of occurrences of a chunk by the probability of occurrences of certain substrings. Specifically, we consider strings of the forms $10^M u 10^M$ or $0^M u 10^M$ ($u \in R_M$): Except the first and the last chunks, the probability of occurrence of chunk $u 10^M$ is greater than the probability of occurrence of a substring $10^M u 10^M$ since the prefix 10^M always marks an ending of the previous chunk; similarly, the probability of occurrence of chunk $u 10^M$ is less than or equal to the probability of occurrence of a substring $0^M u 10^M$ since any occurrences of chunk $u 10^M$ must follow a 0^M which is the ending marker of the previous chunk.

Let $w \in Y_1^B$ denote the event that w appears as a substring of Y_b for some $1 \leq b \leq B$ and let $w \in Y_1^{B/2}$ denote the event that w appears as a substring of Y_b for some $1 \leq b \leq \lceil B/2 \rceil$ ³. We first present in Lemmas 10, 11 and 12 two lower bounds on $w \in Y_1^{B/2}$ and an upper bound on $w \in Y_1^B$.

Lemma 10. Suppose K strings of length n are chosen independently and uniformly from Σ^n . Assume each string produces at least m_1 and at most m_2 descendants by δ -edits. For any string w with $|w| \leq n$, let H_w denote the event that w appears as a substring of one or more descendants. Then,

$$\begin{aligned} \frac{1}{2} \min\left(1, \frac{1}{2} \left\lfloor \frac{n}{|w|} \right\rfloor K \frac{\mathcal{S}_\delta(|w|, m_1)}{2^{|w|}}\right) \\ \leq \Pr(H_w) \leq \\ \min\left(1, (n - |w| + 1) K \frac{\mathcal{S}_\delta(|w|, m_2)}{2^{|w|}}\right). \end{aligned}$$

The proof of Lemma 10 is presented in Appendix D. Similar to Lemma 4, the setting described in Lemma 10 matches the model for the generation of source strings. This time, we allow string w to be any substring of the descendants because chunks can now be in any position of the source string. Note that Lemma 10 is also a generalization of Lemma 4.

Next, we use Lemma 10 to bound the probability of $w \in Y_1^{B/2}$ and $w \in Y_1^B$.

Lemma 11. Consider the source model with edit probability δ . For any string $w \in \Sigma^*$ with $|w| \leq 2L$,

$$\Pr(w \in Y_1^B | \mathcal{E}_u) \leq \min\left(1, 2AL \frac{\mathcal{S}_\delta(|w|, \frac{3B}{2A})}{2^{|w|}}\right).$$

³Here we only consider string/chunk occurrences inside source blocks and leave the study of strings/chunks that occur across the boundaries of source blocks for later.

For any string $w \in \Sigma^*$ with $|w| \leq \lceil \frac{1}{2}L \rceil$,

$$\Pr(w \in Y_1^{B/2} | \mathcal{E}_l) \geq \frac{1}{2} \min\left(1, \frac{AL}{8|w|} \frac{\mathcal{S}_\delta(|w|, \frac{B}{4A})}{2^{|w|}}\right). \quad (17)$$

The proof of Lemma 11 is presented in Appendix D. Although Lemma 11 holds for any string w , we will later restrict w to be of the forms $10^M u 10^M$ or $0^M u 10^M$.

Next, we consider another lower bound as an alternative to (17) for the cases when w is of larger lengths. From the proofs of Lemmas 10 and 11, the lower bound (17) is obtained by only taking into account the possibilities of w appearing in non-overlapping positions of each Y_b . Lemma 12 considers every possible substring of Y_b to be equal to w and gets the lower bound by the inclusion-exclusion principle and turns out to be more accurate for w with large lengths. Note that Lemma 12 directly considers w to be of the form $10^M u 10^M$ and the bound is given in the form of a summation.

Lemma 12. Consider the source model with edit probability $\delta < \frac{1}{2}$. For any n such that $\frac{\log(B/A)-2}{H(\delta)} \leq n + 2M + 2 \leq \frac{L}{4}$,

$$\begin{aligned} \sum_{u \in R_M^n} \Pr(10^M u 10^M \in Y_1^{B/2} | \mathcal{E}_l) \geq \\ \frac{BL}{27 \cdot 2^{2M+2}} \cdot \left(1 - \frac{1}{2^{M-1}}\right)^n - \frac{3B^2 L^2}{2^{n+2M+2}}. \end{aligned}$$

The proof of Lemma 12 is presented in Appendix E.

After characterizing the probabilities of strings (and thus chunks) occurring, we consider in Lemma 13 the number of chunks. Let $C_{VL}^M(s)$ denote the number of chunks of length over 2^{M-4} in $Y_{\lceil B/2 \rceil + 1} \dots Y_B$ for variable-length chunking with marker length M . We show that when $2^M = o(L)$, with high probability, $C_{VL}^M(s)$ is of order $|s|/2^M$.

Lemma 13. Consider the source string $s = Y_1 Y_2 \dots Y_B$. When $2^M = o(L)$, for B, L sufficiently large,

$$\Pr\left(C_{VL}^M(s) \geq \frac{1}{4} \cdot \left\lfloor \frac{B}{2} \right\rfloor \left(\frac{L}{2^{M+8}} - 1\right)\right) \geq \frac{5}{6}.$$

The proof of Lemma 13 is presented in Appendix F. It can be seen from the proof that Lemma 13 can be extended to the case when \mathcal{E}_l holds since each source block Y_b by itself is still a Bernoulli(1/2) process. Therefore, the following corollary holds.

Corollary 7. When $2^M = o(L)$, for B, L sufficiently large,

$$\Pr\left(C_{VL}^M(s) \geq \frac{1}{4} \cdot \left\lfloor \frac{B}{2} \right\rfloor \left(\frac{L}{2^{M+8}} - 1\right) | \mathcal{E}_l\right) \geq \frac{5}{6}.$$

Next, we use Lemmas 11, 12 and Corollary 7 to bound $\mathbb{E}[\mathcal{L}_{VL}(s)]$ from below. As marker length M takes different values, different lower bounds of $\mathbb{E}[\mathcal{L}_{VL}(s)]$ are presented in Theorems 14, 15 and 17. Let $T_{VL}^1(s)$ denote the dictionary when all chunks in s are processed and let $T_{VL}^{1/2}(s)$ denote the dictionary immediately after chunks in $Y_1 \dots Y_{\lceil B/2 \rceil}$ are processed.

We first show in Theorem 14 that similar to the fixed-length schemes, small values for M lead to an oversized dictionary.

Theorem 14. Consider the source model with edit probability δ and the variable-length deduplication algorithm with marker length M . If $2^M = o(\log B)$, then

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] \geq \frac{1}{3 \cdot 2^{16}} BL(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

where the $o(1)$ term is independent of δ .

Proof: We show that with high probability, $|T_{VL}^{1/2}(\mathbf{s})|$ is of the order 2^{2^M} . So encoding each chunk in $Y_{\lceil B/2 \rceil + 1} \cdots Y_B$ takes number of bits either equal to the chunk length or pointer length 2^M . We then show using Lemma 13 that the length of the compressed string is a constant fraction of BL .

If a string \mathbf{w} of the form $\mathbf{w} = 10^M \mathbf{u} 10^M$, $\mathbf{u} \in R_M$, occurs as a substring of some data block Y_b , $b \leq \lceil \frac{B}{2} \rceil$, then $\mathbf{u} 10^M$ must be contained in $T_{VL}^{1/2}(\mathbf{s})$. For any $\mathbf{w} = 10^M \mathbf{u} 10^M$ with $|\mathbf{u}| \leq 2^M$, by Lemma 11,

$$\begin{aligned} \Pr(\mathbf{w} \in Y_1^{B/2} | \mathcal{E}_l) &\geq \frac{1}{2} \min\left(1, \frac{AL}{8|\mathbf{w}|} \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}}\right) \\ &\geq \frac{1}{2} \min\left(1, \frac{AL}{8|\mathbf{w}|}\right) \geq \frac{1}{2}, \end{aligned} \quad (18)$$

where the second inequality follows from $|\mathbf{w}| \leq 2^M + 2M + 2 = o(\log B)$ and the property that $\mathcal{S}_\delta(\ell, m) = 2^\ell$ if $m\delta^\ell \geq 1$.

Denote the set of all M -RLL strings of lengths less than 2^M by $R_M^{\leq 2^M}$. Let $\zeta = \left| \left\{ \mathbf{u} \in R_M^{\leq 2^M} : 10^M \mathbf{u} 10^M \in Y_1^{B/2} \right\} \right|$. Then (18) gives $\mathbb{E}[\zeta | \mathcal{E}_l] \geq |R_M^{\leq 2^M}|/2$ and thus $\mathbb{E}[|R_M^{\leq 2^M}| - \zeta | \mathcal{E}_l] \leq \frac{|R_M^{\leq 2^M}|}{2}$. By Markov inequality, $\Pr(|R_M^{\leq 2^M}| - \zeta \geq 3|R_M^{\leq 2^M}|/4) \leq \frac{2}{3}$ and thus $\Pr(\zeta > |R_M^{\leq 2^M}|/4) \geq \frac{1}{3}$. Noting that $|T_{VL}^{1/2}| \geq \zeta$ and $|R_M^{\leq 2^M}| \geq 2^{2^M-2}$ by Corollary 1, we get

$$\Pr\left(|T_{VL}^{1/2}(\mathbf{s})| \geq 2^{2^M-4} | \mathcal{E}_l\right) \geq \frac{1}{3}. \quad (19)$$

For each chunk in $Y_{\lceil B/2 \rceil + 1} \cdots Y_B$ of length at least 2^{M-4} , we need at least either 2^{M-4} or $\log |T_{VL}^{1/2}(\mathbf{s})|$ bits. So by Corollary 7 and inequality (19),

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s}) | \mathcal{E}_l] \geq \mathbb{E}\left[\min\left(2^{M-4}, \log |T_{VL}^{1/2}(\mathbf{s})|\right) \cdot C_{VL}^M(\mathbf{s}) | \mathcal{E}_l\right]$$

$$\begin{aligned} &\geq \left(1 - \frac{2}{3} - \frac{1}{6}\right) \min(2^{M-4}, 2^M - 4) \\ &\quad \cdot \frac{1}{4} \left\lfloor \frac{B}{2} \right\rfloor \left(\frac{L}{2^{M+8}} - 1 \right) \\ &\geq \frac{BL}{3 \cdot 2^{16}} (1 + o(1)). \end{aligned}$$

The desired result follows from

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] \geq \mathbb{E}[\mathcal{L}_{VL}(\mathbf{s}) | \mathcal{E}_l] \Pr(\mathcal{E}_l)$$

and $\Pr(\mathcal{E}_l) = 1 + o(1)$. \blacksquare

We then show in Theorems 15 and 17 that an oversized M leads to a large number of distinct chunks, each of which needs to be encoded in full and thus compression becomes ineffective. In particular, Theorem 15 covers the case when 2^M is of larger order than $\log B$ but still much smaller than

the expected source symbol length L . Theorem 17 considers the case when $2^M = \Omega(L)$, and therefore a large number of chunks can be of lengths close to or even larger than the expected source symbol length.

Theorem 15. Consider the source model with edit probability δ and the variable-length deduplication algorithm with marker length M . If $2^M = \omega(\log B) \cap o(L)$, then

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] \geq \frac{1}{2^{10} e^2} BL(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

where the $o(1)$ term is independent of δ .

Proof: We show that if 2^M is in $\omega(\log B)$ and $o(L)$, the sum of the lengths of distinct chunks is a constant fraction of $|\mathbf{s}|$.

Each new chunk is encoded as a bit 1 followed by itself. Given \mathcal{E}_l , the expected number of bits needed for encoding distinct chunks is greater than or equal to

$$\begin{aligned} &\mathbb{E}\left[\sum_{\mathbf{v} \in T_{VL}^1(\mathbf{s})} (|\mathbf{v}| + 1) | \mathcal{E}_l\right] \\ &= \sum_{\mathbf{v} \in \Sigma^*} \Pr(\mathbf{v} \in T_{VL}^1(\mathbf{s}) | \mathcal{E}_l) (|\mathbf{v}| + 1) \\ &\geq \sum_{\mathbf{u} \in R_M} \Pr\left(10^M \mathbf{u} 10^M \in Y_1^{B/2} | \mathcal{E}_l\right) (|\mathbf{u}| + M + 2). \end{aligned} \quad (20)$$

As a lower bound, we consider M -RLL strings with lengths in the range $\left[2^M, \left\lceil (2^M L)^{1/2} \right\rceil\right]$. Since asymptotically we have $2^M \geq \frac{\log(B/A) - 2}{H(\delta)}$, we apply Lemma 12 on (20) and get

$$\begin{aligned} &\left[(2^M L)^{1/2} \right] \sum_{\ell=2^M} \sum_{\mathbf{u} \in R_M^\ell} \Pr\left(10^M \mathbf{u} 10^M \in Y_1^{B/2} | \mathcal{E}_l\right) \\ &\quad \cdot (\ell + M + 1) \\ &\geq \left[(2^M L)^{1/2} \right] \left(\frac{BL}{2^7 \cdot 2^{2M+2}} \left(1 - \frac{1}{2^{M-1}}\right)^\ell - \frac{3B^2 L^2}{2^{\ell+2M+2}} \right) \\ &\quad \cdot (\ell + M + 1) \\ &\geq \left[(2^M L)^{1/2} \right] \left(\frac{BL}{2^7 \cdot 2^{2M+2}} \left(1 - \frac{1}{2^{M-1}}\right)^\ell \ell \right) - \frac{3B^2 L^4}{2^{2M}} \\ &\geq \frac{BL}{27 \cdot 2^{2M+2}} 2^{2(M-1)} \left(\frac{2^M - 1}{2^{M-1}} + 1 \right) e^{-2} (1 + o(1)) \\ &\quad - \frac{3B^2 L^4}{2^{2M}} \\ &\geq \frac{BL}{2^{10} e^2} (1 + o(1)) - \frac{3B^2 L^4}{2^{2M}} \\ &= \frac{BL}{2^{10} e^2} (1 + o(1)), \quad \text{as } B \rightarrow \infty, \end{aligned}$$

where the second inequality follows from $(2^M L)^{1/2} + M + 1 \leq L$ and the equality follows from $2^M = \omega(\log B)$. The second to last inequality follows from applying summation (43) in

Appendix G-B with $a = 2^M, b = \lceil (2^M L)^{1/2} \rceil, \beta = 2^{M-1}$ and noting that $\frac{1}{2^{M-1}} \lceil (2^M L)^{1/2} \rceil = \omega(1)$.

Thus,

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})|\mathcal{E}_l] &\geq \mathbb{E} \left[\sum_{\mathbf{v} \in T_{VL}^1(\mathbf{s})} (|\mathbf{v}| + 1) |\mathcal{E}_l \right] \\ &\geq \frac{BL}{2^{10}e^2} (1 + o(1)), \end{aligned}$$

and the desired result follows from

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] \geq \mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})|\mathcal{E}_l] \Pr(\mathcal{E}_l)$$

and $\Pr(\mathcal{E}_l) = 1 + o(1)$. \blacksquare

Next, we present a lemma that will be used in the proof of Theorem 17.

Lemma 16. Consider the source string $\mathbf{s} = Y_1 Y_2 \dots Y_B$, with each Y_b being a descendant of source symbol X_{J_b} . For any integer h and any pairs of integers $(b_1, b_2), (i_1, i_2)$, the probability of Y_{b_1} and Y_{b_2} having identical substrings of length h starting at positions i_1 and i_2 , respectively, is

$$\Pr((Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h}) = \frac{1}{2^h},$$

if $J_{b_1} \neq J_{b_2}$ or $i_1 \neq i_2$.

The proof of Lemma 16 is presented in Appendix F.

Theorem 17. Consider the source model with edit probability δ and the variable-length deduplication algorithm with marker length M . If $2^M = \Omega(L)$, then

$$\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] \geq \frac{1}{360} BL(1 + o(1)), \quad \text{as } B \rightarrow \infty,$$

where the $o(1)$ term is independent of δ .

Proof: Let $q = \min(2^{M-5}, L/2)$. We find a set of distinct M -RLL q -substrings of \mathbf{s} that are encoded in full. In other words, any two such q -substrings are contained in two distinct chunks, or in two chunks that are duplicates, or in a single chunk without overlapping with each other. The total length of these q -substrings thus provides a lower bound on $\mathcal{L}_{VL}(\mathbf{s})$.

Let L_1, \dots, L_A be given and assume \mathcal{E}_l holds. We consider the first $\lceil B/(4A) \rceil$ descendants of each source symbol. Let G_a denote the set of the first $\lceil B/(4A) \rceil$ descendants of X_a . Let Q_a be the set containing all non-overlapping q -substrings of G_a , i.e., $Q_a = \{\mathbf{x}_{1+(c-1)q, q} : \mathbf{x} \in G_a, 1 \leq c \leq c_a\}$, where $c_a = \lfloor L_a/q \rfloor$ and let $Q = \bigcup_{a=1}^A Q_a$. For $\mathbf{w} \in \Sigma^q$, let $\mathbf{w} \in Q$ denote the event that one of the substrings in Q equals \mathbf{w} . Applying Lemma 4 on Q (with substring length equal to descendant length) yields

$$\begin{aligned} \Pr(\mathbf{w} \in Q) &\geq \frac{1}{2} \min \left(1, \frac{1}{2} \left(\sum_{a=1}^A c_a \right) \frac{\mathcal{S}_\delta(q, \lceil \frac{B}{4A} \rceil)}{2^q} \right) \\ &= \frac{1}{4} \left\lceil \frac{B}{4A} \right\rceil \frac{\sum_{a=1}^A c_a}{2^q}, \end{aligned}$$

where the equality follows from $q = \Omega(L)$ and the property that $\mathcal{S}_\delta(\ell, m) = m$ if $m(1 - \delta)^\ell \leq 1$. So the expected number of distinct M -RLL strings in Q is at least

$$\begin{aligned} \sum_{\mathbf{w} \in R_M^q} \frac{1}{4} \left\lceil \frac{B}{4A} \right\rceil \frac{\sum_{a=1}^A c_a}{2^q} &\geq \frac{1}{4} \left(2 - \frac{1}{2^{M-2}} \right)^q \left\lceil \frac{B}{4A} \right\rceil \frac{\sum_{a=1}^A c_a}{2^q} \\ &\geq \frac{1}{5} \cdot \left\lceil \frac{B}{4A} \right\rceil \sum_{a=1}^A c_a, \end{aligned}$$

for all $M > 5$. Since the size of Q is $\lceil \frac{B}{4A} \rceil \sum_{a=1}^A c_a$, by the Markov bound, with probability at least $\frac{1}{9}$, the number of distinct M -RLL q -strings in Q is at least $\frac{1}{10} \lceil \frac{B}{4A} \rceil \sum_{a=1}^A c_a$.

Let $q' = \lceil q/2 \rceil$. Consider the q' -substrings of source blocks Y_1, \dots, Y_B , i.e., $(Y_b)_{i, q'}$ for all $b \in [B], i \in [|Y_b|]$. Define \mathcal{E}_d to be the following event: for every two source blocks Y_{b_1} and Y_{b_2} , the substring of Y_{b_1} starting at position i_1 is different from the substring of Y_{b_2} starting at position i_2 , i.e., $(Y_{b_1})_{i_1, q'} \neq (Y_{b_2})_{i_2, q'}$, as long as $J_{b_1} \neq J_{b_2}$ or $i_1 \neq i_2$. Since there are at most $(2BL)^2$ pairs of such substrings, by the union bound and Lemma 16, \mathcal{E}_d holds with probability at least

$$1 - (2BL)^2/2^{q'}.$$

When \mathcal{E}_d holds, the distinct M -RLL q -substrings in Q are then non-overlapping substrings of the dictionary and it takes q -bits to encode each of them. To see this, we consider the first time such q -strings appear in the source string. Let $(Y_b)_{j, q}$ be one of the M -RLL strings in Q . Given \mathcal{E}_d , the only possible substrings of \mathbf{s} that equal $(Y_b)_{j, q}$ are $(Y_1)_{k, q}, \dots, (Y_B)_{k, q}$. Let b' be the smallest integer such that $(Y_{b'})_{j, q} = (Y_b)_{j, q}$. By the M -RLL property, $(Y_{b'})_{j, q}$ must be fully contained in a chunk. Moreover, this chunk must be a new chunk by the minimality of b' and is entered into the dictionary. Similarly, every distinct M -RLL q -substring corresponds to a q -substring in the dictionary. Since strings in Q do not overlap with each other, the corresponding q -substrings in the dictionary also do not overlap, and each takes q bits to store.

Combining the two arguments, with probability at least $\frac{1}{9} - \frac{(2BL)^2}{2^{q'}} = \frac{1}{10} \lceil \frac{B}{4A} \rceil \sum_{a=1}^A c_a$, there are $\frac{1}{10} \lceil \frac{B}{4A} \rceil \sum_{a=1}^A c_a$ distinct non-overlapping RLL substrings of length q , and each needs q bits to be encoded. It sums up to

$$q \cdot \frac{1}{10} \left\lceil \frac{B}{4A} \right\rceil \sum_{a=1}^A c_a \geq \frac{B}{40A} \sum_{a=1}^A (L_a - q)$$

bits. Therefore,

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})|\mathcal{E}_l] &\geq \left(\frac{1}{9} - \frac{(2BL)^2}{2^{q'}} \right) \frac{B}{40A} \sum_{a=1}^A (L_a - q) \\ &\geq \frac{BL}{360} (1 + o(1)). \end{aligned}$$

The desired result thus follows from (2). \blacksquare

The above three theorems are summarized in Corollary 6 in the Results section to show that poorly choosing M prevents efficient compression by VLD.

In the next theorem, we give our upper bound on $\mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})]$. We consider the case when 2^M is of order

$\Theta(\log B)$ and show that variable-length deduplication achieves high compression ratios.

Theorem 18. Consider the source model with edit probability $\delta < \frac{1}{2}$. For any $\gamma \in (\delta, 1/2)$, the performance of variable-length deduplication with marker length M such that $2^M = \Theta(\log(B/A))$ satisfies

$$\begin{aligned} \mathbb{E}[\mathcal{L}_{VL}(\mathbf{s})] &\leq \\ &\left(12e^{-c_M}(c_M + 1) + 4H(\gamma, \delta)\frac{(1+k_1)}{k_2}c_M\right)BL(1 + o(1)), \end{aligned} \quad (21)$$

as $B \rightarrow \infty$, where $c_M = \frac{\log(B/A)}{H(\gamma, \delta)2^{M+1}}$.

Proof: First, encoding the length $|\mathbf{s}|$ takes $2\log|\mathbf{s}| + 3 \leq 2\log(BL) + 5$ bits. We study next the encoding of chunks. We adopt the same strategy as [13]: dividing chunks into two categories, interior chunks and boundary chunks. Consider all chunks whose first symbols are in Y_b (see Figure 3). Some chunks depend on the values of the neighboring source blocks Y_{b-1} and Y_{b+1} , i.e., it is possible to alter the chunk by replacing Y_{b-1} or Y_{b+1} with other strings. We call these the ‘boundary’ chunks of Y_b . Other chunks are independent of the values of the neighboring source blocks. We call these the ‘interior’ chunks of Y_b . Denote the set of interior chunks in \mathbf{s} by $\mathcal{C}^\circ(\mathbf{s})$. Note that we consider the first chunk and the last chunk of the whole data stream as boundary chunks. It is pointed out in [13] that the number of boundary chunks is upper bounded by $3(B+1)$ and the expected total length of boundary chunks is upper bounded by $B2^{M+2}$ ⁴. Therefore, encoding unique boundary chunks takes at most $3(B+1) + B2^{M+2}$ bits.

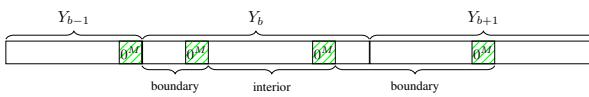


Figure 3. Occurrences of boundary chunks and interior chunks of Y_b in variable-length chunking.

We consider next encoding unique interior chunks. Clearly, every interior chunk follows a 0^M , i.e., the ending marker of the previous chunk. Moreover, this 0^M must also fully lie in the same source block as the chunk since otherwise this chunk is not an interior chunk. Therefore, the probability of occurrence of an interior chunk $u10^M$ is at most the probability of the occurrence of $0^M u10^M$ as a source block

substring. It follows that

$$\begin{aligned} &\mathbb{E}\left[\sum_{\mathbf{w} \in \mathcal{C}^\circ(\mathbf{s})} (|\mathbf{w}| + 1) | \mathcal{E}_u\right] \\ &\leq (M+1) + \sum_{\mathbf{u} \in R_M} \Pr(0^M \mathbf{u} 10^M \in \mathcal{C}^\circ(\mathbf{s}) | \mathcal{E}_u) (|\mathbf{u}| + M + 2) \\ &\leq (M+1) + \sum_{\mathbf{u} \in R_M} \Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) (|\mathbf{u}| + M + 2), \end{aligned} \quad (22)$$

where the term $M+1$ accounts for the chunk 0^M . We compute the summation in (22). Fix $\gamma \in (\delta, 1/2)$ and let $\ell_\gamma = \frac{\log(B/A)}{H(\gamma, \delta)}$.

- For all \mathbf{u} such that $|0^M \mathbf{u} 10^M| \leq \log B$, we trivially bound $\Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u)$ from above by 1. It follows that

$$\begin{aligned} &\sum_{\ell=0}^{\lfloor \log B \rfloor - 2M - 1} \sum_{\mathbf{u} \in R_M^\ell} \Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) (\ell + M + 2) \\ &\leq \sum_{\ell=0}^{\lfloor \log B \rfloor - 2M - 1} \sum_{\mathbf{u} \in R_M^\ell} (\ell + M + 2) \\ &\leq \sum_{\ell=0}^{\lfloor \log B \rfloor - 2M - 1} 2^\ell (\ell + M + 2) \\ &\leq (\lfloor \log B \rfloor - M + 1) 2^{\log B - 2M} \\ &\leq \frac{B \log B}{2^{2M}}. \end{aligned} \quad (23)$$

- For \mathbf{u} such that $|0^M \mathbf{u} 10^M| \geq \ell_\gamma$, we apply Lemma 11 and find

$$\begin{aligned} \Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) &\leq 2AL \frac{\mathcal{S}_\delta(|0^M \mathbf{u} 10^M|, \frac{3B}{2A})}{2^{|0^M \mathbf{u} 10^M|}} \\ &\leq \frac{3BL}{2^{|0^M \mathbf{u} 10^M|}}. \end{aligned}$$

⁴Although in [13], the upper bounds are derived for source strings produced by an edit-free source, the same upper bounds hold when edits exist since every source block is still a Bernoulli(1/2) process by itself.

It follows that

$$\begin{aligned}
& \sum_{\ell=\lceil \ell_\gamma \rceil - 2M-1}^{2L} (\ell + M + 2) \sum_{\mathbf{u} \in R_M^\ell} \Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) \\
& \leq \sum_{\ell=\lceil \ell_\gamma \rceil - 2M-1}^{2L} \sum_{\mathbf{u} \in R_M^\ell} \frac{3BL}{2^{\ell+2M+1}} (\ell + M + 2) \\
& \leq \sum_{\ell=\lceil \ell_\gamma \rceil - 2M-1}^{2L} 2 \left(2 - \frac{1}{2^M}\right)^\ell \frac{3BL}{2^{\ell+2M+1}} (\ell + M + 2) \\
& = \frac{3BL}{2^{2M}} \left(\sum_{\ell=\lceil \ell_\gamma \rceil - 2M-1}^{2L} \left(1 - \frac{1}{2^{M+1}}\right)^\ell (M+2) \right. \\
& \quad \left. + \sum_{\ell=\lceil \ell_\gamma \rceil - 2M-1}^{2L} \left(1 - \frac{1}{2^{M+1}}\right)^\ell \ell \right) \\
& = (1 + o(1)) \frac{3BL}{2^{2M}} \left(2^{M+1} \cdot e^{-\frac{\lceil \ell_\gamma \rceil - 2M-1}{2^{M+1}}} \right. \\
& \quad \left. + 2^{2(M+1)} \cdot e^{-\frac{\lceil \ell_\gamma \rceil - 2M-1}{2^{M+1}}} \left(\frac{\lceil \ell_\gamma \rceil - 2M - 1}{2^{M+1}} + 1 \right) \right) \\
& = 12BL \cdot e^{-\frac{\ell_\gamma}{2^{M+1}}} \left(\frac{\ell_\gamma}{2^{M+1}} + 1 \right) (1 + o(1)), \tag{24}
\end{aligned}$$

Thus,

$$\begin{aligned}
& \sum_{\ell=\lceil \log B \rceil - 2M-1}^{\lfloor \ell_\gamma \rfloor - 2M-1} \sum_{\mathbf{u} \in R_M^\ell} \Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) \\
& \quad \cdot (\ell + M + 2) \\
& \leq \sum_{\ell=\lceil \log B \rceil - 2M-1}^{\lfloor \ell_\gamma \rfloor - 2M-1} \sum_{\mathbf{u} \in R_M^\ell} 5BL 2^{-(\ell+2M+1)(1+D(\gamma||\delta))} \\
& \quad \cdot (\ell + M + 2) \\
& \leq \frac{5BL}{2^{2M}} \sum_{\ell=\lceil \log B \rceil - 2M-1}^{\lfloor \ell_\gamma \rfloor - 2M-1} \left(1 - \frac{1}{2^{M+1}}\right)^\ell 2^{-(\ell+2M+1)D(\gamma||\delta)} \\
& \quad \cdot (\ell + M + 2) \\
& \leq \frac{5BL\ell_\gamma^2}{2^{2M}} \left(1 - \frac{1}{2^{M+1}}\right)^{\log B - 2M - 1} 2^{-D(\gamma||\delta)\log B} \\
& = \Theta(B^{1-D(\gamma||\delta)} L) \\
& = o(BL), \tag{25}
\end{aligned}$$

where the first equality follows from the fact that $\frac{\ell_\gamma^2}{2^{2M}}$ and $\left(1 - \frac{1}{2^{M+1}}\right)^{\log B - 2M - 1}$ are both $\Theta(1)$ since 2^M and ℓ_γ are $\Theta(\log(B/A))$.

Plugging (23), (24) and (25) in (22), we find that as $B \rightarrow \infty$ (also $A, L \rightarrow \infty$),

$$\mathbb{E} \left[\sum_{\mathbf{w} \in \mathcal{C}^\circ(\mathbf{s})} (|\mathbf{w}| + 1) | \mathcal{E}_u \right] \leq 12e^{-c_M} (c_M + 1) BL + o(BL),$$

where $c_M = \frac{\ell_\gamma}{2^{M+1}}$.

If the complement of \mathcal{E}_u holds, then the number of bits needed for encoding interior chunks at their first occurrences is at most $4BL$, since the total length of interior chunks is at most $2BL$ and the total number of chunks is at most $2BL$. By noting that $\Pr(\bar{\mathcal{E}}_u) \leq Ae^{-\frac{B}{10A}}$,

$$\begin{aligned}
\Pr(0^M \mathbf{u} 10^M \in Y_1^B | \mathcal{E}_u) & \leq 2AL \frac{\mathcal{S}_\delta(|0^M \mathbf{u} 10^M|, \frac{3B}{2A})}{2^{|0^M \mathbf{u} 10^M|}} \\
& \leq 5BL 2^{-|0^M \mathbf{u} 10^M|(1+D(\gamma||\delta))},
\end{aligned}$$

where the second inequality follows from (8) and the fact that $2^{nH(\gamma)} \leq \frac{B}{A} 2^{-nD(\gamma||\delta)}$ if $n \leq \frac{\log(B/A)}{H(\gamma, \delta)}$.

$$\begin{aligned}
& \mathbb{E} \left[\sum_{\mathbf{w} \in \mathcal{C}^\circ(\mathbf{s})} (|\mathbf{w}| + 1) \right] \\
& \leq 12e^{-c_M} (c_M + 1) BL + o(BL) + 4BL Ae^{-\frac{B}{10A}} \\
& = 12e^{-c_M} (c_M + 1) BL (1 + o(1)). \tag{26}
\end{aligned}$$

The number of bits needed for encoding pointers of repeated chunks can be bounded from above in a trivial way. Note that there are at most $\frac{|\mathbf{s}|}{M} + 1$ strings in the dictionary T . So a pointer takes at most $\log\left(\frac{|\mathbf{s}|}{M} + 1\right) + 1 \leq \log|\mathbf{s}|$ bits. Moreover, the total number of chunks is less than the number of occurrences of 0^M plus 1 since every chunk except possibly the last one ends with 0^M . On average, the number of occurrences of 0^M in Y_b is at most $\frac{|Y_b|}{2^M}$. So given $|\mathbf{s}|$, the expected number of

chunks in s is at most $\frac{|s|}{2^M} + B + 1$. Therefore the expected number of bits used by pointers is at most

$$\begin{aligned} & \mathbb{E} \left[(\log |s| + 1) \cdot \left(\frac{|s|}{2^M} + B + 1 \right) \right] \\ & \leq \log(2BL + 1) \left(\frac{2BL}{2^M} + B + 1 \right) \\ & \leq 2BL \frac{\log(BL)}{2^M} (1 + o(1)) \\ & \leq 4H(\gamma, \delta) \frac{(1 + k_1)}{k_2} c_M \cdot BL(1 + o(1)), \end{aligned} \quad (27)$$

where the last inequality follows from $\frac{\log(BL)}{\log(B/A)} \leq \frac{1+k_1}{k_2}(1+o(1))$.

The desired result follows from summing (26) and (27) and noting that the number of bits used for encoding the length of s and the unique boundary chunks are $o(BL)$. \blacksquare

A detailed analysis in the Results section shows that as δ approaches 0, by appropriately choosing M , the compression ratio $\frac{\mathbb{E}[|s|]}{\mathbb{E}[\mathcal{L}_{VL}(s)]}$ can get arbitrarily large.

VIII. CONCLUSION

In this paper, we studied the performance of deduplication algorithms on data streams with approximate repeats, a situation that is common in practice. For simplicity, we modeled the process producing approximate repeats as independent bit-wise Bernoulli substitutions. We showed, in particular, that correctly choosing the chunk lengths is critical to the success of deduplication. With optimally chosen chunk lengths, deduplication in the fixed-length scheme is shown to achieve performance within a constant factor of optimal for a specific family of source models and with the knowledge of source parameters. Additionally, appropriately choosing the length of the marker leads to suitable chunk lengths for variable-length deduplication, resulting in arbitrarily large compression ratios as source entropy gets smaller.

While this work sheds light on certain important aspects of the problem, the information-theoretic analysis of data deduplication provides a wealth of open problems. For example, while VLD was shown to achieve high compression ratios, it is not known whether it is order optimal. Moreover, the source model proposed in this paper only included independent substitution edits. However, in practice, insertions, deletions and substitutions of single symbols, as well as longer strings, occur frequently. The probabilistic description of the source models can also be further refined based on experiments. Therefore, to gain a fuller understanding, it is important to study deduplication algorithms under more general source models and edit processes.

APPENDIX A PROOF OF LEMMA 1

Lemma 1. Let k be a positive integer. The number of binary k -RLL strings of length n , $|R_k^n|$, satisfies

$$(2 - \frac{1}{2^{k-2}})^n \leq |R_k^n| \leq 2(2 - \frac{1}{2^k})^n.$$

Proof: Clearly, if $0 \leq n \leq k-1$, then any string of length n is a k -RLL string (we consider the empty string as the only string of length 0). Therefore, for all $0 \leq n \leq k-1$,

$$|R_k^n| = 2^n \geq (2 - \frac{1}{2^{k-2}})^n,$$

and

$$\begin{aligned} |R_k^n| &= 2^n = 2^{n+1}2^{-1} \leq 2^{n+1}(1 - \frac{1}{2^{k+1}})^{k-1} \\ &\leq 2^{n+1}(1 - \frac{1}{2^{k+1}})^n = 2(2 - \frac{1}{2^k})^n. \end{aligned}$$

For $n \geq k$, we prove the lemma by induction on n . Suppose the desired results hold for all $n' < n$. It is shown in [15, Chapter 8] that $|R_k^N| = \sum_{i=1}^k |R_k^{N-i}|$ for all $N \geq k$. Therefore,

$$\begin{aligned} |R_k^n| &= \sum_{i=1}^k |R_k^{n-i}| \geq \sum_{i=1}^k (2 - \frac{1}{2^{k-2}})^{n-i} \\ &= \frac{(2 - \frac{1}{2^{k-2}})^n - (2 - \frac{1}{2^{k-2}})^{n-k}}{1 - \frac{1}{2^{k-2}}} \\ &= \left(2 - \frac{1}{2^{k-2}}\right)^n \\ &\quad + \frac{(2 - \frac{1}{2^{k-2}})^{n-k}2^k}{2^{k-2} - 1} \left((1 - \frac{1}{2^{k-1}})^k - \frac{1}{4} \right) \\ &\geq (2 - \frac{1}{2^{k-2}})^n, \end{aligned}$$

and

$$\begin{aligned} |R_k^n| &= \sum_{i=1}^k |R_k^{n-i}| \leq \sum_{i=1}^k 2(2 - \frac{1}{2^k})^n \\ &= 2 \frac{(2 - \frac{1}{2^k})^n - (2 - \frac{1}{2^k})^{n-k}}{1 - \frac{1}{2^k}} \\ &= 2(2 - \frac{1}{2^k})^n + \frac{2(2 - \frac{1}{2^k})^{n-k}}{1 - \frac{1}{2^k}} \left(\left(\frac{2 - \frac{1}{2^k}}{2} \right)^k - 1 \right) \\ &\leq 2(2 - \frac{1}{2^k})^n. \end{aligned}$$

APPENDIX B PROOF OF LEMMA 3

Lemma 3. Let r be a string drawn uniformly at random from Σ^ℓ . Let r_1, r_2, \dots, r_m be m iid descendants of r by δ -edit and let $r_{[m]} = \{r_1, r_2, \dots, r_m\}$. For any $w \in \Sigma^\ell$, let $w \in r_{[m]}$ denote the event that $w = r_i$ for some i . Then

$$\frac{1}{2} \frac{\mathcal{S}_\delta(\ell, m)}{2^\ell} \leq \Pr(w \in r_{[m]}) \leq \frac{\mathcal{S}_\delta(\ell, m)}{2^\ell}, \quad (4)$$

and thus the expected number of unique strings in $r_{[m]}$ is bounded between $\frac{1}{2}\mathcal{S}_\delta(\ell, m)$ and $\mathcal{S}_\delta(\ell, m)$.

Furthermore, $\mathcal{S}_\delta(\ell, m)$ takes the following values for different values of ℓ and m :

- If $\ell \geq \frac{\log m}{H(\delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \geq \frac{1}{4}m. \quad (5)$$

In particular if $\ell \geq \frac{\log m}{\log(\frac{1}{1-\delta})}$, then

$$\mathcal{S}_\delta(\ell, m) = m. \quad (6)$$

- If $\ell \leq \frac{\log m}{H(\frac{1}{2}, \delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \geq 2^{\ell-1}. \quad (7)$$

In particular if $\ell \leq \frac{\log m}{\log(\frac{1}{\delta})}$, then

$$\mathcal{S}_\delta(\ell, m) = 2^\ell.$$

- For any $\delta < \delta' < \frac{1}{2}$,

$$\mathcal{S}_\delta(\ell, m) \leq 2^{\ell H(\delta')} + m2^{-\ell D(\delta' || \delta)}. \quad (8)$$

In particular if $\ell = \frac{\log m}{H(\delta', \delta)}$, then

$$\mathcal{S}_\delta(\ell, m) \leq 2^{\ell H(\delta')} + m2^{-\ell D(\delta' || \delta)} = 2^{\ell H(\delta')+1}. \quad (9)$$

- For any values of ℓ and m ,

$$\mathcal{S}_\delta(\ell, m) \leq \min(2^\ell, m). \quad (10)$$

Proof: We first prove inequality (4). Given \mathbf{r} , the probability of a δ -edit descendant being equal to \mathbf{w} is $\delta^{d_{\mathbf{w}, \mathbf{r}}}(1-\delta)^{\ell-d_{\mathbf{w}, \mathbf{r}}}$, where $d_{\mathbf{w}, \mathbf{r}}$ denotes the Hamming distance between \mathbf{w} and \mathbf{r} . Therefore,

$$\begin{aligned} \Pr(\mathbf{w} \in \mathbf{r}_{[m]}) &= 1 - \Pr(\mathbf{w} \notin \mathbf{r}_{[m]}) \\ &= 1 - \sum_{\mathbf{r} \in \Sigma^\ell} \Pr(\mathbf{r}) \Pr(\mathbf{w} \neq \mathbf{r}_1 | \mathbf{r})^m \\ &= 1 - \sum_{\mathbf{r} \in \Sigma^\ell} \Pr(\mathbf{r}) (1 - \delta^{d_{\mathbf{w}, \mathbf{r}}}(1-\delta)^{\ell-d_{\mathbf{w}, \mathbf{r}}})^m \\ &= 1 - \sum_{t=0}^{\ell} \left(\frac{\binom{\ell}{t}}{2^\ell} (1 - \delta^t(1-\delta)^{\ell-t})^m \right), \end{aligned}$$

where the second equality follows from the fact that $\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_m$ are iid given \mathbf{r} and the last equality follows from the fact that there are $\binom{\ell}{t}$ strings of length ℓ that are at Hamming distance t from \mathbf{w} . The desired inequalities then follow directly from applying inequalities (1) on $1 - \left(1 - \delta^t(1-\delta)^{\ell-t}\right)^m$.

The expected number of unique strings in $\mathbf{r}_{[m]}$ equals

$$\mathbb{E} \left[\sum_{\mathbf{w} \in \Sigma^\ell} I_{\mathbf{w} \in \mathbf{r}_{[m]}} \right] = \sum_{\mathbf{w} \in \Sigma^\ell} \Pr(\mathbf{w} \in \mathbf{r}_{[m]}).$$

So the upper bound $\mathcal{S}_\delta(\ell, m)$ and the lower bound $\frac{1}{2}\mathcal{S}_\delta(\ell, m)$ follow from replacing $\Pr(\mathbf{w} \in \mathbf{r}_{[m]})$ with its upper and lower bounds, respectively.

We show that $\mathcal{S}_\delta(\ell, m)$ takes the given values for different m and ℓ :

- When $\ell \geq \frac{\log m}{H(\delta)}$, $m\delta^{\ell\delta}(1-\delta)^{(1-\delta)\ell} \leq 1$. It follows that

$$\begin{aligned} \mathcal{S}_\delta(\ell, m) &\geq \sum_{t=\lceil \delta\ell \rceil}^{\ell} \binom{\ell}{t} \min(1, m\delta^t(1-\delta)^{\ell-t}) \\ &= \sum_{t=\lceil \delta\ell \rceil}^{\ell} \binom{\ell}{t} m\delta^t(1-\delta)^{\ell-t} \geq \frac{1}{4}m, \end{aligned}$$

where the equality follows from the fact that $m\delta^t(1-\delta)^{\ell-t}$ is decreasing in t so $m\delta^t(1-\delta)^{\ell-t} \leq 1$ for all $t \geq \delta\ell$ and the second inequality follows from the result shown in [4] that for a binomial random variable X with parameters n and p , $\Pr(X \geq np) > \frac{1}{4}$ if $p \geq 1/n$.

Moreover, when $\ell \geq \frac{\log m}{\log(\frac{1}{1-\delta})}$, $m\delta^t(1-\delta)^{\ell-t} \leq 1$ for all t . Hence,

$$\mathcal{S}_\delta(\ell, m) = \sum_{t=0}^{\ell} \binom{\ell}{t} m\delta^t(1-\delta)^{\ell-t} = m.$$

- When $\ell \leq \frac{\log m}{H(\frac{1}{2}, \delta)}$, $m\delta^{\frac{\ell}{2}}(1-\delta)^{\frac{\ell}{2}} \geq 1$. It follows that

$$\begin{aligned} \mathcal{S}_\delta(\ell, m) &\geq \sum_{t=0}^{\lfloor \frac{\ell}{2} \rfloor} \binom{\ell}{t} \min(1, m\delta^t(1-\delta)^{\ell-t}) = \sum_{t=0}^{\lfloor \frac{\ell}{2} \rfloor} \binom{\ell}{t} \\ &\geq 2^{\ell-1}, \end{aligned}$$

where the first inequality follows from the fact that $m\delta^t(1-\delta)^{\ell-t} \geq 1$ for all $t \leq \frac{\ell}{2}$.

Moreover, when $\ell \leq \frac{\log m}{\log(\frac{1}{\delta})}$, $m\delta^t \geq 1$ for all t . Hence,

$$\mathcal{S}_\delta(\ell, m) = \sum_{t=0}^{\ell} \binom{\ell}{t} \cdot 1 = 2^\ell.$$

- For any $\delta < \delta' < 1/2$,

$$\begin{aligned} \mathcal{S}_\delta(\ell, m) &\leq \sum_{t=0}^{\lfloor \delta' \ell \rfloor} \binom{\ell}{t} + \sum_{t=\lceil \delta' \ell \rceil}^{\ell} \binom{\ell}{t} m\delta^t(1-\delta)^{\ell-t} \\ &\leq 2^{\ell H(\delta')} + m2^{-\ell D(\delta' || \delta)}, \end{aligned}$$

where the second inequality follows from applying the Chernoff bound on a binomial distribution with parameters ℓ and δ .

When $\ell = \frac{\log m}{H(\delta', \delta)}$, $2^{\ell H(\delta')} = m2^{-\ell D(\delta' || \delta)}$. So $2^{\ell H(\delta')} + m2^{-\ell D(\delta' || \delta)} = 2^{\ell H(\delta')+1}$ and

$$\mathcal{S}_\delta(\ell, m) \leq 2^{\ell H(\delta')+1}.$$

- The upper bounds 2^ℓ and m follow from:

$$\mathcal{S}_\delta(\ell, m) \leq \sum_{t=0}^{\ell} \frac{\binom{\ell}{t}}{2^\ell} = 1,$$

$$\mathcal{S}_\delta(\ell, m) \leq \sum_{t=0}^{\ell} \frac{\binom{\ell}{t}}{2^\ell} m\delta^t(1-\delta)^{\ell-t} = \frac{m}{2^\ell}.$$

APPENDIX C PROOFS OF LEMMA 4 AND LEMMA 5

Lemma 4. Suppose K strings of length n are chosen independently and uniformly from Σ^n . Assume each string produces at least m_1 and at most m_2 descendants by δ -edits. For any

string \mathbf{w} with $|\mathbf{w}| = n$, let $G_{\mathbf{w}}$ denote the event that \mathbf{w} equals one or more descendants. Then

$$\begin{aligned} \frac{1}{2} \min\left(1, \frac{1}{2} K \frac{\mathcal{S}_{\delta}(n, m_1)}{2^n}\right) \\ \leq \Pr(G_{\mathbf{w}}) \leq \\ \min\left(1, K \frac{\mathcal{S}_{\delta}(n, m_2)}{2^n}\right). \end{aligned}$$

Proof: Let the K strings be denoted $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_K$. Let $G_{\mathbf{w}}(i)$ denote the event that \mathbf{w} equals one of the descendants of \mathbf{y}_i . Clearly, $G_{\mathbf{w}}(1), G_{\mathbf{w}}(2), \dots, G_{\mathbf{w}}(K)$ are independent and

$$G_{\mathbf{w}} = \bigcup_{i=1}^K G_{\mathbf{w}}(i). \quad (28)$$

Note that by Lemma 3 and the fact that $\mathcal{S}_{\delta}(n, m)$ is non-decreasing in m ,

$$\frac{1}{2} \frac{\mathcal{S}_{\delta}(n, m_1)}{2^n} \leq \Pr(G_{\mathbf{w}}(i)) \leq \frac{\mathcal{S}_{\delta}(n, m_2)}{2^n}.$$

Applying the union bound on (28) gives

$$\Pr(G_{\mathbf{w}}) \leq \sum_{i=1}^K \Pr(G_{\mathbf{w}}(i)) \leq K \frac{\mathcal{S}_{\delta}(n, m_2)}{2^n}.$$

The desired upper bound follows by noting that 1 is a trivial upper bound.

We then prove the lower bound. By independence,

$$\begin{aligned} \Pr(G_{\mathbf{w}}) &= \Pr\left(\bigcup_{i=1}^K G_{\mathbf{w}}(i)\right) \\ &= 1 - \prod_{i=1}^K (1 - \Pr(G_{\mathbf{w}}(i))) \\ &\geq 1 - \left(1 - \frac{1}{2} \frac{\mathcal{S}_{\delta}(n, m_1)}{2^n}\right)^K \\ &\geq \frac{1}{2} \min\left(1, \frac{1}{2} K \frac{\mathcal{S}_{\delta}(n, m_1)}{2^n}\right), \end{aligned}$$

where the last inequality follows from inequality (1) that $1 - (1 - x)^n \geq \frac{1}{2} \min(1, nx)$ for $x \in (0, 1)$ and integer n . ■

Lemma 5. Consider the two-stage fixed-length chunking process with first-stage parsing length $D = L$ and chunk length ℓ . The dictionary sizes $T_F^1(\mathbf{s})$ and $T_F^{1/2}(\mathbf{s})$ satisfy

$$\begin{aligned} \mathbb{E}[|T_F^1(\mathbf{s})| | \mathcal{E}_u] &\leq \min\left(2^\ell, ACS_{\delta}\left(\ell, \frac{3B}{2A}\right)\right) + B, \\ \mathbb{E}[|T_F^{1/2}(\mathbf{s})| | \mathcal{E}_l] &\geq \frac{1}{2} \min\left(2^\ell, \frac{1}{2} ACS_{\delta}\left(\ell, \frac{B}{4A}\right)\right). \quad (12) \end{aligned}$$

Proof: The size of $T_F^1(\mathbf{s})$ equals the number of distinct strings among chunks $Z_c^b, 1 \leq c \leq C+1, 1 \leq b \leq B$. Clearly, chunks of length ℓ are δ -edit descendants of the AC source symbol substrings $U_c^a, 1 \leq c \leq C, 1 \leq a \leq A$, which are independent and uniformly distributed in Σ^ℓ . Given \mathcal{E}_u , each U_c^a has at most $\frac{3B}{2A}$ descendants. Moreover, since we assume that the source symbols X_1, \dots, X_A are chosen uniformly and independently, it follows directly from Lemma 4 that for any ℓ -string \mathbf{w} ,

$$\Pr(\mathbf{w} \in T_F^1(\mathbf{s}) | \mathcal{E}_u) \leq \min\left(1, AC \frac{\mathcal{S}_{\delta}(\ell, \frac{3B}{2A})}{2^\ell}\right).$$

Hence

$$\begin{aligned} \mathbb{E}[|T_F^1(\mathbf{s})| | \mathcal{E}_u] &\leq \sum_{\mathbf{w} \in \Sigma^\ell} \Pr(\mathbf{w} \in T_F^1(\mathbf{s}) | \mathcal{E}_u) + B \\ &\leq \min\left(2^\ell, ACS_{\delta}\left(\ell, \frac{3B}{2A}\right)\right) + B, \end{aligned}$$

where the addend B accounts for the chunks of lengths less than ℓ at the end of each source block, if any.

The lower bound on $|T_F^{1/2}(\mathbf{s})|$ given \mathcal{E}_l follows similarly from Lemma 4. ■

APPENDIX D PROOF OF LEMMA 10 AND LEMMA 11

Lemma 10. Suppose K strings of length n are chosen independently and uniformly from Σ^n . Assume each string produces at least m_1 and at most m_2 descendants by δ -edits. For any string \mathbf{w} with $|\mathbf{w}| \leq n$, let $H_{\mathbf{w}}$ denote the event that \mathbf{w} appears as a substring of one or more descendants. Then,

$$\begin{aligned} \frac{1}{2} \min\left(1, \frac{1}{2} \left\lfloor \frac{n}{|\mathbf{w}|} \right\rfloor K \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, m_1)}{2^{|\mathbf{w}|}}\right) \\ \leq \Pr(H_{\mathbf{w}}) \leq \\ \min\left(1, (n - |\mathbf{w}| + 1) K \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, m_2)}{2^{|\mathbf{w}|}}\right). \end{aligned}$$

Proof: Let the K strings be denoted $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_K$. We use \mathcal{D}_i to denote the set of δ -edit descendants of \mathbf{y}_i . Let $H_{\mathbf{w}}(i, j)$ denote the event that $\mathbf{w} = \mathbf{x}_{j, |\mathbf{w}|}$ for some $\mathbf{x} \in \mathcal{D}_i$. Clearly,

$$H_{\mathbf{w}} = \bigcup_{i=1}^K \bigcup_{j=1}^{n-|\mathbf{w}|+1} H_{\mathbf{w}}(i, j). \quad (29)$$

Note that the strings $\{\mathbf{x}_{j, |\mathbf{w}|}\}_{\mathbf{x} \in \mathcal{D}_i}$ are iid δ -edit descendants of $(\mathbf{y}_i)_{j, |\mathbf{w}|}$. Hence by Lemma 3

$$\begin{aligned} \frac{1}{2} \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, m_1)}{2^{|\mathbf{w}|}} &\leq \frac{1}{2} \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, |\mathcal{D}_i|)}{2^{|\mathbf{w}|}} \\ &\leq \Pr(H_{\mathbf{w}}(i, j)) \leq \\ &\frac{\mathcal{S}_{\delta}(|\mathbf{w}|, |\mathcal{D}_i|)}{2^{|\mathbf{w}|}} \leq \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, m_2)}{2^{|\mathbf{w}|}}. \end{aligned}$$

where the first and the last inequalities follow from $m_1 \leq |\mathcal{D}_i| \leq m_2$.

Applying the union bound on (29) gives

$$\begin{aligned} \Pr(H_{\mathbf{w}}) &\leq \bigcup_{i=1}^K \bigcup_{j=1}^{n-|\mathbf{w}|+1} \Pr(H_{\mathbf{w}}(i, j)) \\ &\leq (n - |\mathbf{w}| + 1) K \frac{\mathcal{S}_{\delta}(|\mathbf{w}|, m_2)}{2^{|\mathbf{w}|}}. \end{aligned}$$

The desired upper bound follows by noting that 1 is a trivial upper bound.

We next prove the lower bound. For each i , non-overlapping substrings of \mathbf{r}_i are independent and so are their descendants.

Hence, events $H_{\mathbf{w}}(i, j)$, $j = 1, 1 + |\mathbf{w}|, \dots, 1 + (p - 1)|\mathbf{w}|$, where $p = \left\lfloor \frac{n}{|\mathbf{w}|} \right\rfloor$, are mutually independent. It follows that

$$\begin{aligned} & \Pr\left(\bigcup_{i=1}^K \bigcup_{a=1}^p H_{\mathbf{w}}(i, 1 + (a - 1)|\mathbf{w}|)\right) \\ &= 1 - \prod_{i=1}^K \prod_{a=1}^p (1 - \Pr(H_{\mathbf{w}}(i, 1 + (a - 1)|\mathbf{w}|))) \\ &\geq 1 - \left(1 - \frac{1}{2} \frac{\mathcal{S}_\delta(|\mathbf{w}|, m_1)}{2^{|\mathbf{w}|}}\right)^{Kp} \\ &\geq \frac{1}{2} \min\left(1, \frac{1}{2} Kp \frac{\mathcal{S}_\delta(|\mathbf{w}|, m_1)}{2^{|\mathbf{w}|}}\right), \end{aligned}$$

where the last inequality follows from inequality (1) that $1 - (1 - x)^n \geq \frac{1}{2} \min(1, nx)$ for $x \in (0, 1)$ and integer n . The desired lower bound thus follows by noting that

$$\bigcup_{i=1}^K \bigcup_{a=1}^p H_{\mathbf{w}}(i, 1 + (a - 1)|\mathbf{w}|) \subseteq H_{\mathbf{w}}.$$

■

Lemma 11. Consider the source model with edit probability δ . For any string $\mathbf{w} \in \Sigma^*$ with $|\mathbf{w}| \leq 2L$,

$$\Pr(\mathbf{w} \in Y_1^B | \mathcal{E}_u) \leq \min\left(1, 2AL \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{3B}{2A})}{2^{|\mathbf{w}|}}\right).$$

For any string $\mathbf{w} \in \Sigma^*$ with $|\mathbf{w}| \leq \lceil \frac{1}{2}L \rceil$,

$$\Pr(\mathbf{w} \in Y_1^{B/2} | \mathcal{E}_l) \geq \frac{1}{2} \min\left(1, \frac{AL}{8|\mathbf{w}|} \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}}\right). \quad (17)$$

Proof: Recall that we assume every source symbol (and thus every source block) is of length at least $\frac{1}{2}L$ and at most $2L$. So we can get a lower bound on $\Pr(\mathbf{w} \in Y_1^{B/2} | \mathcal{E}_l)$ by assuming every source block is of length $\frac{L}{2}$. Similarly, we get an upper bound on $\Pr(\mathbf{w} \in Y_1^B | \mathcal{E}_u)$ by assuming every source block is of length $2L$.

Now that the B source blocks are independent and each is a δ -edit descendant of one of the A source symbols. Moreover, each random string (source symbol) has at most $\frac{3B}{2A}$ descendants given \mathcal{E}_u . Therefore, by directly applying Lemma 10,

$$\begin{aligned} \Pr(\mathbf{w} \in Y_1^B | \mathcal{E}_u) &\leq \min\left(1, (2L - |\mathbf{w}| + 1)A \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{3B}{2A})}{2^{|\mathbf{w}|}}\right) \\ &\leq \min\left(1, 2LA \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{3B}{2A})}{2^{|\mathbf{w}|}}\right). \end{aligned}$$

The lower bound can be obtained similarly:

$$\begin{aligned} \Pr(\mathbf{w} \in Y_1^{B/2} | \mathcal{E}_l) &\geq \frac{1}{2} \min\left(1, \frac{1}{2} \left\lfloor \frac{L/2}{|\mathbf{w}|} \right\rfloor A \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}}\right) \\ &\geq \frac{1}{2} \min\left(1, \frac{1}{8} \frac{L}{|\mathbf{w}|} A \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}}\right). \end{aligned}$$

■

APPENDIX E PROOF OF LEMMA 12

Lemma 12. Consider the source model with edit probability $\delta < \frac{1}{2}$. For any n such that $\frac{\log(B/A)-2}{H(\delta)} \leq n + 2M + 2 \leq \frac{L}{4}$,

$$\begin{aligned} \sum_{\mathbf{u} \in R_M^n} \Pr\left(10^M \mathbf{u} 10^M \in Y_1^{B/2} | \mathcal{E}_l\right) &\geq \\ \frac{BL}{2^7 \cdot 2^{2M+2}} \cdot \left(1 - \frac{1}{2^{M-1}}\right)^n &- \frac{3B^2 L^2}{2^{n+2M+2}}. \end{aligned}$$

Proof: Let $\mathbf{w} = 10^M \mathbf{u} 10^M$. By assumption, $|\mathbf{w}| = |\mathbf{u}| + 2M + 2 \geq \frac{\log(B/A)-2}{H(\delta)}$.

For definiteness, we assume $|Y_{1/2}(a)| = \frac{B}{4A}$ for all a and all source symbols are of length $\frac{L}{2}$. With these assumptions, we have a similar setting to that in Lemma 10. So we adopt the same notation. Let $H_{\mathbf{w}}$ denote $\mathbf{w} \in Y_1^{B/2}$ and $H_{\mathbf{w}}(a, j)$ denote the event that $\mathbf{w} = \mathbf{x}_{j, |\mathbf{w}|}$ for some $\mathbf{x} \in Y_{1/2}(a)$. Similar to (29):

$$H_{\mathbf{w}} = \bigcup_{a=1}^A \bigcup_{j=1}^{\lceil L/2 \rceil - |\mathbf{w}| + 1} H_{\mathbf{w}}(a, j). \quad (30)$$

Moreover,

$$\frac{1}{2} \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}} \leq \Pr(H_{\mathbf{w}}(a, j)) \leq \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}}.$$

In Lemma 11, an upper bound on $\Pr(\mathbf{w} \in Y_1^{B/2} | \mathcal{E}_l)$ is obtained by applying the union bound on (30). Here, we get a lower bound by the inclusion-exclusion principle:

$$\Pr(H_{\mathbf{w}}) \geq \sum_{a=1}^A \sum_{i=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, i)) \quad (31)$$

$$- \sum_{1 \leq a_1 \neq a_2 \leq A} \sum_{j=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \sum_{k=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a_1, j) \cap H_{\mathbf{w}}(a_2, k)) \quad (32)$$

$$- \sum_{a=1}^A \sum_{1 \leq j, k \leq \lceil L/2 \rceil - |\mathbf{w}|, j \neq k} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)). \quad (33)$$

We compute the three terms on the right-hand side of the inequality above as follows.

For the term in (31), since $|\mathbf{w}| \geq \frac{\log(B/A)-2}{H(\delta)}$,

$$\begin{aligned} \Pr(H_{\mathbf{w}}(a, i)) &\geq \frac{1}{2} \frac{\mathcal{S}_\delta(|\mathbf{w}|, \lceil \frac{B}{4A} \rceil)}{2^{|\mathbf{w}|}} \geq \frac{1}{2} \frac{\mathcal{S}_\delta(|\mathbf{w}|, \frac{B}{4A})}{2^{|\mathbf{w}|}} \\ &\geq \frac{B}{32A \cdot 2^{|\mathbf{w}|}}, \end{aligned}$$

where the last inequality follows from (5). It follows that

$$\begin{aligned} \sum_{a=1}^A \sum_{i=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, i)) &\geq A(\lceil L/2 \rceil - |\mathbf{w}|) \frac{B}{32A \cdot 2^{|\mathbf{w}|}} \\ &\geq \frac{BL}{2^7 \cdot 2^{|\mathbf{w}|}}. \end{aligned} \quad (34)$$

For the term in (32), since for all $a_1 \neq a_2$, \mathbf{y}_{a_1} and \mathbf{y}_{a_2} are independent and so are their descendants, we get

$$\begin{aligned}
& \sum_{1 \leq a_1 \neq a_2 \leq A} \sum_{j=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \sum_{k=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a_1, j) \cap H_{\mathbf{w}}(a_2, k)) \\
&= \sum_{1 \leq a_1 \neq a_2 \leq A} \sum_{j=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \sum_{k=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a_1, j)) \Pr(H_{\mathbf{w}}(a_2, k)) \\
&\leq \sum_{1 \leq a_1 \neq a_2 \leq A} \sum_{j=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \sum_{k=1}^{\lceil L/2 \rceil - |\mathbf{w}|} \left(\frac{\mathcal{S}_{\delta}(|\mathbf{w}|, \lceil \frac{B}{4A} \rceil)}{2^{|\mathbf{w}|}} \right)^2 \\
&\leq \sum_{1 \leq a_1 \neq a_1 \leq A} \frac{B^2 L^2}{A^2 2^{2|\mathbf{w}|}} \\
&\leq \frac{B^2 L^2}{2^{2|\mathbf{w}|}}, \tag{35}
\end{aligned}$$

where the second inequality follows from (10) that $\mathcal{S}_{\delta}(\ell, m) \leq m$ and the inequalities $\lceil L/2 \rceil - |\mathbf{w}| \leq L$, $\lceil \frac{B}{4A} \rceil \leq \frac{B}{A}$.

We then consider the term in (33), where the two occurrences of \mathbf{w} are among the descendants of a single source symbol, and thus might not be independent. Unlike the previous two terms, we consider lower bounding the sum of probabilities $\Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k))$ over all \mathbf{w} of the form $10^M \mathbf{u} 10^M$, $\mathbf{u} \in R_M^n$. For clarity of presentation, we first claim (to be proved later) that for any a ,

$$\begin{aligned}
& \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \sum_{\substack{1 \leq j, k \leq \lceil L/2 \rceil - |\mathbf{w}| \\ j \neq k}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
&\leq \frac{B^2 L^2}{A^2 2^{|\mathbf{w}|}} \left(1 + \frac{n+M+1}{L} \right). \tag{36}
\end{aligned}$$

It follows that

$$\begin{aligned}
& \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \sum_{a=1}^A \sum_{\substack{1 \leq j, k \leq \lceil L/2 \rceil - |\mathbf{w}| \\ j \neq k}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
&\leq \frac{B^2 L^2}{A^2 2^{|\mathbf{w}|}} \left(1 + \frac{n+M+1}{L} \right). \tag{37}
\end{aligned}$$

Thus, combining (33), (34), (35) and (37) gives

$$\begin{aligned}
& \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \Pr(H_{\mathbf{w}}) \\
&\geq \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \left(\frac{BL}{2^7 \cdot 2^{|\mathbf{w}|}} - \frac{B^2 L^2}{2^{2|\mathbf{w}|}} \right) \\
&\quad - \frac{B^2 L^2}{A^2 2^{|\mathbf{w}|}} \left(1 + \frac{n+M+1}{L} \right) \\
&\geq \frac{BL}{2^7 \cdot 2^{|\mathbf{w}|}} \cdot |R_M^n| \\
&\quad - \frac{B^2 L^2}{2^{2|\mathbf{w}|}} \cdot |R_M^n| - \frac{B^2 L^2}{A^2 2^{|\mathbf{w}|}} \left(1 + \frac{n+M+1}{L} \right) \\
&\geq \frac{BL}{2^7 \cdot 2^{|\mathbf{w}|}} \cdot |R_M^n| - \frac{3B^2 L^2}{2^{|\mathbf{w}|}},
\end{aligned}$$

where the last inequality follows from $|R_M^n| \leq 2^{|\mathbf{w}|}$, $\frac{n+M+1}{L} \leq 1$ and $A \geq 1$. The desired lower bound thus follows from bounding $|R_M^n|$ by Lemma 1.

Finally, we prove inequality (36). Fix a . That $H_{\mathbf{w}}(a, j)$ and $H_{\mathbf{w}}(a, k)$ both hold means there exist descendants $\mathbf{x}_1, \mathbf{x}_2$ (possibly the same one) of \mathbf{y}_a such that $(\mathbf{x}_1)_{j,|\mathbf{w}|} = (\mathbf{x}_2)_{k,|\mathbf{w}|} = \mathbf{w}$. Assume $j < k$ without loss of generality. We compute $\Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k))$ for different values of (j, k) :

- $|j - k| \geq |\mathbf{w}|$. The two occurrences of \mathbf{w} in \mathbf{x}_1 and \mathbf{x}_2 are plotted in Figure 4. In this case, they are produced by two non-overlapping substrings of \mathbf{y}_a and thus are independent. It follows that

$$\begin{aligned}
& \sum_{|j-k| \geq |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
&= \sum_{|j-k| \geq |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, j)) \Pr(H_{\mathbf{w}}(a, k)) \\
&\leq L^2 \left(\frac{\mathcal{S}_{\delta}(|\mathbf{w}|, \lceil \frac{B}{4A} \rceil)}{2^{|\mathbf{w}|}} \right)^2 \\
&\leq \frac{L^2 B^2}{A^2 2^{2|\mathbf{w}|}}. \tag{38}
\end{aligned}$$

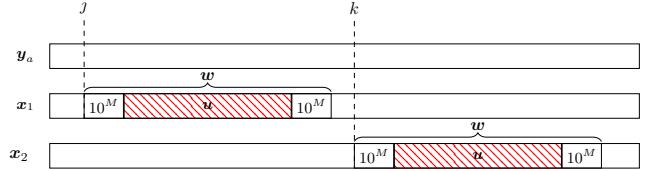


Figure 4. Relative position of the two occurrences of \mathbf{w} at position j and k when $|j - k| \geq |\mathbf{w}|$.

- $1 \leq |j - k| < |\mathbf{w}|$. The two occurrences of \mathbf{w} in \mathbf{x}_1 and \mathbf{x}_2 are plotted in Figure 5. In this case, the two occurrences of \mathbf{w} are descendants of two overlapping substrings of \mathbf{y}_a . Recall that $\mathbf{w} = 10^M \mathbf{u} 10^M$. We write the string \mathbf{u} in \mathbf{x}_1 as $\mathbf{u}_1 \mathbf{u}_2$, and write the string \mathbf{u} in \mathbf{x}_2 as $\mathbf{u}'_2 \mathbf{u}_3$, so that \mathbf{u}_2 and \mathbf{u}'_2 have the same ancestors, denoted \mathbf{r}_2 . Denote the ancestor of \mathbf{u}_1 and \mathbf{u}_3 by \mathbf{r}_1 and \mathbf{r}_3 , respectively. Denote the ancestor of 10^M at the beginning of \mathbf{w} in \mathbf{x}_1 by \mathbf{r}_0 , and the ancestor of 10^M at the end of \mathbf{w} in \mathbf{x}_2 by \mathbf{r}_4 . We have $|\mathbf{r}_1| = |\mathbf{u}_1| = |\mathbf{r}_3| = |\mathbf{u}_3| = k - j$, $|\mathbf{r}_2| = |\mathbf{u}_2| = |\mathbf{u}'_2| = |\mathbf{u}| - (k - j)$. Write $\mathbf{r} = \mathbf{r}_0 \mathbf{r}_1 \mathbf{r}_2 \mathbf{r}_3 \mathbf{r}_4$.

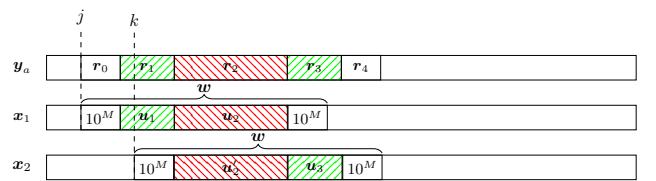


Figure 5. Relative position of the two occurrences of \mathbf{w} at position j and k when $1 \leq |j - k| < |\mathbf{w}|$.

For a single descendant \mathbf{x} of \mathbf{y}_a , \mathbf{x} can not have \mathbf{w} as substrings at positions j and k simultaneously since \mathbf{u} is

M -RLL. In other words, either exactly one of $x_{j,|\mathbf{w}|}$ and $x_{k,|\mathbf{w}|}$ equals \mathbf{w} or none of them does. So given \mathbf{r} , we can get an upper bound on the probability of $H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)$ by assuming they are independent, i.e.,

$$\begin{aligned} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k) | \mathbf{r}) \\ \leq \Pr(H_{\mathbf{w}}(a, j) | \mathbf{r}) \Pr(H_{\mathbf{w}}(a, k) | \mathbf{r}). \end{aligned} \quad (39)$$

We prove (39) rigorously by Lemma 19 at the end of this section.

Denote the Hamming distance between \mathbf{r}_0 and 10^M by d_0 , \mathbf{r}_1 and \mathbf{u}_1 by d_1 , \mathbf{r}_2 and \mathbf{u}_2 by d_2 , \mathbf{r}_2 and \mathbf{u}'_2 by d'_2 , \mathbf{r}_3 and \mathbf{u}_3 by d_3 , and \mathbf{r}_4 and 10^M by d_4 . Let $\mathbf{w}_l = 10^M \mathbf{u}$ and $\mathbf{w}_r = \mathbf{u} 10^M$. The probability of occurrences increases if we only consider substrings \mathbf{w}_l or \mathbf{w}_r . We have

$$\begin{aligned} \Pr(H_{\mathbf{w}}(a, j) | \mathbf{r}) \\ \leq \Pr(H_{\mathbf{w}_l}(a, j) | \mathbf{r}) \\ = 1 - \left(1 - \delta^{d_0+d_1+d_2} (1 - \delta)^{|\mathbf{w}|-M-1-(d_0+d_1+d_2)}\right)^{\lceil \frac{B}{4A} \rceil} \\ \leq \frac{B}{A} \delta^{d_0+d_1+d_2} (1 - \delta)^{|\mathbf{w}|-M-1-(d_0+d_1+d_2)}, \end{aligned}$$

and

$$\begin{aligned} \Pr(H_{\mathbf{w}}(a, k) | \mathbf{r}) \\ \leq \Pr(H_{\mathbf{w}_r}(a, k+M+1) | \mathbf{r}) \\ = 1 - \left(1 - \delta^{d'_2+d_3+d_4} (1 - \delta)^{|\mathbf{w}|-M-1-(d'_2+d_3+d_4)}\right)^{\lceil \frac{B}{4A} \rceil} \\ \leq \frac{B}{A} \delta^{d'_2+d_3+d_4} (1 - \delta)^{|\mathbf{w}|-M-1-(d'_2+d_3+d_4)}. \end{aligned}$$

It follows from (39) that $\Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k))$ is less

than or equal to

$$\begin{aligned} & \sum_{\mathbf{r} \in \Sigma^{|\mathbf{w}|}} \Pr(\mathbf{r}) \Pr(H_{\mathbf{w}}(a, j) | \mathbf{r}) \Pr(H_{\mathbf{w}}(a, k) | \mathbf{r}) \\ &= \left(\frac{B}{A}\right)^2 \cdot \left(\sum_{\mathbf{r}_0 \in \Sigma^{M+1}} \frac{1}{2^{|\mathbf{r}_0|}} \delta^{d_0} (1 - \delta)^{|\mathbf{r}_0|-d_0} \right) \\ & \quad \cdot \left(\sum_{\mathbf{r}_1 \in \Sigma^{k-j}} \frac{1}{2^{|\mathbf{r}_1|}} \delta^{d_1} (1 - \delta)^{|\mathbf{r}_1|-d_1} \right) \\ & \quad \cdot \left(\sum_{\mathbf{r}_3 \in \Sigma^{k-j}} \frac{1}{2^{|\mathbf{r}_3|}} \delta^{d_3} (1 - \delta)^{|\mathbf{r}_3|-d_3} \right) \\ & \quad \cdot \left(\sum_{\mathbf{r}_4 \in \Sigma^{M+1}} \frac{1}{2^{|\mathbf{r}_4|}} \delta^{d_4} (1 - \delta)^{|\mathbf{r}_4|-d_4} \right) \\ & \quad \cdot \left(\sum_{\mathbf{r}_2 \in \Sigma^{|\mathbf{u}|-(k-j)}} \frac{1}{2^{|\mathbf{r}_2|}} \delta^{d_2+d'_2} (1 - \delta)^{2|\mathbf{r}_2|-(d_2+d'_2)} \right) \\ &= \left(\frac{B}{A}\right)^2 \cdot \frac{1}{2^{2M+2+2(k-j)}} \\ & \quad \cdot \left(\sum_{\mathbf{r}_2 \in \Sigma^{|\mathbf{u}|-(k-j)}} \frac{1}{2^{|\mathbf{r}_2|}} \delta^{d_2+d'_2} (1 - \delta)^{2|\mathbf{r}_2|-(d_2+d'_2)} \right). \end{aligned}$$

Let d° denote the Hamming distance between \mathbf{u}_2 and \mathbf{u}'_2 . Among the $|\mathbf{u}_2| - d^\circ$ positions where \mathbf{u}_2 and \mathbf{u}'_2 are the same, suppose \mathbf{u}_2 differs from \mathbf{r}_2 in v of them. Among the d° positions where \mathbf{u}_2 differs from \mathbf{u}'_2 , suppose \mathbf{u}_2 differs from \mathbf{r}_2 in t of them. It follows that $d_2 = v + t$ and $d'_2 = d^\circ + v - t$. Thus, we further have

$$\begin{aligned} & \sum_{\mathbf{r}_2 \in \Sigma^{|\mathbf{u}|-(k-j)}} \frac{1}{2^{|\mathbf{r}_2|}} \delta^{d_2+d'_2} (1 - \delta)^{2|\mathbf{r}_2|-(d_2+d'_2)} \\ &= \sum_{\mathbf{r}_2 \in \Sigma^{|\mathbf{u}|-(k-j)}} \frac{1}{2^{|\mathbf{r}_2|}} \delta^{2v+d^\circ} (1 - \delta)^{2|\mathbf{r}_2|-2v-d^\circ} \\ &= \sum_{v=0}^{|r_2|-d^\circ} \binom{|r_2| - d^\circ}{v} \frac{2^{d^\circ}}{2^{|\mathbf{r}_2|}} \delta^{2v+d^\circ} (1 - \delta)^{2|\mathbf{r}_2|-2v-d^\circ} \\ &= \frac{(2\delta(1-\delta))^{d^\circ}}{2^{|\mathbf{r}_2|}} \sum_{v=0}^{|r_2|-d^\circ} \binom{|r_2| - d^\circ}{v} \\ & \quad \cdot (\delta^2)^v ((1-\delta)^2)^{(|r_2|-d^\circ)-v} \\ &= \frac{1}{2^{|\mathbf{r}_2|}} (2\delta(1-\delta))^{d^\circ} (\delta^2 + (1-\delta)^2)^{|\mathbf{r}_2|-d^\circ}. \end{aligned}$$

Since $|\mathbf{r}_2| = |\mathbf{w}| - (k - j)$,

$$\begin{aligned} & \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\ & \leq \left(\frac{B}{A}\right)^2 \frac{(2\delta(1-\delta))^{d^\circ} (\delta^2 + (1-\delta)^2)^{|\mathbf{r}_2|-d^\circ}}{2^{|\mathbf{w}|+(k-j)}}. \end{aligned}$$

Note that \mathbf{u}_2 is the $|\mathbf{r}_2|$ -suffix of \mathbf{u} and \mathbf{u}'_2 is the $|\mathbf{r}_2|$ -prefix of \mathbf{u} . With $|\mathbf{u}| = n$, the number of n -strings whose $|\mathbf{r}_2|$ -suffix and $|\mathbf{r}_2|$ -prefix are at Hamming distance d°

is $2^{n-|\mathbf{r}_2|} \binom{|\mathbf{r}_2|}{d^\circ}$ since an n -string can be uniquely determined by its $|\mathbf{r}_2|$ -prefix and the mismatches. Therefore,

$$\begin{aligned}
& \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \sum_{1 \leq |j-k| < |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& \leq \sum_{1 \leq |j-k| < |\mathbf{w}|} \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in \Sigma^n}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& \leq L|\mathbf{u}| \cdot \sum_{d^\circ=0}^{|\mathbf{r}_2|} 2^{|\mathbf{r}_1|} \binom{|\mathbf{r}_2|}{d^\circ} \cdot \left(\frac{B}{A}\right)^2 \\
& \cdot \frac{(2\delta(1-\delta))^{d^\circ} (\delta^2 + (1-\delta)^2)^{|\mathbf{r}_2|-d^\circ}}{2^{|\mathbf{w}|+k-j}} \\
& = \left(\frac{B}{A}\right)^2 \frac{Ln}{2^{|\mathbf{w}|}}. \tag{40}
\end{aligned}$$

- $|\mathbf{u}| \leq |j-k| < |\mathbf{w}|$. The two occurrences of \mathbf{w} in \mathbf{x}_1 and \mathbf{x}_2 are plotted in Figure 6.

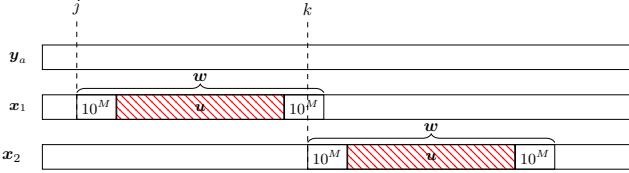


Figure 6. Relative position of the two occurrences of \mathbf{w} at position j and k when $|\mathbf{u}| \leq |j-k| < |\mathbf{w}|$.

It can be seen that the prefix $10^M \mathbf{u}$ of \mathbf{w} in \mathbf{x}_1 and \mathbf{w} in \mathbf{x}_2 are descendants of non-overlapping substrings of \mathbf{y}_a and thus independent. We can write

$$\begin{aligned}
& \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& \leq \Pr(H_{\mathbf{w}_l}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& = \Pr(H_{\mathbf{w}_l}(a, j)) \Pr(H_{\mathbf{w}}(a, k)) \\
& \leq \frac{\mathcal{S}_\delta(|\mathbf{w}_l|, \lceil \frac{B}{4A} \rceil) \mathcal{S}_\delta(|\mathbf{w}|, \lceil \frac{B}{4A} \rceil)}{2^{|\mathbf{w}_l|+|\mathbf{w}|}} \\
& \leq \left(\frac{B}{A}\right)^2 \frac{1}{2^{2|\mathbf{w}|-M-1}}.
\end{aligned}$$

It follows that

$$\begin{aligned}
& \sum_{|\mathbf{u}| \leq |j-k| < |\mathbf{w}|} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& \leq \left(\frac{B}{A}\right)^2 \frac{L \cdot 2(M+1)}{2^{2|\mathbf{w}|-M-1}}. \tag{41}
\end{aligned}$$

Thus, combining (38), (40), (41) gives

$$\begin{aligned}
& \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \sum_{\substack{1 \leq j, k \leq \lceil L/2 \rceil - |\mathbf{w}| \\ j \neq k}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \\
& \leq \sum_{\substack{\mathbf{w}: \mathbf{w} = 10^M \mathbf{u} 10^M \\ \mathbf{u} \in R_M^n}} \left(\sum_{\substack{|j-k| \geq |\mathbf{w}|}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \right. \\
& \quad \left. + \sum_{\substack{1 \leq |j-k| < |\mathbf{w}|}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \right. \\
& \quad \left. + \sum_{\substack{|\mathbf{u}| \leq |j-k| \leq |\mathbf{w}|}} \Pr(H_{\mathbf{w}}(a, j) \cap H_{\mathbf{w}}(a, k)) \right) \\
& \leq \frac{L^2 B^2}{A^2 \cdot 2^{2|\mathbf{w}|}} \cdot |R_M^n| + \frac{B^2 L n}{A^2 \cdot 2^{2|\mathbf{w}|}} + \frac{2B^2 L(M+1)}{A^2 \cdot 2^{2|\mathbf{w}|-M-1}} \cdot |R_M^n| \\
& \leq \frac{B^2 L^2}{A^2 2^{|\mathbf{w}|}} \left(1 + \frac{n+M+1}{L} \right).
\end{aligned}$$

■

We present a lemma from which inequality (39) follows directly.

Lemma 19. *Let \mathbf{r} be any string of length n with m iid δ -edit descendants. For a string \mathbf{v} , $|\mathbf{v}| < n$ and $1 \leq j < k \leq n - |\mathbf{v}| + 1$, let $\mathcal{J}(\mathbf{v}), \mathcal{K}(\mathbf{v})$ denote the events that there exists a descendant of \mathbf{r} whose j -th, k -th $|\mathbf{v}|$ -substring equal \mathbf{v} , respectively. We have*

$$\Pr(\mathcal{J}(\mathbf{v}) \cap \mathcal{K}(\mathbf{v})) \leq \Pr(\mathcal{J}(\mathbf{v})) \Pr(\mathcal{K}(\mathbf{v})),$$

if the $(|\mathbf{v}| - (k-j))$ -suffix and $(|\mathbf{v}| - (k-j))$ -prefix of \mathbf{v} are not the same.

Proof: If the $(|\mathbf{v}| - (k-j))$ -suffix and $(|\mathbf{v}| - (k-j))$ -prefix of \mathbf{v} are not the same, then in any descendant \mathbf{x} , \mathbf{v} can not be both the j -th and the k -th substring. Therefore, in \mathbf{x} , exactly one of the following three mutually exclusive events holds: i) $\mathbf{x}_{j,|\mathbf{v}|} = \mathbf{v}$, ii) $\mathbf{x}_{k,|\mathbf{v}|} = \mathbf{v}$, iii) $\mathbf{x}_{j,|\mathbf{v}|} \neq \mathbf{v}$ and $\mathbf{x}_{k,|\mathbf{v}|} \neq \mathbf{v}$. Let p_j denote the probability of $\mathbf{x}_{j,|\mathbf{v}|} = \mathbf{v}$ and p_k denote the probability of $\mathbf{x}_{k,|\mathbf{v}|} = \mathbf{v}$. We have

$$\Pr(\mathbf{x}_{j,|\mathbf{v}|} \neq \mathbf{v} \cap \mathbf{x}_{k,|\mathbf{v}|} \neq \mathbf{v}) = 1 - p_j - p_k.$$

Therefore, among the m iid descendants of \mathbf{r} ,

$$\begin{aligned}
& \Pr(\mathcal{J}(\mathbf{v}) \cap \mathcal{K}(\mathbf{v})) \\
& = \Pr(\mathcal{J}(\mathbf{v})) + \Pr(\mathcal{K}(\mathbf{v})) + \Pr(\bar{\mathcal{J}}(\mathbf{v}) \cap \bar{\mathcal{K}}(\mathbf{v})) - 1 \\
& = (1 - (1 - p_j)^m) + (1 - (1 - p_k)^m) \\
& \quad + (1 - p_j - p_k)^m - 1 \\
& = 1 - (1 - p_j)^m - (1 - p_k)^m + (1 - p_j - p_k)^m.
\end{aligned}$$

On the other hand,

$$\begin{aligned}
& \Pr(\mathcal{J}(\mathbf{v})) \Pr(\mathcal{K}(\mathbf{v})) \\
& = (1 - (1 - p_j)^m)(1 - (1 - p_k)^m) \\
& = 1 - (1 - p_j)^m - (1 - p_k)^m \\
& \quad + (1 - p_j)^m (1 - p_k)^m.
\end{aligned}$$

The desired inequality thus follows by noting that $1 - p_j - p_k \leq (1 - p_j)(1 - p_k)$. ■

Inequality (39) can be obtained by replacing $\mathcal{J}(v)$ and $\mathcal{K}(v)$ with $H_w(a, j)$ and $H_w(a, k)$, respectively.

APPENDIX F PROOFS OF LEMMA 13 AND LEMMA 16

Lemma 13. Consider the source string $s = Y_1 Y_2 \dots Y_B$. When $2^M = o(L)$, for B, L sufficiently large,

$$\Pr\left(C_{VL}^M(s) \geq \frac{1}{4} \cdot \left\lfloor \frac{B}{2} \right\rfloor \left(\frac{L}{2^{M+8}} - 1 \right)\right) \geq \frac{5}{6}.$$

Proof: Equally parse each of $Y_{\lceil B/2 \rceil + 1}, \dots, Y_B$ into segments of length 2^{M+7} . So that every Y_b contains $\left\lfloor \frac{|Y_b|}{2^{M+7}} \right\rfloor$ segments. We show that among these $\sum_{b=\lceil \frac{B}{2} \rceil + 1}^B \left\lfloor \frac{|Y_b|}{2^{M+7}} \right\rfloor$ segments, a constant fraction of them contain a chunk of length over 2^{M-4} .

Pick an arbitrary segment, denoted z . Consider the two halves of z . The second half of z , which is of length 2^{M+6} , is by itself a Bernoulli(1/2) process going forward. We study the first time a run of M 0's appears in this process. By the union bound, with probability at least $1 - \frac{2^{M-5}}{2^M}$, there exist no runs of M 0s in the first 2^{M-5} bits. Moreover, the average position of the end of the first run of M 0s in a Bernoulli(1/2) process is $2^{M+1} - 2$ [15]. Therefore, by Markov's inequality, with probability at least $1 - \frac{2^{M+1}-2}{2^{M+6}}$, there is a 0^M within the first 2^{M+6} bits. So the first time we see 0^M is after 2^{M-5} bits and before 2^{M+6} bits (i.e., the first 0^M is within the last $2^{M+6} - 2^{M-5}$ bits) with probability at least

$$1 - \frac{2^{M-5}}{2^M} - \frac{2^{M+1}-2}{2^{M+6}} \geq 1 - \frac{1}{24}.$$

Similarly, the first half of z can be regarded as a reversed Bernoulli(1/2) process. So we also have with probability at least $1 - \frac{1}{24}$, the first 0^M (counting backwards) is within the first $2^{M+6} - 2^{M-5}$ bits. Clearly, a chunk exists between these two occurrences of 0^M . So with probability at least $1 - \frac{1}{23}$, z contains a chunk of length at least 2^{M-4} . Since this property holds for all such segments of length 2^{M+7} , by the Markov inequality, with probability at least $1 - \frac{1}{6}$, at least $\frac{1}{4}$ of the segments in $Y_{\lceil B/2 \rceil + 1} \dots Y_B$ contain a chunk of length at least 2^{M-4} . The desired result is derived by noting $|Y_b| \geq L/2$. ■

Lemma 16. Consider the source string $s = Y_1 Y_2 \dots Y_B$, with each Y_b being a descendant of source symbol X_{J_b} . For any integer h and any pairs of integers $(b_1, b_2), (i_1, i_2)$, the probability of Y_{b_1} and Y_{b_2} having identical substrings of length h starting at positions i_1 and i_2 , respectively, is

$$\Pr\left((Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h}\right) = \frac{1}{2^h},$$

if $J_{b_1} \neq J_{b_2}$ or $i_1 \neq i_2$.

Proof: We compute $\Pr((Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h})$ as $(b_1, b_2), (i_1, i_2)$ take different values in the following three cases:

- $J_{b_1} \neq J_{b_2}$ or $|i_1 - i_2| \geq h$. If $J_{b_1} \neq J_{b_2}$, then Y_{b_1} and Y_{b_2} have different ancestors and are thus independent. It follows that their substrings are also independent. If $|i_1 - i_2| \geq h$, then $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are descendants of non-overlapping substrings of the source

alphabet and are thus also independent. The desired result follows from the fact that $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are both Bernoulli(1/2) processes by themselves.

- $b_1 = b_2, |i_1 - i_2| < h$. In this case, $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are overlapping substrings of a single source block. Again, Y_{b_1} is Bernoulli(1/2) by itself. So the probability of $(Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h}$ is the same as that when $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are independent.
- $J_{b_1} = J_{b_2}, b_1 \neq b_2, |i_1 - i_2| < h$. Let $J_{b_1} = J_{b_2} = a$. Assume $i_1 < i_2$ without loss of generality. In this case, $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are two independent δ -edit descendants of $(X_a)_{i_1, h}$ and $(X_a)_{i_2, h}$, respectively. So $\Pr((Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h})$ is uniquely determined by the Hamming distance between $(X_a)_{i_1, h}$ and $(X_a)_{i_2, h}$. Moreover, the distribution of the Hamming distance between $(X_a)_{i_1, h}$ and $(X_a)_{i_2, h}$ is the same as the distribution of the Hamming distance between two independent Bernoulli(1/2) process of length h . Therefore, we can assume $(Y_{b_1})_{i_1, h}$ and $(Y_{b_2})_{i_2, h}$ are independent and thus $\Pr((Y_{b_1})_{i_1, h} = (Y_{b_2})_{i_2, h}) = \frac{1}{2^h}$.

APPENDIX G SUMMATIONS

For integers $b \geq a$ and $\beta > 1$, summations of the forms $\sum_{n=a}^b \left(1 - \frac{1}{\beta}\right)^n$ and $\sum_{n=a}^b n \left(1 - \frac{1}{\beta}\right)^n$ appear in the proofs of Theorem 15 and Theorem 18. Let $x = 1 - \frac{1}{\beta}$. The limits of these sums in a certain asymptotic regime is discussed below.

A. Asymptotic behavior of $\sum_{n=a}^b x^n$

We have

$$\begin{aligned} \sum_{n=a}^b x^n &= \frac{x^a (1 - x^{b-a+1})}{1 - x} \\ &= \beta \left(1 - \frac{1}{\beta}\right)^a \left(1 - \left(1 - \frac{1}{\beta}\right)^{b-a+1}\right). \end{aligned}$$

If $b - a = \omega(\beta)$, then as $\beta \rightarrow \infty$,

$$\left(1 - \frac{1}{\beta}\right)^{b-a+1} = \left(\left(1 - \frac{1}{\beta}\right)^{\beta}\right)^{\frac{b-a+1}{\beta}} = o(1).$$

It follows that

$$\begin{aligned} \sum_{n=a}^b \left(1 - \frac{1}{\beta}\right)^n &= \beta \left(1 - \frac{1}{\beta}\right)^a (1 + o(1)) \\ &= \beta e^{-a/\beta} (1 + o(1)). \end{aligned} \tag{42}$$

B. Asymptotic behavior of $\sum_{n=a}^b nx^n$

We have

$$\begin{aligned}
 \sum_{n=a}^b nx^n &= x \sum_{n=a}^b nx^{n-1} = x \left(\sum_{n=a}^b x^n \right)' \\
 &= x \left(\frac{x^a (1 - x^{b-a+1})}{1 - x} \right)' \\
 &= x \frac{(ax^{a-1} - (b+1)x^b)(1-x) + (x^a - x^{b+1})}{(1-x)^2} \\
 &= \beta^2 \left(\left(\frac{a-1}{\beta} + 1 \right) \left(1 - \frac{1}{\beta} \right)^a + \left(\frac{b}{\beta} + 1 \right) \left(1 - \frac{1}{\beta} \right)^{b+1} \right).
 \end{aligned}$$

If $\frac{b}{\beta} = \omega(1)$, then as $\beta \rightarrow \infty$,

$$\begin{aligned}
 \left(\frac{b}{\beta} + 1 \right) \left(1 - \frac{1}{\beta} \right)^{b+1} &= \left(\frac{b}{\beta} + 1 \right) \left(\left(1 - \frac{1}{\beta} \right)^\beta \right)^{\frac{b+1}{\beta}} \\
 &= o(1).
 \end{aligned}$$

It follows that

$$\begin{aligned}
 \sum_{n=a}^b n \left(1 - \frac{1}{\beta} \right)^n &= \beta^2 \left(\left(\frac{a-1}{\beta} + 1 \right) \left(1 - \frac{1}{\beta} \right)^a + o(1) \right) \\
 &= \beta^2 \left(\frac{a-1}{\beta} + 1 \right) e^{-a/\beta} (1 + o(1)). \quad (43)
 \end{aligned}$$

REFERENCES

- [1] L. Conde-Canencia, T. Condie, and L. Dolecek, “Data deduplication with edit errors”, in *2018 IEEE Global Communications Conference (GLOBECOM)*, IEEE, 2018, pp. 1–6.
- [2] P. Elias, “Universal codeword sets and representations of the integers”, *IEEE transactions on information theory*, vol. 21, no. 2, pp. 194–203, 1975.
- [3] J. Gantz and D. Reinsel, “The digital universe in 2020: Big data, bigger digital shadows, and biggest growth in the far east”, *IDC iView: IDC Analyze the future*, vol. 2007, no. 2012, pp. 1–16, 2012.
- [4] S. Greenberg and M. Mohri, “Tight lower bound on the probability of a binomial exceeding its expectation”, *Statistics & Probability Letters*, vol. 86, pp. 91–98, 2014.
- [5] K. Jin and E. L. Miller, “The effectiveness of deduplication on virtual machine disk images”, in *Proceedings of SYSTOR 2009: The Israeli Experimental Systems Conference*, 2009, pp. 1–12.
- [6] E. S. Lander, L. M. Linton, B. Birren, C. Nusbaum, M. C. Zody, J. Baldwin, K. Devon, K. Dewar, M. Doyle, W. FitzHugh, *et al.*, “Initial sequencing and analysis of the human genome”, *Nature*, vol. 409, no. 6822, pp. 860–921, 2001.
- [7] H. Lou and F. Farnoud, “Data deduplication with random substitutions”, in *2020 IEEE International Symposium on Information Theory (ISIT)*, IEEE, 2020, pp. 2377–2382.
- [8] U. Manber, “Finding similar files in a large file system”, in *Usenix Winter*, vol. 94, 1994, pp. 1–10.
- [9] B. H. Marcus, R. M. Roth, and P. H. Siegel, “An introduction to coding for constrained systems”, *Lecture notes*, 2001.
- [10] D. T. Meyer and W. J. Bolosky, “A study of practical deduplication”, *ACM Transactions on Storage*, vol. 7, no. 4, p. 14, 2012.
- [11] M. Mitzenmacher and E. Upfal, *Probability and computing: randomization and probabilistic techniques in algorithms and data analysis*. Cambridge university press, 2017.
- [12] A. Muthitacharoen, B. Chen, and D. Mazieres, “A low-bandwidth network file system”, in *ACM SIGOPS Operating Systems Review*, ACM, vol. 35, 2001, pp. 174–187.
- [13] U. Niesen, “An information-theoretic analysis of deduplication”, *IEEE Transactions on Information Theory*, vol. 65, no. 9, pp. 5688–5704, Sep. 2019.
- [14] S. Quinlan and S. Dorward, “Venti: A new approach to archival storage”, in *FAST*, vol. 2, 2002, pp. 89–101.
- [15] R. Sedgewick and P. Flajolet, *An introduction to the analysis of algorithms*. Pearson Education India, 2013.
- [16] P. Shilane, G. Wallace, M. Huang, and W. Hsu, “Delta compressed and deduplicated storage using stream-informed locality.”, in *HotStorage*, 2012.
- [17] P. Shilane, M. Huang, G. Wallace, and W. Hsu, “Wan-optimized replication of backup datasets using stream-informed delta compression”, *ACM Transactions on Storage (ToS)*, vol. 8, no. 4, pp. 1–26, 2012.
- [18] A. El-Shimi, R. Kalach, A. Kumar, A. Ottean, J. Li, and S. Sengupta, “Primary data deduplication: large scale study and system design”, in *Presented as part of the 2012 USENIX Annual Technical Conference (USENIX ATC 12)*, 2012, pp. 285–296.
- [19] R. Vestergaard, Q. Zhang, and D. E. Lucani, “Generalized deduplication: Bounds, convergence, and asymptotic properties”, *arXiv preprint arXiv:1901.02720*, 2019.
- [20] W. Xia, H. Jiang, D. Feng, F. Dougulis, P. Shilane, Y. Hua, M. Fu, Y. Zhang, and Y. Zhou, “A comprehensive study of the past, present, and future of data deduplication”, *Proceedings of the IEEE*, vol. 104, no. 9, pp. 1681–1710, 2016.
- [21] W. Xia, H. Jiang, D. Feng, and L. Tian, “Combining deduplication and delta compression to achieve low-overhead data reduction on backup datasets”, in *2014 Data Compression Conference*, IEEE, 2014, pp. 203–212.
- [22] W. Xia, H. Jiang, D. Feng, L. Tian, M. Fu, and Y. Zhou, “Ddelta: A deduplication-inspired fast delta compression approach”, *Performance Evaluation*, vol. 79, pp. 258–272, 2014.