

Coding for Deletion Channels with Multiple Traces

Mahed Abroshan
University of Cambridge
ma675@cam.ac.uk

Ramji Venkataramanan
University of Cambridge
rv285@cam.ac.uk

Lara Dolecek
ECE Department, UCLA
dolecek@ee.ucla.edu

Albert Guillén i Fàbregas
ICREA & Universitat Pompeu Fabra
University of Cambridge
guillen@ieee.org

Abstract—Motivated by the sequence reconstruction problem from traces in DNA-based storage, we consider the problem of designing codes for the deletion channel when multiple observations (or traces) are available to the decoder. We propose simple binary and non-binary codes based on Varshamov-Tenengolts (VT) codes. The proposed codes split the codeword in blocks and employ a VT code in each block. The availability of multiple traces helps the decoder to identify deletion-free copies of a block, and to avoid mis-synchronization while decoding. The encoding complexity of the proposed scheme is linear in the codeword length; the decoding complexity is linear in the codeword length, and quadratic in the number of deletions and the number of traces. The proposed scheme offers an explicit low-complexity technique for correcting deletions using multiple traces.

I. INTRODUCTION

We consider the problem of coding for a deletion channel which, given an input sequence, produces t output sequences. Each output sequence, known as a trace, is produced by deleting k symbols from the length n input sequence. The deletion pattern for each trace is assumed to be independent, and the locations of the deletions within a trace are uniformly random. (Therefore it is possible that two or more traces are identical.)

The problem of recovering coded information from multiple traces is relevant in DNA-based storage systems [1]. When retrieving information by sequencing stored DNA, each trace may contain errors that are a combination of deletions, insertions, and substitutions. In this paper, using the stylized model of a channel that introduces only deletions, we aim to understand the coding advantage obtained by having multiple traces. In particular, we show how one can use simple codes with efficient encoding and decoding to achieve small probability of error under suitable assumptions.

The problem of reconstructing a sequence using traces from a deletion channel has been studied in several works, with the goal being either exact recovery of the sequence [2], [3], [4], or an estimate [5]. In these papers, the sequence can be an arbitrary one from the underlying alphabet, i.e., it need not originate from a codebook. A few recent works study the reconstruction of a *coded* sequence under different trace models. The paper [6] analyzes the minimum number of deletion channel traces required to recover a sequence drawn from a single-deletion correcting code, and [7] considers a

similar problem for the insertion channel. The problem of reconstructing a coded sequence from the multiset of its substrings is studied in [8]. From an information-theoretic perspective, [9] characterizes the capacity of the multiple trace i.i.d. deletion channel as the deletion probability $p \rightarrow 0$.

Overview of the coding scheme: Our code construction is based on the family of Varshamov-Tenengolts (VT) codes [10], [11]. VT codes are single deletion correcting codes, and can be constructed for any finite alphabet and any block length. (VT codes will be reviewed in the next section). Each codeword in our code is a concatenation of blocks, with each block drawn from a predetermined VT codebook.

We illustrate the idea using the following binary example, which shows a codeword x of length 15 with three blocks, each of which is a sequence from a length 5 binary VT code. The channel produces two traces, y_1 and y_2 , by deleting the underlined bits:

$$\begin{aligned} 10001 \ 110\underline{11} \ 01010 &\longrightarrow y_1 = 1001110101010 \\ 10001 \ 1\underline{1}011 \ 0101\underline{0} &\longrightarrow y_2 = 100011110101 \end{aligned}$$

The decoder operates in two phases. In phase 1, it identifies blocks that are deletion-free in at least one of the traces. Each block for which a deletion-free copy is identified in one of the traces is recovered by inserting the required bits in the other traces. In the example above, block 1 has no deletions in y_2 , so y_2 is used to correct the first block of y_1 ; similarly, block 3 has no deletions in y_1 . Assuming that there were no errors, at the end of this phase the decoder has corrected all blocks which are deletion-free in at least one trace. We call the remaining blocks ‘congested’. In the example, block 2 is congested as both traces have bits deleted in this block.

In phase 2, the decoder attempts to correct the congested blocks, i.e., blocks for which no clean copy was found in phase 1. In the example, since y_1 has one deletion in block 2, the block can be corrected using the VT decoder. Since blocks 1 and 3 were corrected in phase 1, x is recovered.

However, decoding errors may occur in either phase. In phase 1, we may wrongly identify a block as deletion-free in a trace, which leads to errors in the starting positions of other blocks. In phase 2, a congested block (or set of consecutive congested blocks) may not be correctable with the VT code, because of too many deletions in each trace. In phase 1, wrongly identifying a trace as having a deletion-free copy of a block will lead to an unusually large number of insertions when correcting the other traces using this copy. This can be

This work was supported in part by the European Research Council under Grant 725411, by the Spanish Ministry of Economy and Competitiveness under Grant TEC2016-78434-C3-1-R, and by NSF grants CCF-1718389 and CCF-1527130.

used to discard accidental matches in phase 1. We show via numerical simulations that the probability of phase 1 error decreases with n_b , the length of each block. For the phase 2 error, under the assumption that the locations of the deletions within each trace are uniformly random, we obtain a bound that decreases exponentially with the number of traces.

The rate of the code is equal to the rate of a VT code of length n_b , which is close to $\log q - \log n_b/n_b$, where $q \geq 2$ is the alphabet size. (The precise values are given in Section II.) The decoding complexity is $\mathcal{O}(l^2 k^2 n)$. Therefore the proposed scheme offers an explicit, efficient technique for recovering from deletions using multiple traces. Due to its low-complexity, it can be well suited to a variety of applications, including DNA-based storage.

Notation: Sequences are denoted using bold letters, and scalars with plain font. For $\mathbf{x} = x_1 x_2 \cdots x_n$, the subsequence $x_i x_{i+1} \cdots x_j$ is denoted by $\mathbf{x}(i : j)$.

II. CODE CONSTRUCTION

We first review the family of VT codes, and then present the code construction.

Binary VT codes: The VT syndrome of a binary sequence $\mathbf{x} = x_1 x_2 \cdots x_n$ is defined as

$$\mathbf{syn}(\mathbf{x}) = \sum_{j=1}^n j x_j \pmod{(n+1)}. \quad (1)$$

For positive integers n and $0 \leq s \leq n$, we define the VT code of length n and syndrome s , denoted by

$$\mathcal{VT}_s(n) = \{\mathbf{x} \in \{0, 1\}^n : \mathbf{syn}(\mathbf{x}) = s\}, \quad (2)$$

as the set of sequences \mathbf{x} of length n for which $\mathbf{syn}(\mathbf{x}) = s$. The $(n+1)$ sets $\mathcal{VT}_s(n) \subset \{0, 1\}^n$, for $0 \leq s \leq n$, partition the set of all binary sequences of length n . Each of these sets $\mathcal{VT}_s(n)$ is a zero-error single-deletion correcting code. The VT encoding and decoding complexity is linear in the code length n [10], [12].

Non-binary VT codes: VT codes were extended to non-binary alphabets in [11]. Let the alphabet be $\mathcal{X} = \{0, \dots, q-1\}$, with $q > 2$. For each sequence $\mathbf{x} \in \mathcal{X}^n$, define a length $(n-1)$ auxiliary binary sequence $\alpha(\mathbf{x}) = \alpha_2, \dots, \alpha_n$ as follows. For $2 \leq j \leq n$,

$$\alpha_j = \begin{cases} 1 & \text{if } x_j \geq x_{j-1} \\ 0 & \text{if } x_j < x_{j-1} \end{cases} \quad (3)$$

We also define the modular sum as

$$\text{sum}(\mathbf{x}) = \sum_{j=1}^n x_j \pmod{q}. \quad (4)$$

A q -ary VT code with length n and parameters (a, c) is defined as [11]

$$\mathcal{VT}_{a,c}(n) = \{\mathbf{x} \in \mathcal{X}^n : \mathbf{syn}(\alpha(\mathbf{x})) = a, \text{sum}(\mathbf{x}) = c\}, \quad (5)$$

for $0 \leq a \leq n-1$ and $c \in \mathcal{X}$. Similarly to the binary case, the sets $\mathcal{VT}_{a,c}(n)$ partition the space \mathcal{X}^n of all q -ary sequences of length n into qn sets. Each set is a single deletion correcting

code. The complexity of the encoding and decoding q -ary VT codes is linear in the code length n [11], [13].

A. Code construction

Codewords of length n are constructed by concatenating l blocks of VT codewords from the relevant alphabet. Each block has length n_b (therefore $n = l n_b$).

Binary code: Each block i ($1 \leq i \leq l$) is a binary VT codeword with a predetermined VT syndrome a_i , known to both the encoder and the decoder. To encode each block, one can use the systematic VT encoder in [12] that maps $n_b - \lceil \log(n_b + 1) \rceil$ bits to a length n_b VT sequence with the desired syndrome. The rate of the code will be

$$R = 1 - \lceil \log(n_b + 1) \rceil / n_b. \quad (6)$$

Non-binary code: The code construction is very similar to the binary case. Each block is encoded separately, and belongs to a known non-binary VT class, as defined in (5). There are qn_b non-binary VT classes, so there exists a class with at least $\frac{q^{n_b}}{qn_b}$ sequences. Using this class for encoding each block induces the following lower bound for the rate of the code:

$$R \geq \log q - (\log n_b + \log q) / n_b. \quad (7)$$

The VT syndrome/class for each of the l blocks can be arbitrarily chosen. Choosing the largest of the VT classes of length n_b sequences maximises the rate of the code.

III. DECODING

The goal of the decoder is to reconstruct \mathbf{x} using the traces $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_t$, each obtained by deleting k symbols from \mathbf{x} . We describe the binary decoder, and then outline the main differences for the non-binary case. We explain the main ideas using examples. A detailed description of the decoding algorithm is given in [14]. We remark that the same decoding algorithm can be applied when the number of deletions in one or more traces is less than k .

A. Phase 1

Consider block i of the codeword, for $1 \leq i \leq l$. If the starting position of block i within each trace is known, then the decoder can compute the VT syndrome of the length n_b sequence from the starting position, for each trace. If there is a trace for which the computed syndrome matches with a_i (the correct syndrome for block i), then the trace can be used to correct this block within other traces. The following example illustrates this idea.

Example 1: Consider a codeword with $l = 3$ blocks, with each block of length $n_b = 5$. Thus $n = 15$. Let

$$\mathbf{x} = 01001 \ 11001 \ 11111$$

be the transmitted codeword. The VT syndromes of the blocks are $a_1 = 1$, $a_2 = 2$, and $a_3 = 3$. Suppose that the decoder receives two traces, each with $k = 2$ deletions. The underlined bits are deleted from \mathbf{x} to produce \mathbf{y}_1 and \mathbf{y}_2 :

$$01001 \ 11001 \ 11111 \longrightarrow \mathbf{y}_1 = 0111100111111$$

$$01001 \underline{11001} \ 11111 \longrightarrow \mathbf{y}_2 = 0100110011111$$

The decoder first computes the VT syndromes of $\mathbf{y}_1(1 : 5)$ and $\mathbf{y}_2(1 : 5)$. We have

$$\text{syn}(\mathbf{y}_1(1 : 5)) = 2 \quad \text{and} \quad \text{syn}(\mathbf{y}_2(1 : 5)) = 1.$$

Since $a_1 = 1$, the decoder assumes that $\mathbf{y}_2(1 : 5)$ is the first block of \mathbf{x} , and uses it to correct the first block of \mathbf{y}_1 by inserting the two missing bits. The decoder then considers the second block, whose starting position is now known for each trace. Finding that $\text{syn}(\mathbf{y}_1(4 : 8)) = a_2 = 2$, it assumes this sequence is the second block of \mathbf{x} , and uses it to correct \mathbf{y}_2 by inserting two bits. Since there are no deletions in the third block, the decoder finds $\text{syn}(\mathbf{y}_1(9 : 13)) = \text{syn}(\mathbf{y}_2(9 : 13)) = a_3 = 3$, and stops.

A decoding error may occur if the VT syndrome of a block in a trace accidentally matches the correct value. In this case, an incorrect sequence will be used to correct the block in all other traces, potentially introducing multiple errors. The following example shows that how other traces can help to identify and discard such accidental matches.

Example 2: Consider a codeword with $l = 2$ blocks, with each block of length $n_b = 10$. The VT syndromes of the blocks are $a_1 = a_2 = 5$. Let the transmitted codeword be

$$\mathbf{x} = 1000111100 \ 0011101100.$$

There are two traces, each with $k = 2$ deletions (underlined bits are deleted):

$$\begin{aligned} 100\underline{0}111\underline{1}00 \ 0011011100 &\longrightarrow \mathbf{y}_1 = 100111000011011100 \\ \underline{1}000111100 \ 001\underline{1}011100 &\longrightarrow \mathbf{y}_2 = 000111100001011100 \end{aligned}$$

The decoder finds that $\text{syn}(\mathbf{y}_1(1 : 10)) = 5$ (this is an instance of an accidental match), and $\text{syn}(\mathbf{y}_2(1 : 10)) = 0$. It assumes that $\mathbf{y}_1(1 : 10)$ is the correct block, and uses it to fix \mathbf{y}_2 . It does this by comparing $\mathbf{y}_1(1 : 10)$ with \mathbf{y}_2 , and inserting the required bits to get $\hat{\mathbf{y}}_2(1 : 10) = \underline{1}00\underline{1}1\underline{1}0000$. Since there are two deletions in \mathbf{y}_2 , exactly two inserted bits are required to recover the codeword. However, since 7 bits need to be inserted into \mathbf{y}_2 to get $\mathbf{y}_1(1 : 10)$ and $k = 2$, the decoder realizes that $\mathbf{y}_1(1 : 10)$ is an accidental match.

The above example shows that when an accidentally matched block is used as the model to correct other traces, the number of inserted bits is likely to be large. Hence the decoder can distinguish between an accidental match and a correct match in most cases.

Congested blocks and resynchronization. There may be blocks that have undergone at least one deletion in each of the traces. These blocks are called *congested*, as a correct match for them cannot be found in any of the traces. In Example 2, the first block is congested as there are deleted bits in both the traces. As the decoder proceeds from left to right in phase 1, it needs to resynchronize whenever it identifies a congested block. It does so by testing all possible starting positions for the next block in each trace.

Assume block i is congested, and consider a trace for which that the decoder has inferred that there are $d < k$ deletions up

to block $(i - 1)$. The decoder needs to test $(k - d)$ possible starting positions for block $(i + 1)$ in this trace. It computes the VT syndromes of the length n_b sequences starting from each of these positions, and checks for a match with the correct syndrome a_{i+1} . If a match is found, it is used to correct the other traces. It repeats this process for each trace, testing all possible starting positions for block $(i + 1)$, and checking whether a match is found for the correct VT syndrome. If the decoder finds one or more syndrome matches among those tested, it chooses the one that requires the minimum number of insertions (across all traces) for correcting block $(i + 1)$.

When block i is identified as congested, it is possible that block $(i + 1)$ is also congested (i.e., has deletions in all the traces). In this case, no matches may be found among all the tested starting positions for block $(i + 1)$. The decoder then tries to synchronize by testing all possible starting positions for block $(i + 2)$.

B. Phase 2

At the end of phase 1, if there are no errors, the decoder has corrected all blocks for which there is at least one trace with a deletion-free copy of the block. Each remaining block is congested, and is either: i) isolated, i.e., the bits corresponding to the block in each trace are known, or ii) part of an isolated set of consecutive congested blocks.

In the second phase, the decoder uses the VT syndromes to correct as many congested blocks as possible. For each congested set of r consecutive blocks ($r \geq 1$, with $r = 1$ corresponding to a single congested block), the decoder can infer the number of deleted bits within each trace. It uses this information, and attempts to correct the congested blocks as follows. For a congested set of r consecutive blocks ($r \geq 1$), the decoder looks for a trace with exactly r deletions. If such a trace exists, then this set of blocks can be corrected using that trace and the known VT syndromes of the r blocks. On the other hand, a congested set of r consecutive blocks cannot be corrected if it has at least $(r + 1)$ deletions in each trace.

Example 3: Consider a codeword with $l = 4$ blocks, each of length $n_b = 5$. The VT syndromes are $a_1 = a_2 = a_3 = a_4 = 0$, and

$$\mathbf{x} = 11100 \ 10001 \ 10001 \ 01010.$$

There are two traces, with 4 deletions in the first and 3 in the second:

$$\begin{aligned} 1110\underline{0} \ 1000\underline{1} \ 1000\underline{1} \ 0\underline{1}0\underline{1}0 &\rightarrow \mathbf{y}_1 = 1110100110001101 \\ 11100 \ 100\underline{0}1 \ 1000\underline{1} \ 0\underline{1}0\underline{1}0 &\rightarrow \mathbf{y}_2 = 11100100110001000. \end{aligned}$$

The first block is recovered using \mathbf{y}_2 , using which the block is corrected in \mathbf{y}_1 . The second block is congested, and neither trace provides a match for its VT syndrome. The decoder therefore tests the possible starting positions for the third block. Consider the first trace, which has a total of 4 deletions. Since there was one deletion in the first block, there are three possible starting positions for the third block: bits 7, 8 and 9 of \mathbf{y}_1 . Similarly, bits 8, 9 and 10 of \mathbf{y}_2 are the possible starting positions for the third block.

The decoder therefore computes the VT syndrome of $\mathbf{y}_1(9 : 13)$, $\mathbf{y}_1(8 : 12)$, $\mathbf{y}_1(7 : 12)$, and $\mathbf{y}_2(10 : 14)$, $\mathbf{y}_2(9 : 13)$, $\mathbf{y}_2(8 : 12)$. Among these, the only one that satisfies the correct syndrome $a_3 = 0$ is $\mathbf{y}_1(9 : 13) = \mathbf{y}_2(10 : 14) = 10001$. This indicates that there is one deletion in the second block, in each of the traces. Thus the second block can be recovered using the VT decoder in phase 2. With the first three blocks synchronized, the decoder attempts to correct the fourth. The fourth block has two deletions in both traces. As the VT decoder can only correct a single deletion, the decoder declares an error due to an unresolvable congestion.

In the next section, we derive a bound (Proposition 1) on the probability of phase 2 error, caused by an unresolvable congestion like the one above.

C. Non-binary alphabet

The decoding is similar to the binary case. The only difference is that the VT syndrome of a non-binary sequence is a pair of numbers. Therefore, when we comparing VT syndromes of two sequences in the first phase, both numbers in the pair should be compared. In the second phase, the decoder uses the non-binary VT decoder from [11] to recover a block with a single deletion.

D. Decoding complexity

In phase 1, for a block for which the starting position is unknown, the decoder computes at most k VT syndromes of length n_b sequences in each of the t traces. For each matched syndrome, the decoder needs to check inserted bits in at most k blocks in the other $(t - 1)$ traces. Since there are l blocks, and $n = n_b l$, the complexity for the first phase is $\mathcal{O}(t^2 k^2 n)$.

In phase 2, the VT decoder is used in at most l blocks (each of length n_b), and then uses the recovered sequence to correct the block in the other traces. Since the VT decoder has linear complexity, the complexity for phase 2 is $\mathcal{O}(tn)$.

IV. ERROR PROBABILITY AND SIMULATION RESULTS

A. Phase 1 errors

In the first phase of decoding, an error can occur in two ways. First, an accidental match may lead to a block being wrongly identified as deletion-free in a trace; this is then used to correct the block in other traces. Second, when a congested block is identified, the decoder may pick a wrong starting position for the next block. As shown in Example 2, an accidental match in a trace can be often detected by the decoder when it leads to a large number of inserted bits in the other traces. This detection feature makes it hard to derive a rigorous bound for the phase 1 error.

Without the detection feature, the probability of an accidental VT match in the binary case will be inversely proportional to n_b , the length of the block. Indeed, the family of $(n_b + 1)$ VT codes partitions the space of length- n_b binary sequences into approximately equal-sized sets of size $\sim 2^{n_b}/(n_b + 1)$. Hence the probability that a binary sequence picked uniformly at random will match a given VT syndrome is close to $\frac{1}{(n_b + 1)}$.

B. Phase 2 errors

Errors in the second phase of the decoding are due to unresolvable congestion. Recall that unresolvable congestion occurs if, for some $1 \leq r \leq k$, there is a set of r consecutive congested blocks with at least $(r + 1)$ deletions in each trace. The following proposition bounds the probability of phase 2 error, denoted by P_{e_2} .

Proposition 1: Consider a code with l blocks, and a channel that introduces at most k deletions in each of the t traces. If $k < l$ and the locations of deletions within each trace are uniformly random, the probability of phase 2 error satisfies

$$P_{e_2} \leq l \left((1 - p_0 - p_1)^t + ((1 - p_0)^2 - p_1 p_1')^t + \frac{(1 - p_0)^{3t}}{1 - (1 - p_0)^t} \right)$$

where, for $s = 0, 1$,

$$p_s = \frac{\binom{(k-s)+(l-1)-1}{k-s}}{\binom{k+l-1}{k}}, \quad \text{and} \quad p_1' = \frac{\binom{(k-2)+(l-2)-1}{k-2}}{\binom{(k-1)+(l-1)-1}{k-1}}. \quad (8)$$

We note that the phase 2 error probability (and the bound) depends only on k and l , and on neither n_b nor the alphabet size. The restriction $k < l$ is natural, as otherwise the expected number of deletions per block would be greater than 1.

Proof: For $1 \leq i \leq l$ and $1 \leq r \leq k$, let $Z_{i,r}$ be an indicator random variable with $Z_{i,r} = 1$ if the i th block is in an unresolvable congestion of exactly r consecutive blocks, and $Z_{i,r} = 0$ otherwise. Let $Z = \sum_{r=1}^{k-1} \sum_{i=1}^l \frac{1}{r} Z_{i,r}$. For each r , the inner sum in Z counts the number of distinct sets of r consecutive congested blocks. Therefore Z is the total number of distinct congested sets, where a congested set is a set of r consecutive congested blocks, for some $r \geq 1$. Hence $P_{e_2} = \mathbb{P}(Z \geq 1)$. Using Markov's inequality,

$$\mathbb{P}(Z \geq 1) \leq \mathbb{E}[Z] = \sum_{i=1}^l \sum_{r=1}^{k-1} \frac{1}{r} \mathbb{E}[Z_{i,r}] \quad (9)$$

The probability that a given block has exactly s deletions ($0 \leq s \leq k$) is given by p_s in (8), which is the proportion of non-negative integer solutions of $x_1 + \dots + x_l = k$ with $x_1 = s$. A block can be in an unresolvable congestion only if it has more than one deletion in each of the traces. Therefore,

$$\mathbb{E}[Z_{i,1}] \leq (1 - p_0 - p_1)^t. \quad (10)$$

We then obtain the following upper bounds for $r \geq 2$:

$$\mathbb{E}[Z_{i,2}] \leq 2 \left((1 - p_0)^2 - p_1 p_1' \right)^t, \quad (11)$$

$$\mathbb{E}[Z_{i,r}] \leq r(1 - p_0)^{tr}, \quad r > 2. \quad (12)$$

The proofs of these bounds is given in [14]. Using (10), (11), and (12) in (9) yields the result of the proposition. ■

C. Numerical Simulations

Figure 1 shows the empirical error probability of the code for different values of n_b , for $q = 2$ (binary) and for $q = 4$. Each codeword consists of $l = 7$ blocks, each of length n_b . There are $t = 5$ traces, each with $k = 4$ deletions at uniformly

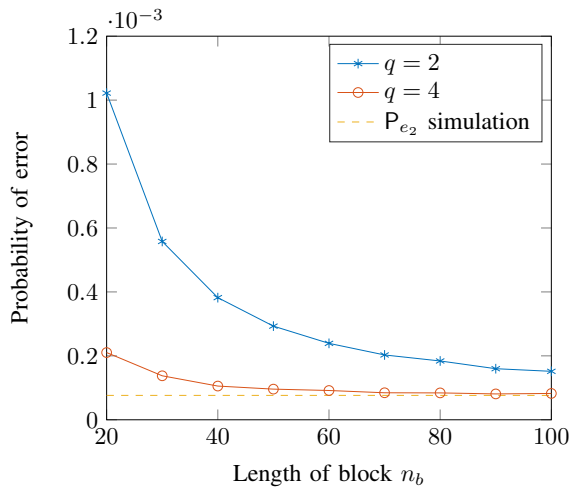


Fig. 1: Probability of error for different values of block length n_b , with $l = 7$, $k = 4$, and $t = 5$.

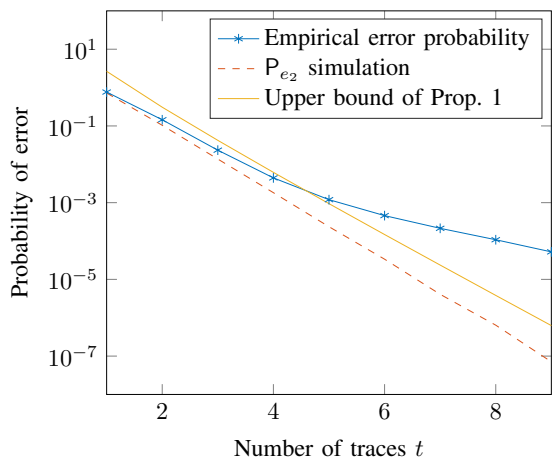


Fig. 2: Probability of error for a binary code for different values of t when $l = 6$, $k = 4$, and $n_b = 30$. The code length $n = 180$, and the rate is $5/6$.

random locations. We note that both the code length and the rate (cf. (6), (7)) increase with n_b .

Figure 1 also shows the empirical phase 2 error (dashed line), which does not depend on either n_b or the alphabet. The difference between the overall and the phase 2 error probabilities can be interpreted as the phase 1 error probability. The phase 1 error caused by wrong matches of VT syndrome decreases with n_b , as explained in Sec. IV-A. Furthermore, we observe that the phase 1 error is smaller (and decreases faster with n_b) for $q = 4$ than for $q = 2$. There are two reasons for this. First, the number of potential VT syndromes for $q > 2$ is qn , in contrast to the binary case where there are $(n + 1)$ VT syndromes. Thus the probability of an accidental match is smaller for the non-binary code. Second, as q increases we expect an accidental match to produce more insertions in the other traces, making it is less likely to be accepted as the correct block. This is because matching of two symbols is less likely in a larger alphabet. The error probability for $q = 4$ approaches P_{e_2} as n_b increases. This is also the case for the binary alphabet. E.g., for $n_b = 1000$ the binary error probability is 7.71×10^{-5} , while P_{e_2} is 7.63×10^{-5} .

Figure 2 shows how the error probability decreases with the number of traces t , for a rate $\frac{5}{6}$ binary code with code parameters held fixed. Each trace has 4 deletions. As shown in Proposition 1, the phase 2 error decays exponentially with t . The overall error probability decays more slowly. Hence for larger values of t , phase 1 error becomes the dominant contribution to the overall error probability. In [14], we report additional simulation results for a varying number of deletions, with the code parameters and the number of traces held fixed.

V. DISCUSSION AND FUTURE WORK

The coding scheme demonstrates that single-deletion correcting codes can be effective in correcting multiple deletions when several traces are available. Having multiple traces helps the decoder in two ways: to identify deletion-free copies of a block, and to avoid mis-synchronization.

The error performance of the coding scheme can be improved by using more complex decoding in either phase. In phase 1, when there are multiple syndrome matches for a block, the decoder could run the decoding procedure for a fixed number of these matches in parallel, using the “fewest insertions” rule to decide which candidates to retain. In phase 2, when unresolvable congestion occurs, the decoder could guess a subset of the bits in the block, use the VT decoder to recover the rest, and choose a guess that produces the correct number of insertions in other traces. Investigating these improved decoders is part of ongoing work. Another idea for future work is to replace the VT code in each block with a lower-rate code capable of correcting more than one deletion.

REFERENCES

- [1] H. M. Kiah, G. J. Puleo, and O. Milenkovic, “Codes for DNA sequence profiles,” *IEEE Trans Inf. Theory*, vol. 62, no. 6, pp. 3125–3146, 2016.
- [2] V. I. Levenshtein, “Efficient reconstruction of sequences,” *IEEE Trans Inf. Theory*, vol. 47, no. 1, pp. 2–22, 2001.
- [3] T. Holenstein, M. Mitzenmacher, R. Panigrahy, and U. Wieder, “Trace reconstruction with constant deletion probability and related results,” in *Proc. ACM-SIAM SODA*, 2008.
- [4] Y. Peres and A. Zhai, “Average-case reconstruction for the deletion channel: subpolynomially many traces suffice,” in *Proc. FOCS*, 2017.
- [5] S. R. Srinivasavaradhan, M. Du, S. Diggavi, and C. Fragouli, “On maximum likelihood reconstruction over multiple deletion channels,” in *Proc. IEEE ISIT*, 2018.
- [6] R. Gabrys and E. Yaakobi, “Sequence reconstruction over the deletion channel,” *IEEE Trans. Inf. Theory*, vol. 64, pp. 2924 – 2931, Apr. 2018.
- [7] F. Sala, R. Gabrys, C. Schoeny, and L. Dolecek, “Exact reconstruction from insertions in synchronization codes,” *IEEE Trans Inf. Theory*, vol. 63, no. 4, pp. 2428–2445, 2017.
- [8] R. Gabrys and O. Milenkovic, “Unique reconstruction of coded strings from multiset substring spectra,” in *Proc. IEEE ISIT*, 2018.
- [9] B. Haeupler and M. Mitzenmacher, “Repeated deletion channels,” in *Proc. IEEE ITW*, 2014.
- [10] R. R. Varshamov and G. M. Tenengolts, “Codes which correct single asymmetric errors,” *Automatica i Telemekhanika*, vol. 26, no. 2, pp. 288–292, 1965.
- [11] G. Tenengolts, “Nonbinary codes, correcting single deletion or insertion,” *IEEE Trans. Inf. Theory*, vol. 30, no. 5, pp. 766–769, 1984.
- [12] K. A. S. Abdel-Ghaffar and H. C. Ferreira, “Systematic encoding of the Varshamov-Tenengolts codes and the Constantin-Rao codes,” *IEEE Trans. Inf. Theory*, vol. 44, pp. 340–345, Jan 1998.
- [13] M. Abroshan, R. Venkataramanan, and A. Guillén i Fàbregas, “Efficient systematic encoding of non-binary VT codes,” in *Proc. IEEE ISIT*, 2018.
- [14] M. Abroshan, R. Venkataramanan, L. Dolecek, and A. Guillén i Fàbregas, “Coding for deletion channels with multiple traces,” 2019. (Online) <http://link.eng.cam.ac.uk/foswiki/pub/Main/RV285/trace.pdf>.