

Adversarial Torn-paper Codes

Daniella Bar-Lev, *Student Member, IEEE*, Sagi Marcovich, *Student Member, IEEE*
Eitan Yaakobi, *Senior Member, IEEE*, and Yonatan Yehezkeally, *Member, IEEE*

Abstract—We study the *adversarial torn-paper channel*. This problem is motivated by applications in DNA data storage where the DNA strands that carry information may break into smaller pieces which are received out of order. Our model extends the previously researched probabilistic setting to the worst-case. We develop code constructions for any parameters of the channel for which non-vanishing asymptotic rate is possible and show our constructions achieve asymptotically optimal rate while allowing for efficient encoding and decoding. Finally, we extend our results to related settings included multi-strand storage, presence of substitution errors, or incomplete coverage.

Index Terms—Sequence reconstruction, DNA sequences, Error correction codes, Worst-case analysis

I. INTRODUCTION

High density and extreme longevity make DNA an appealing medium for data storage, especially for archival purposes [4], [9], [10], [36]. Advances in DNA synthesis and sequencing technologies and recent proofs of concept [5], [9], [13], [16], [17], [25] have ignited active research into the capacity and challenges of data storage in this medium.

An aspect of this medium is that typically only short DNA sequences may be read; information molecules are therefore broken up into pieces and then read out of order, such as in shotgun sequencing [6], [14], [23], [27]. Multiple channel models have recently been suggested and studied based on this property. An assumption of overlap in read substrings

and (near) uniform coverage leads to the problem of string reconstruction from substring composition [3], [6], [15], [22], [23], [29], [31]; on the contrary, assuming no overlap in read substrings leads to the *torn-paper problem* [24], [26], [32], a problem closely related to the shuffling channel [18], [19], [30], [35]. This problem is motivated by DNA-based storage systems, where the information is stored in synthesized strands of DNA molecules. However, during and after synthesis, the DNA strands may break into smaller segments and due to the lack of ordering among the strands in these systems, all broken segments can only be read out of order [32]. Thus, the goal is to successfully retrieve the data from this collection of read segments of the broken DNA strands.

In the *torn-paper channel* [26], [32], also known as the *chop-and-shuffle channel* [24], a long information string is segmented into non-overlapping substrings and their length has some known distribution. The channel outputs an unordered collection of these substrings, preserving their left-to-right orientation. Given the lengths' distribution, the goal is to determine the channel capacity and devise efficient coding techniques. The geometric distribution was first studied in [32], and later in [24] using the Varshamov-Tenengolts (VT) codes [34]. Subsequently, [26] considered almost arbitrary distributions while, additionally, extending the problem by introducing incomplete coverage, i.e., assuming some of the substrings are deleted with some probability.

The torn-paper channel was studied so far only in the probabilistic setting. The goal of this paper is to extend this channel to the worst case, referred to herein as the *adversarial torn-paper channel*. Namely, it is assumed that an information string is adversarially segmented into non-overlapping substrings, where the length of each substring is between L_{\min} and L_{\max} , for some given L_{\min} and L_{\max} . We show that the capacity of this channel is determined by L_{\min} , whereas the capacity of the probabilistic channel was shown to depend on the *average* substring length; nevertheless, we choose this adversarial model here for ease of analysis, and observe that under this setting the average substring length might indeed approach L_{\min} . For further discussion of an average-restricted adversary, see Section V.

We study the noiseless adversarial torn-paper channel for a single information string, as well as multiple strings, which is motivated by DNA sequencing technologies where multiple strings are sequenced simultaneously [8], [21], [28]. We also extend the model to either allow for substitution errors affecting the information string prior to segmentation, or for incomplete coverage due to deletion of several segments after the segmentation. In all cases we investigate the values of L_{\min} and L_{\max} that permit codes with non-vanishing asymptotic rates, and develop constructions of codes with efficient

Manuscript received 26 August 2022; revised 20 February 2023; accepted 29 June 2023. This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 801434). It was also funded by the European Union (ERC, DNASStorage, 865630). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. The work of Daniella Bar-Lev, Sagi Marcovich, and Eitan Yaakobi was supported in part by the U.S.-Israel Binational Science Foundation (BSF) under Grant 2018048. The work of Yonatan Yehezkeally was supported the Alexander von Humboldt Foundation under a Carl Friedrich von Siemens Post-Doctoral Research Fellowship. An earlier version of this paper was presented in part at the 2022 IEEE International Symposium on Information Theory (ISIT) [DOI: 10.1109/ISIT50566.2022.9834766]. (Daniella Bar-Lev, Sagi Marcovich, and Yonatan Yehezkeally contributed equally to this work.) (Corresponding author: Yonatan Yehezkeally.)

Daniella Bar-Lev, Sagi Marcovich, and Eitan Yaakobi are with the Department of Computer Science, Technion—Israel Institute of Technology, Haifa 3200003, Israel (e-mail: daniellalev@cs.technion.ac.il; sagimar@cs.technion.ac.il; yaakobi@cs.technion.ac.il). Yonatan Yehezkeally is with the Institute for Communications Engineering, School of Computation, Information and Technology, Technical University of Munich, 80333 Munich, Germany (e-mail: yonatan.yehezkeally@tum.de).

Copyright (c) 2023 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

encoding and decoding algorithms, asymptotically achieving optimal rates.

The rest of this paper is organized as follows. In Section II, the definitions and notations that will be used throughout the paper are presented, as well as a lower bound on L_{\min} required for the existence of codes for the adversarial torn-paper channel with non-vanishing asymptotic rates. In Section III we first study the application of a known code construction to the adversarial channel, and observe its limitations in that setting; then, we present the basic construction used throughout the paper for the noiseless case of the single-strand adversarial torn-paper channel, and extend it to the multi-strand case. In Section IV we extend our construction to two noisy settings, including substitution errors or incomplete coverage. We conclude with a summary and remarks in Section V.

II. DEFINITIONS AND PRELIMINARIES

Let Σ be a finite alphabet of size q . For convenience of presentation, we assume Σ is equipped with a ring structure, and in particular identify elements $0, 1 \in \Sigma$. For a positive integer n , let $[n]$ denote the set $[n] \triangleq \{0, 1, \dots, n-1\}$. Let Σ^* denote the set of all finite strings over Σ . The length of a string $\mathbf{x} \in \Sigma^*$ is denoted by $|\mathbf{x}|$. We also denote, for $\mathbf{x} = (x_i)_{i \in [n]} \in \Sigma^n$, its *support* $\text{supp}(\mathbf{x}) \triangleq \{i \in [n] : x_i \neq 0\}$, and $\|\mathbf{x}\| \triangleq |\text{supp}(\mathbf{x})|$. For strings $\mathbf{x}, \mathbf{y} \in \Sigma^*$, we denote their concatenation by $\mathbf{x} \circ \mathbf{y}$. We say that \mathbf{v} is a *substring*, or *segment*, of \mathbf{x} if there exist strings \mathbf{u}, \mathbf{w} (perhaps empty) such that $\mathbf{x} = \mathbf{u} \circ \mathbf{v} \circ \mathbf{w}$. If $|\mathbf{v}| = \ell$, we specifically say that \mathbf{v} is an ℓ -*substring* (ℓ -*segment*) of \mathbf{x} . If $|\mathbf{u}| = i$ then it is said that \mathbf{v} is the substring (similarly, ℓ -substring) of \mathbf{x} at *location* i . We say that \mathbf{v} appears *cyclically* in \mathbf{x} , at location i , if $\mathbf{x} = \mathbf{u} \circ \mathbf{w}$ and \mathbf{v} is the substring of $\mathbf{w} \circ \mathbf{u}$ at location $(i - |\mathbf{u}|)$. For example, 010 is the 3-substring of 00101 at location 1, and also its 3-substring at location 3, where the latter is a cyclic appearance. We avoid using the term *index* as it is reserved to elements of presented constructions.

In our setting, information is stored in an unordered collection of strings over Σ ; it might be allowed for the same string to appear with multiplicity in the collection, which is encapsulated in the following formal definition:

$$\mathcal{X}_{n,k} \triangleq \{S = \{\{\mathbf{x}_0, \dots, \mathbf{x}_{k-1}\}\} : \forall i, \mathbf{x}_i \in \Sigma^n\}.$$

Here, $\{\{a, a, b, \dots\}\}$ denotes a multiset; i.e., elements appear with multiplicity (but no order). Note that $|\mathcal{X}_{n,k}| = \binom{k+q^n-1}{k}$. It is assumed that a message $S \in \mathcal{X}_{n,k}$ is read by segmenting all elements of S into non-overlapping substrings of lengths between some fixed values L_{\min} and L_{\max} , and all segments are received, possibly with multiplicity, without order or information on which element they originated from. More formally, a *segmentation* of the string \mathbf{x} is a multiset $\{\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{m-1}\}\}$, where \mathbf{x} can be presented as $\mathbf{x} = \mathbf{u}_0 \circ \mathbf{u}_1 \circ \dots \circ \mathbf{u}_{m-1}$. In case $L_{\min} \leq |\mathbf{u}_i| \leq L_{\max}$ for $0 \leq i < m-1$ and $|\mathbf{u}_{m-1}| \leq L_{\max}$, then the segmentation is called an (L_{\min}, L_{\max}) -*segmentation*. The set of all (L_{\min}, L_{\max}) -segmentations of \mathbf{x} is denoted by $\mathcal{T}_{L_{\min}}^{L_{\max}}(\mathbf{x})$ and

is referred as the (L_{\min}, L_{\max}) -*segmentation spectrum* of \mathbf{x} . For example,

$$\mathcal{T}_2^3(00101) = \left\{ \{\{001, 01\}\}, \{\{00, 101\}\}, \{\{00, 10, 1\}\} \right\}.$$

These definitions are naturally extended for a multiset $S \in \mathcal{X}_{n,k}$, so a *segmentation* of S is a union (as a multiset) of segmentations of all the strings in S (and the same holds for an (L_{\min}, L_{\max}) -segmentation), and $\mathcal{T}_{L_{\min}}^{L_{\max}}(S)$, the (L_{\min}, L_{\max}) -*segmentation spectrum* of S , is the set of all (L_{\min}, L_{\max}) -segmentations of S .

Note that our channel model only restricts the length of the last segment to be at most L_{\max} . Such a relaxation is motivated in applications where segmentation of the strings occurs sequentially, so that it might happen that the last segment is shorter than L_{\min} , but not larger than L_{\max} .

A code $\mathcal{C} \subseteq \mathcal{X}_{n,k}$ is said to be an (L_{\min}, L_{\max}) -*multistrand torn-paper code* if for all $S, S' \in \mathcal{C}$, $S \neq S'$, it holds that all possible (L_{\min}, L_{\max}) -segmentations of S, S' are distinct. That is, $\mathcal{T}_{L_{\min}}^{L_{\max}}(S) \cap \mathcal{T}_{L_{\min}}^{L_{\max}}(S') = \emptyset$. For $k = 1$, we simply refer to (L_{\min}, L_{\max}) -*single strand torn-paper codes*.

In case $L_{\min} = L_{\max} = \ell$, then for convenience, we let $\mathcal{T}_\ell(\mathbf{x}) \triangleq \mathcal{T}_\ell^\ell(\mathbf{x})$ and $\mathcal{T}_\ell(S) \triangleq \mathcal{T}_\ell^\ell(S)$ and note that in this case $|\mathcal{T}_\ell(\mathbf{x})| = |\mathcal{T}_\ell(S)| = 1$. For example, if $S = \{\{01010, 00101, 11101\}\}$ (which may be thought of as a multiset), then

$$\mathcal{T}_2(S) = \left\{ \{\{01, 01, 0, 00, 10, 1, 11, 10, 1\}\} \right\}.$$

Note that $\mathcal{T}_\ell(S)$ is only one possible channel output given input S . Nevertheless, $\mathcal{T}_{L_{\min}}(S) \subseteq \mathcal{T}_{L_{\min}}^{L_{\max}}(S)$ for all S and $L_{\min} \leq L_{\max}$, hence every (L_{\min}, L_{\max}) -multistrand torn-paper code $\mathcal{C} \subseteq \mathcal{X}_{n,k}$ satisfies

$$|\mathcal{C}| \leq |\{\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}|. \quad (1)$$

For all $\mathcal{C} \subseteq \mathcal{X}_{n,k}$ we denote the *rate*, *redundancy* of \mathcal{C} by $R(\mathcal{C}) \triangleq \frac{\log|\mathcal{C}|}{\log|\mathcal{X}_{n,k}|}$, $\text{red}(\mathcal{C}) \triangleq \log|\mathcal{X}_{n,k}| - \log|\mathcal{C}|$, respectively. Throughout the paper, we use the base- q logarithms.

For two non-negative functions f, g of a common variable n , denoting $L \triangleq \limsup_{n \rightarrow \infty} \frac{f(n)}{g(n)}$ (in the wide sense, i.e., $L = \infty$ if $\frac{f(n)}{g(n)}$ is unbounded) we say that $f = o_n(g)$ if $L = 0$, $f = \Omega_n(g)$ if $L > 0$, $f = O_n(g)$ if $L < \infty$, and $f = \omega_n(g)$ if $L = \infty$. If f is not positive, we say $f = O_n(g)$ ($f = o_n(g)$) if $|f| = O_n(g)$ (respectively, $|f| = o_n(g)$). We say that $f = \Theta_n(g)$ if $f = \Omega_n(g)$ and $f = O_n(g)$. If clear from context, we omit the subscript from aforementioned notations.

We conclude this section by observing a lower bound on the required segment length L_{\min} for multi-strand torn-paper codes to achieve non-vanishing rates, and in particular rates approaching one.

Lemma 1 *If $\log(k) = o(n)$ and $L_{\min} = a \log(nk) + O_{nk}(1)$ for some $a \geq 1$, then*

$$\begin{aligned} & \log|\mathcal{X}_{n,k}| - \log|\{\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}| \\ & \geq nk \left(\frac{1}{a} - a \frac{\log(k)}{n} - O\left(\frac{\log \log(nk)}{\log(nk)}\right) \right). \end{aligned}$$

Proof: First, note that

$$|\mathcal{X}_{n,k}| = \binom{k+q^n-1}{k} \geq \frac{q^{nk}}{k!} \geq \frac{q^{nk}}{k^k},$$

and hence $\log|\mathcal{X}_{n,k}| \geq (n - \log(k))k$. Next, since $\{|\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}|$ is monotonically non-decreasing in n , we have that

$$\begin{aligned} |\{\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}| &\leq \binom{k\lceil n/L_{\min} \rceil + q^{L_{\min}} - 1}{q^{L_{\min}} - 1} \\ &\leq \binom{k\lceil n/L_{\min} \rceil + q^{L_{\min}}}{q^{L_{\min}}}. \end{aligned}$$

Now, for $v \geq u \geq 0$ we observe

$$\begin{aligned} \log \binom{u+v}{u} &\leq \log \frac{1}{u!} (u+v)^u \leq u \log \left(e \left(1 + \frac{v}{u} \right) \right) \\ &\leq u \left(\left(1 + \frac{u}{v} \right) \log(e) + \log \left(\frac{v}{u} \right) \right) \\ &\leq u \left(2 \log(e) + \log \left(\frac{v}{u} \right) \right), \end{aligned}$$

where we used $\log(1+x) \leq \frac{\log(e)}{x} + \log(x)$. Setting $u \triangleq k\lceil n/L_{\min} \rceil \leq \frac{nk}{L_{\min}} + k$ and $v \triangleq q^{L_{\min}} = \Theta((nk)^a)$, we have $\frac{v}{u} = \Theta((nk)^{a-1} L_{\min}) = \Theta((nk)^{a-1} \log(nk))$, and therefore $\log \left(\frac{v}{u} \right) = (a-1) \log(nk) + \log \log(nk) + O(1)$. We then conclude

$$\begin{aligned} &\log|\{\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}| \\ &\leq \left(\frac{nk}{L_{\min}} + k \right) \left((a-1) \log(nk) + \log \log(nk) + O(1) \right) \\ &= (a-1) \frac{nk \log(nk)}{L_{\min}} + k(a-1) \log(nk) + \left(\frac{nk}{L_{\min}} + k \right) (\log \log(nk) + O(1)) \\ &= (a-1) \frac{nk \log(nk)}{L_{\min}} + k(a-1) \log(nk) + O \left(\frac{nk \log \log(nk)}{\log(nk)} \right) \\ &= nk \left((a-1) \frac{\log(nk)}{L_{\min}} + (a-1) \frac{\log(k)}{n} + O \left(\frac{\log(n)}{n} \right) + O \left(\frac{\log \log(nk)}{\log(nk)} \right) \right) \\ &= nk \left(\frac{a-1}{a + O(1/\log(nk))} + (a-1) \frac{\log(k)}{n} + O \left(\frac{\log \log(nk)}{\log(nk)} \right) \right) \\ &= nk \left(\frac{a-1}{a} + (a-1) \frac{\log(k)}{n} + O \left(\frac{\log \log(nk)}{\log(nk)} \right) \right), \end{aligned}$$

which verifies the lemma's statement. \blacksquare

We note that throughout this paper, we perform redundancy analysis to the second-most-significant term, and retain the order or magnitude for the remainder; since proofs demonstrate that this asymptotic notation does not in fact hide significant coefficients, we believe this representation is faithful for the purpose of finite-length analysis, as well.

The implications of Lemma 1 are more clearly stated in the next corollary.

Corollary 2 *Let \mathcal{C} be any (L_{\min}, L_{\max}) -multistrand torn-paper code. Assuming $\log(k) = o(n)$, if $L_{\min} = (a + o_{nk}(1)) \log(nk)$, for some $a \geq 1$, then $R(\mathcal{C}) \leq 1 - \frac{1}{a} + o_{nk}(1)$.*

Proof: From Eq. (1) and Lemma 1 we have

$$\begin{aligned} R(\mathcal{C}) &\leq \frac{\log|\{\mathcal{T}_{L_{\min}}(S) : S \in \mathcal{X}_{n,k}\}|}{\log|\mathcal{X}_{n,k}|} \\ &\leq 1 - \frac{nk}{\log|\mathcal{X}_{n,k}|} \left(\frac{1}{a} - a \frac{\log(k)}{n} - O \left(\frac{\log \log(nk)}{\log(nk)} \right) \right), \end{aligned}$$

which, together with $\log|\mathcal{X}_{n,k}| \leq nk$, justifies the claim. \blacksquare

III. CONSTRUCTIONS OF TORN-PAPER CODES

In this section we study constructions of torn-paper codes, in context of the bound of Corollary 2.

A. Related works: pilot-based construction

An explicit and efficient coding scheme was presented in [32] for the probabilistic torn-paper channel. Therein, it was argued that an indexing approach to coding is challenging due to the a priori unknown locations of segmentation by the channel, hence this construction relied on interleaving a *pilot* (or *phase-detection sequence*). We describe this scheme below to study its performance in the adversarial channel.

Construction P [32, Sec. VII] Fix an integer $m > 1$. Let n be a multiple of m , to be determined later, and s an integer satisfying $s \geq \log(n/m)$. Let $\mathbf{p} \in \Sigma^{n/m}$ be any (n/m) -segment of a de Bruijn sequence [11] of order s , which we refer to as the *pilot*.

For $\mathbf{x}, \mathbf{y} \in \Sigma^{n/m}$, denote $\mathbf{x} \perp^s \mathbf{y}$ if \mathbf{x}, \mathbf{y} have no common s -segment, i.e., if for all $i, j \in [n/m - s + 1]$ it holds that $\mathbf{x}^{(i)} \neq \mathbf{y}^{(j)}$, where $\mathbf{x}^{(i)}$ ($\mathbf{y}^{(j)}$) is the s -segment of \mathbf{x} (respectively, \mathbf{y}) at location i (respectively, j). Then, we denote $\mathcal{O}_{\mathbf{p}} \triangleq \{\mathbf{c} \in \Sigma^{n/m} : \mathbf{c} \perp^s \mathbf{p}\}$.

For any code $\mathcal{C} \subseteq \Sigma^{n/m}$, we construct a code $\mathcal{C}_{\text{pilot}} \subseteq \Sigma^n$ as follows: for every choice of $m-1$ elements $(\mathbf{c}_j)_{j \in [m-1]} \subseteq \mathcal{C} \cap \mathcal{O}_{\mathbf{p}}$ (allowing for repetition), we interleave a single symbol from each $\mathbf{p}, \mathbf{c}_0, \mathbf{c}_1, \dots, \mathbf{c}_{m-2}$, in order, to construct a codeword $\mathbf{c} \in \mathcal{C}_{\text{pilot}}$. \square

Example 3 Let $q = 2, m = 2, n = 12, s = 3$. We choose 00010111 as the binary de Bruijn sequence of order s , and let $\mathbf{p} \triangleq 000101$ be its (n/m) -prefix. Then,

$$\begin{aligned} \mathcal{O}_{\mathbf{p}} = \{ &011100, 011110, \\ &011111, 111100, \\ &111110, 111111 \}. \end{aligned}$$

Letting $\mathcal{C} \triangleq \Sigma^5$, and for any choice of $m-1 = 1$ element of $\mathcal{C} \cap \mathcal{O}_{\mathbf{p}} = \mathcal{O}_{\mathbf{p}}$, we interleave \mathbf{p} with that element to derive the code

$$\begin{aligned} \mathcal{C}_{\text{pilot}} = \{ &000101110010, 000101110110, \\ &000101110111, 010101110010, \\ &010101110110, 010101110111 \}. \end{aligned}$$

Lemma 4 [32, Sec. VII-B] For all $s \geq \log(n/m)$ it holds that $\mathcal{C}_{\text{pilot}}$ is an (ms, L_{max}) -single strand torn-paper code, for any $L_{\text{max}} \geq ms$.

Proof: We replicate the proof for completeness. Observe that every (ms) -segment \mathbf{u} of $\mathbf{c} \in \mathcal{C}_{\text{pilot}}$ contains s consecutive symbols from each $\mathbf{p}, \mathbf{c}_0, \mathbf{c}_1, \dots, \mathbf{c}_{m-2}$; since $\mathbf{c}_j \perp^s \mathbf{p}$ for every $j \in [m-1]$, the s -segment of \mathbf{p} can be uniquely identified. Since \mathbf{p} is a segment of a de Bruijn sequence of order s , the location in \mathbf{p} of the observed segment can be deduced, and hence the location of \mathbf{u} in \mathbf{c} can readily be obtained. ■

Example 5 Continuing Example 3, assume $010101110110 \in \mathcal{C}_{\text{pilot}}$ is passed through an adversarial torn-paper channel with $L_{\text{min}} = ms = 6$ and, say, $L_{\text{max}} = 8$. The received segments are

$$010101, 110110.$$

Taking the first segment, we decompose the two interleaved strings

$$\bar{\mathbf{c}}_0 = 000, \bar{\mathbf{c}}_1 = 111;$$

we identify $\bar{\mathbf{c}}_0$ as the s -substring of \mathbf{p} at location 0, implying that $\bar{\mathbf{c}}_1$ is the substring of \mathbf{c}_0 at location 0. Similarly, we decompose the second segment into

$$\tilde{\mathbf{c}}_0 = 101, \tilde{\mathbf{c}}_1 = 110;$$

since 101 is the s -substring of \mathbf{p} at location 3, we also have that $\tilde{\mathbf{c}}_1$ is the substring of \mathbf{c}_0 at location 3, i.e., $\mathbf{c}_0 = 111110 \in \mathcal{O}_{\mathbf{p}}$, confirming $010101110110 \in \mathcal{C}_{\text{pilot}}$ was the transmitted sequence.

For the probabilistic channel studied in [32], \mathcal{C} in Construction P was chosen to be an error-correcting code. Note from the proof of Lemma 4 that in our chosen adversarial setting, this is redundant; that element of the construction is preserved in our presentation to support the discussion in Section V regarding alternate models.

Next, we turn to find the achievable rates of Construction P.

Corollary 6 $R(\mathcal{C}_{\text{pilot}}) = (1 - \frac{1}{m})R(\mathcal{C} \cap \mathcal{O}_{\mathbf{p}})$.

Proof: Observe that $|\mathcal{C}_{\text{pilot}}| = |\mathcal{C} \cap \mathcal{O}_{\mathbf{p}}|^{m-1} = q^{n(1-\frac{1}{m})\log(|\mathcal{C} \cap \mathcal{O}_{\mathbf{p}}|)/\frac{n}{m}} = q^{n(1-\frac{1}{m})R(\mathcal{C} \cap \mathcal{O}_{\mathbf{p}})}$. ■

The following lemma was implied by [32, Sec. VII-A].

Lemma 7 For all $\mathcal{C} \subseteq \Sigma^{n/m}$ there exists $\mathbf{z} \in \Sigma^{n/m}$ such that

$$R((\mathbf{z} + \mathcal{C}) \cap \mathcal{O}_{\mathbf{p}}) \geq R(\mathcal{C}) - (1 - R(\mathcal{O}_{\mathbf{p}})),$$

where $\mathbf{z} + \mathcal{C} \triangleq \{\mathbf{z} + \mathbf{c} : \mathbf{c} \in \mathcal{C}\}$.

Proof: Observe that

$$\begin{aligned} \sum_{\mathbf{z} \in \Sigma^{n/m}} |(\mathbf{z} + \mathcal{C}) \cap \mathcal{O}_{\mathbf{p}}| &= \sum_{\mathbf{z} \in \Sigma^{n/m}} \sum_{\substack{\mathbf{c}_1 \in \mathcal{C} \\ \mathbf{c}_2 \in \mathcal{O}_{\mathbf{p}}}} \mathbb{1}_{\mathbf{z} + \mathbf{c}_1 = \mathbf{c}_2} \\ &= \sum_{\substack{\mathbf{c}_1 \in \mathcal{C} \\ \mathbf{c}_2 \in \mathcal{O}_{\mathbf{p}}}} \sum_{\mathbf{z} \in \Sigma^{n/m}} \mathbb{1}_{\mathbf{z} = \mathbf{c}_2 - \mathbf{c}_1} \\ &= \sum_{\substack{\mathbf{c} \in \mathcal{C} \\ \mathbf{o} \in \mathcal{O}_{\mathbf{p}}}} 1 = |\mathcal{C}| \cdot |\mathcal{O}_{\mathbf{p}}|. \end{aligned}$$

It follows from the pigeonhole principle that there exists $\mathbf{z} \in \Sigma^{n/m}$ such that $|(\mathbf{z} + \mathcal{C}) \cap \mathcal{O}_{\mathbf{p}}| \geq q^{-n/m} |\mathcal{C}| \cdot |\mathcal{O}_{\mathbf{p}}|$, which concludes the proof. ■

In the rest of the section, it remains to analyze what values of s assure that $1 - R(\mathcal{O}_{\mathbf{p}}) = o_n(1)$; we also discuss the implications of these available choices.

Lemma 8 [32, Sec. VII-A] If $s \triangleq \lceil (2 + \delta) \log(n/m) \rceil$ for some $\delta > 0$, then, using $\mathcal{C} \triangleq \Sigma^{n/m}$ in Construction P,

$$\begin{aligned} R(\mathcal{C}_{\text{pilot}}) &\geq 1 - \frac{1}{m} - \frac{m-1}{n} \cdot \frac{1}{(n/m)^\delta - 1} \\ &= 1 - \frac{1}{m} - o_n(1). \end{aligned}$$

Proof: Again, we replicate the proof here. Denote for a uniformly chosen $\mathbf{c} \in \Sigma^{n/m}$ the event $A_{i,j}$ that $\mathbf{c}^{(i)} = \mathbf{p}^{(j)}$. Clearly $\Pr(A_{i,j}) = q^{-s}$; using the union bound, $\Pr(\mathbf{c} \perp^s \mathbf{p}) \geq 1 - (n/m)^2 q^{-s} \geq 1 - (n/m)^{-\delta}$, i.e.,

$$|\mathcal{O}_{\mathbf{p}}| \geq q^{n/m} (1 - (n/m)^{-\delta}).$$

It follows from Corollary 6 that

$$\begin{aligned} R(\mathcal{C}_{\text{pilot}}) &= \left(1 - \frac{1}{m}\right) R(\mathcal{O}_{\mathbf{p}}) \\ &\geq \left(1 - \frac{1}{m}\right) \left(1 + \frac{m}{n} \log(1 - (n/m)^{-\delta})\right) \\ &\geq 1 - \frac{1}{m} - \frac{(m-1)(n/m)^{-\delta}}{n(1 - (n/m)^{-\delta})}. \end{aligned}$$

Unfortunately, Lemma 8 doesn't match the upper bound of Corollary 2; asymptotically, it produces rate $1 - \frac{2+\delta}{a}$, where $a \triangleq \frac{ms}{\log(n)}$. Further, the construction may only be applied when a is (approximately) an even integer ≥ 4 . The former can be remedied by replacing the union bound in the analysis of [32, Sec. VII-A] with the Lovász local lemma [33] (similarly to techniques used independently in [38] and [12]), as follows.

Lemma 9 Let $s \triangleq \lceil \log(n/m) + \log \log(n/m) + \log(3e) \rceil$. Then, using $\mathcal{C} \triangleq \Sigma^{n/m}$ in Construction P,

$$\begin{aligned} R(\mathcal{C}_{\text{pilot}}) &\geq \left(1 - \frac{1}{m}\right) \cdot \left(1 - \frac{\log(e)}{2 \log(n/m)}\right) \\ &= 1 - \frac{1}{m} - O\left(\frac{1}{\log(n)}\right). \end{aligned}$$

Proof: Denote for a uniformly chosen $\mathbf{c} \in \Sigma^{n/m}$ the event $A_{i,j}$ that $\mathbf{c}^{(i)} = \mathbf{p}^{(j)}$. Clearly $p \triangleq \Pr(A_{i,j}) = q^{-s}$, and $A_{i,j}$

is jointly independent of $\{A_{i',j'} : |i - i'| \geq s\}$, i.e., all except $(n/m - s)(2s - 1) - 1 \leq 2sn/m - 1$ distinct events.

For sufficiently large n , observe that

$$\begin{aligned} sq^{-s} &\leq \frac{\log(n/m) + \log \log(n/m) + \log(3e)}{(n/m) \log(n/m) 3e} \\ &= \frac{m}{2en} \cdot \frac{2}{3} \left(1 + \frac{\log \log(n/m) + \log(3e)}{\log(n/m)} \right) < \frac{m}{2en}, \end{aligned}$$

where the first inequality is justified by $(s+r)q^{-(s+r)} \leq sq^{-s}$ for $r \geq 0$ and $s \geq \log(e)$. Rearranging, we have $m \geq 2epsn$. Therefore, letting $x \triangleq \frac{ep}{1+ep}$ (hence, $\frac{x}{1-x} = ep$), and recalling for all $x \in (0, 1)$ that $1 - x \geq \exp(\frac{-x}{1-x})$, we have

$$\begin{aligned} x(1-x)^{2sn/m-1} &= ep(1-x)^{2sn/m} \\ &\geq p \exp\left(1 - 2epsn/m\right) > p. \end{aligned}$$

It therefore follows from the local lemma that

$$\begin{aligned} \Pr(\mathbf{c} \perp^s \mathbf{p}) &\geq (1-x)^{(n/m)^2} \geq \exp(-ep(n/m)^2) \\ &= e^{-e(n/m)^2 q^{-s}} \geq e^{-n/2sm}, \end{aligned}$$

where again we used the fact that $m \geq 2epsn$. That is, $|\mathcal{O}_p| \geq q^{n/m} e^{-n/2sm} = (qe^{-1/2s})^{n/m}$, and

$$R(\mathcal{O}_p) \geq 1 - \frac{\log(e)}{2s}.$$

Hence, Corollary 6 concludes the proof. \blacksquare

Based on Lemma 9, Construction P achieves $1 - \frac{1}{a} - o_n(1)$ rate, where $a \triangleq \frac{ms}{\log(n)}$, asymptotically matching the bound of Corollary 2. It also expands the values of a for which the construction may be applied; however, unfortunately a is still restricted to be (approximately) an integer ≥ 2 . Moreover, encoding $\mathcal{C}_{\text{pilot}}(n)$ involves a choice of \mathbf{p} , and the authors are not aware of a straightforward way to make this choice while optimizing $R(\mathcal{O}_p)$; it further requires encoding into (potentially, a sub-code of) \mathcal{O}_p , which is also, to the best of our knowledge, not readily done in an efficient manner. To bridge that gap, we present in the next section a construction based on an indexing approach, which can be applied for any $a > 1$, asymptotically matching Corollary 2 for all choices.

B. Index-based construction

In this section, an index-based construction of single-strand torn-paper codes is presented and is then extended for multiple strands.

It is assumed from here on out that $L_{\min} = \lceil a \log(n) \rceil$, for some $a > 1$ which is fixed throughout this section. We propose the following construction of length- n (L_{\min}, L_{\max})-single strand torn-paper codes. The construction is based on the following components.

Definition 10 For an integer I , let $(\mathbf{c}_i)_{i \in [q^I]}$, $\mathbf{c}_i \in \Sigma^I$ be codewords of a q -ary Gray code, in order. Denote by \mathbf{c}'_i the concatenation of \mathbf{c}_i with a single parity symbol (i.e., the sum of the entries in \mathbf{c}'_i is zero). Further, denote by \mathbf{c}''_i the result of inserting '1's into \mathbf{c}'_i at every location divisible by $f(n)$ (since the locations of substrings start with 0, the first bit of




Figure 1. Index generation (best viewed in online colored version)

\mathbf{c}'_i is always '1'). The process is illustrated in Figure 1. Note that $\alpha \triangleq |\mathbf{c}''_i| = \left\lceil \frac{f(n)}{f(n)-1} (I+1) \right\rceil$ for all $i \in [q^I]$. We refer to \mathbf{c}_i (or simply i) as an index in the construction and to \mathbf{c}''_i as an encoded index.

This index structure is motivated by the property indicated in the following lemma.

Lemma 11 Let \mathbf{c} be an α -substring of $\mathbf{c}''_i \circ \mathbf{c}''_{i+1}$, for some $i \in [q^I - 1]$. Then i can uniquely be recovered from \mathbf{c} .

Proof: Since \mathbf{c}''_i and \mathbf{c}''_{i+1} differ only at the parity symbol and one additional coordinate (which corresponds to the only position where \mathbf{c}_i and \mathbf{c}_{i+1} differ), \mathbf{c} is either \mathbf{c}''_i or a copy of \mathbf{c}''_{i+1} with an erroneous parity symbol. To obtain i it suffices to distinguish these two cases, which may be done by calculating the parity symbol of \mathbf{c}'' ; If the parity symbol is correct then i equals to the decoding of \mathbf{c} (with the Gray-code decoder), and otherwise i equals to the decoding of \mathbf{c} minus one. \blacksquare

Definition 12 Let f, N be integers. The Run-length limited (RLL) encoder E_N^{RLL} receives strings of length $m(N)$ and returns strings of length N that do not contain zero runs of length f . Constructions of such encoders can be taken from [20] or [37, Lem. 4].

Construction A The main idea of the construction is that every codeword should constitute a concatenation of length- L_{\min} segments with the following structure: an index, followed by a marker, then encoded data. Let $f(n)$ be any integer-valued function satisfying $f(n) = \omega(1)$ and $f(n) = o(\log(n))$ (see Theorem 18 for a choice optimizing the redundancy of this construction). Further assume $n \geq L_{\min} \geq \alpha + f(n) + 2$. Let $I \triangleq \lceil \log(n/L_{\min}) \rceil$, $K \triangleq \lfloor n/L_{\min} \rfloor - 1$ and $N \triangleq L_{\min} - \alpha - f(n) - 2$. The constructed (L_{\min}, L_{\max})-single strand torn-paper code, denoted by $\mathcal{C}_A(n)$, is defined by the encoder $\text{Enc}_A : \Sigma^{Km(N)} \rightarrow \Sigma^n$ in Algorithm 1, and illustrated in Figure 2. \square

In the rest of the paper, we call the strings \mathbf{x}_i (respectively, \mathbf{y}_i) in the constructions an *information block (encoded block)*; the strings $10^{f(n)}1$ are called *markers*; finally, a string \mathbf{z}_i will

Algorithm 1: Encoder for Construction A**Input:** $\mathbf{x} = (x_0, x_1, \dots, x_{Km(N)-1}) \in \Sigma^{Km(N)}$ **Output:** $\text{Enc}_A(\mathbf{x})$ **for** $i \leftarrow 0$ **to** $K - 1$ **do**

$$\mathbf{x}_i \leftarrow (x_{im(N)}, x_{im(N)+1}, \dots, x_{(i+1)m(N)-1})$$


// $|\mathbf{x}_i| = m(N)$

$$\mathbf{y}_i \leftarrow E_N^{\text{RLL}}(\mathbf{x}_i) \text{ // } \mathbf{y}_i \text{ contains no zero runs}$$

of length $f(n)$

$$\mathbf{z}_i \leftarrow \mathbf{c}_i'' \circ 10^{f(n)}1 \circ \mathbf{y}_i \text{ // } |\mathbf{z}_i| = L_{\min}$$
end

$$\mathbf{z}_K \leftarrow \mathbf{c}_K'' \circ 10^{f(n)}10^N \text{ // } |\mathbf{z}_K| = L_{\min}$$

$$\mathbf{z} \leftarrow \mathbf{z}_0 \circ \mathbf{z}_1 \circ \dots \circ \mathbf{z}_K \circ 0^{n \bmod L_{\min}} \text{ // } |\mathbf{z}| = n$$
return \mathbf{z} **Figure 2.** Illustration of Algorithm 1 (best viewed in online colored version)

simply be referred to as a segment of \mathbf{z} . Note that the last segment \mathbf{z}_K of \mathbf{z} deliberately does not contain data, to account for the possibility that a part of $\mathbf{z}_K \circ 0^{n \bmod L_{\min}}$ might be partitioned by an adversarial channel in such a way that it does not contain, at its suffix, a prefix of an index. We observe that once the encoded blocks \mathbf{y}_i 's are obtained, encoding (including the generation of the Gray code) then requires a number of operations linear in n . By [20], [37], encoding each \mathbf{x}_i into \mathbf{y}_i may also be achieved with a linear number of operations. Hence, the complexity of Construction A is linear with n .

Example 13 We demonstrate the operation of Enc_A . Let $q = 2, n = 45, L_{\min} = 14, f(n) = 2$. For index generation, we utilize the binary Gray code (00, 01, 11, 10), whose encoded indices are (in order)

$$(101010, 101111, 111110, 111011)$$

(observe $\alpha = 6$). Let $N = L_{\min} - f(n) - 2 - \alpha = 4$, and observe an encoder E_N^{RLL} exists with $m(N) = 3$, defined by the lexicographic ordering of the $f(n)$ -run-length-limited sequences of length N

$$\{0101, 0110, 0111, 1010, \\ 1011, 1101, 1110, 1111\}.$$

Noting that $K = 2$, we demonstrate, e.g., the encoding of the information sequence 001110. Observe, $\mathbf{x}_0 = 001, \mathbf{x}_1 = 110$, hence $\mathbf{y}_0 = 0110, \mathbf{y}_1 = 1110$. We then have

$$\mathbf{z}_0 = 101010 \ 1001 \ 0110$$

$$\mathbf{z}_1 = 101111 \ 1001 \ 1110$$

$$\mathbf{z}_2 = 111110 \ 1001 \ 0000,$$

and $\mathbf{z} = \mathbf{z}_0 \circ \mathbf{z}_1 \circ \mathbf{z}_2 \circ 000$.

Next, it is shown that the constructed code $\mathcal{C}_A(n)$ is an (L_{\min}, L_{\max}) -single strand torn-paper code.

Theorem 14 For all $L_{\max} \geq L_{\min}$, $\mathcal{C}_A(n)$ is an (L_{\min}, L_{\max}) -single strand torn-paper code with a linear-runtime decoder.

The proof of Theorem 14 is carried by presenting an explicit decoder to $\mathcal{C}_A(n)$ as follows. Let $\mathbf{z} \in \mathcal{C}_A(n)$ and let $\mathbf{z} = \mathbf{u}_0 \circ \mathbf{u}_1 \circ \dots \circ \mathbf{u}_{s-1}$ so that $\{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_{s-1}\}$ is an (L_{\min}, L_{\max}) -segmentation of \mathbf{z} . The main task of the decoding algorithm is to successfully retrieve the location within \mathbf{z} of each of the s segments of the (L_{\min}, L_{\max}) -segmentation. For every segment \mathbf{u}_j , $j \in [s]$, the decoder first finds the location i such that the first (maybe partial) occurrence of an encoded index in the segment \mathbf{u}_j is of \mathbf{c}_i'' (see below for a proof that this is possible). Given i and the location of \mathbf{c}_i'' in \mathbf{u}_j , the location of the segment \mathbf{u}_j within \mathbf{z} can be calculated. Then, according to the location in \mathbf{z} for each segment in the (L_{\min}, L_{\max}) -segmentation, one can simply concatenate the segments in the correct order to obtain the codeword \mathbf{z} . Finally, by removing the markers and the encoded indices and applying the RLL decoder for each of the strings \mathbf{y}_i 's, the information string \mathbf{x} is retrieved.

Consider the case where a segment \mathbf{u} is a proper substring of the suffix of \mathbf{z} of length $(n \bmod L_{\min}) + N + f(n)$, i.e., $\mathbf{z}_K 0^{n \bmod L_{\min}}$ (note that this does not imply that \mathbf{u} is itself a suffix of \mathbf{z}). Then, \mathbf{u} does not intersect \mathbf{y}_i for any $i \in [K]$, and may safely be discarded. We see next that these cases may be identified efficiently.

Lemma 15 Let $\mathbf{z} \in \mathcal{C}_A(n)$ and let \mathbf{u} be a proper substring of $\mathbf{z}_K 0^{n \bmod L_{\min}}$. If n is sufficiently large (specifically, if $(a - 1)\lceil \log(n) \rceil > 2f(n) + 1$), then this fact can efficiently be identified.

Proof: Observe that either $|\mathbf{u}| < L_{\min}$ or \mathbf{u} contains a suffix of '0's of length at least

$$L_{\min} - \alpha - f(n) - 1 \geq (a - 1)\lceil \log(n) \rceil - f(n) - 1,$$

i.e., longer than $f(n)$, which can easily be identified. ■

By Lemma 15, it is sufficient to retrieve the location of any segment which is not a substring of the suffix of length $(n \bmod L_{\min}) + N + f(n)$ of \mathbf{z} . For any such \mathbf{u} , the calculation of the index i such that \mathbf{c}_i'' is the first (perhaps partial) occurrence of an encoded index within \mathbf{u} , is given in Algorithm 2.

Any L -segment \mathbf{u} of $\mathbf{z} \in \mathcal{C}_A(n)$, such that $L \geq L_{\min}$, contains at least part of one of the encoded indices \mathbf{c}_i'' . If \mathbf{c}_i'' is the first encoded index to intersect \mathbf{u} , we denote by $\text{Ind}(\mathbf{u}) \triangleq i$

Algorithm 2: Index retrieval from a segment

Input: An L -segment \mathbf{u} of a codeword of $\mathcal{C}_A(n)$,
where $L \geq L_{\min}$.

Output: The index of \mathbf{u} within \mathbf{z} , $\text{Ind}(\mathbf{u})$
 $\mathbf{u}' \leftarrow$ the L_{\min} -length prefix of \mathbf{u}
 $j \leftarrow$ the starting index of the unique occurrence of
 $10^{f(n)}1$ within \mathbf{u}' ; if none exists, of the cyclic
occurrence
 $\mathbf{c}'' \leftarrow$ the (cyclic) α -substring of \mathbf{u} strictly preceding j
 $\mathbf{c}' \leftarrow$ the non-padded subsequence of \mathbf{c}''
 $\mathbf{c} \leftarrow$ the I -prefix of \mathbf{c}'
 $\text{Ind} \leftarrow$ the index of \mathbf{c} in the Gray code
if the last symbol of \mathbf{c}' is not the parity of \mathbf{c} **then**
| $\text{Ind} \leftarrow \text{Ind} - 1$
end
return Ind

the index of \mathbf{u} . Note that this index does not depend on the information that was encoded in the construction, but rather, only on the location of \mathbf{u} in \mathbf{z} . Algorithm 2 ensures that it is possible to determine the index of every L -segment \mathbf{u} of \mathbf{z} , where $L \geq L_{\min}$.

The correctness of Algorithm 2 follows from the next lemma.

Lemma 16 Let $\mathbf{z} \in \mathcal{C}_A(n)$, $L \geq L_{\min}$, and let \mathbf{u} be an L -segment of \mathbf{z} which is not a substring of the suffix of length $(n \bmod L_{\min}) + N + f(n)$ of \mathbf{z} . Then, Algorithm 2 successfully returns the index $\text{Ind}(\mathbf{u})$ of \mathbf{u} .

Proof: Let \mathbf{u} be a substring of \mathbf{z} and w.l.o.g. assume that $|\mathbf{u}| = L_{\min}$. From the RLL encoding of the strings \mathbf{x}_i 's, observe that \mathbf{u} does not contain any occurrences of $10^{f(n)}1$ except those explicitly added to the encoded indices by Construction A. Since $|\mathbf{z}_j| = L_{\min}$ for all j , either \mathbf{u} contains an occurrence of $10^{f(n)}1$ or it has a suffix-prefix pair whose concatenation is $10^{f(n)}1$ (this follows from Construction A and the assumption that \mathbf{u} does not begin with a proper suffix of $\mathbf{z}_K 0^{n \bmod L_{\min}}$). In both cases, we will show that the precise location of the (perhaps incomplete) occurrence of \mathbf{c}_i'' in \mathbf{u} can be deduced, for some i .

Let j be the (unique) location in \mathbf{u} of the substring $10^{f(n)}1$. If $j \geq \alpha$, then \mathbf{u} contains a complete occurrence of the encoded index \mathbf{c}_i'' , and so the index \mathbf{c}_i , and therefore i , are readily obtained. Otherwise, $j < \alpha$ and let \mathbf{c}'' be the cyclic α -substring of \mathbf{u} strictly preceding the substring $10^{f(n)}1$ which starts at location $L_{\min} - (\alpha - j)$. The substring \mathbf{c}'' is obtained by the concatenation of the $(\alpha - j)$ -suffix of \mathbf{u} with the j -prefix of \mathbf{u} . The proof is now concluded by Lemma 11. ■

We remark that the described procedure operates in run-time which is linear in the substring length. In addition, if \mathbf{z} can be reconstructed from its non-overlapping substrings, then the strings \mathbf{y}_i 's are readily obtained, and \mathbf{x} may be decoded (again, see [20], [37]). These algorithms also require a linear number of operations. This completes the proof of Theorem 14.

Example 17 We return to Example 13, to demonstrate the operation of Theorem 14. Recall, for $q = 2, n = 45, L_{\min} = 14, f(n) = 2$, that we have constructed the following code-word

$$\mathbf{z} = 101010100101101011111001111011111010010000000.$$

Suppose that we receive the following (14, 20)-segmentation of \mathbf{z} :

$$\{\{10101010010110101, 1111001111011111, 010010000000\}.$$

Note since $|010010000000| = 12 < L_{\min}$, it might readily be inferred that it is the suffix of \mathbf{z} . We therefore only need identify the locations of the other two segments.

- The segment 101010 1001 0110101 contains the marker $10^{f(n)}1 = 1001$, and we therefore conclude that 101010 (given $\alpha = 6$) is an encoded-index, which as we recall corresponds to the Gray-code element $\mathbf{c}_0 = 00$. It follows that $\mathbf{y}_0 = 0110$, and 101 is a prefix of \mathbf{z}_1 (observe that the following segment of \mathbf{z}_1 could not have been immediately identify, if more segments were received).
- Next, the segment 111 1001 111011111 also contains a marker, implying that 111 is the suffix of \mathbf{c}_i'' , and 111 the prefix of \mathbf{c}_{i+1}'' . Concatenating, we have the index $\mathbf{c}'' = 111111$ and $\mathbf{c}' = 11 1$, which is an instance of \mathbf{c}_2 containing an erroneous parity symbol. Hence we deduce $i = 1$, and $\mathbf{y}_1 = 1110$.

Together, the decoding $\mathbf{x}_0 = 001$ and $\mathbf{x}_1 = 110$ may now be performed, reconstructing the original information sequence.

Lastly the redundancy of Construction A is analyzed.

Theorem 18 Using the RLL encoders of [20], [37] in Construction A, it holds that

$$\begin{aligned} \text{red}(\mathcal{C}_A(n)) &\leq \frac{n}{a} \left(1 + \frac{f(n)}{\log(n)} + \frac{1}{f(n) - 1} + \right. \\ &\quad \left. \frac{9 + 2/(f(n) - 1)}{\log(n)} + \frac{4a}{q^{f(n)}} + \frac{2a^2 + 2}{n} \right) \\ &= \frac{n}{a} \left(1 + (1 + o(1)) \left(\frac{f(n)}{\log(n)} + \frac{1}{f(n)} \right) \right). \end{aligned}$$

In particular, the redundancy is optimized for $f(n) = (1 + o(1))\sqrt{\log(n)}$, i.e.,

$$\text{red}(\mathcal{C}_A(n)) \leq \frac{n}{a} \left(1 + \frac{2 + o(1)}{\sqrt{\log(n)}} \right).$$

Proof: From Construction A, observe that $\text{red}(\mathcal{C}_A(n)) = (n \bmod L_{\min}) + L_{\min} + K(L_{\min} - m(N))$ and

$$\begin{aligned} L_{\min} - N &= \alpha + f(n) + 2 \\ &= \left\lceil \frac{f(n)}{f(n) - 1} (I + 1) \right\rceil + f(n) + 2 \\ &\leq \frac{f(n)}{f(n) - 1} (I + 1) + f(n) + 3 \\ &\leq \frac{f(n)}{f(n) - 1} (\log(n/L_{\min}) + 2) + f(n) + 3 \\ &\leq \log(n) + f(n) + \frac{\log(n)}{f(n) - 1} + 5 + \frac{2}{f(n) - 1}. \end{aligned}$$

Further, by [37, Lem. 4] one may efficiently encode $\mathbf{x} \mapsto \mathbf{y}$ such that $N - m(N) \leq \lceil \frac{q}{q-2} \cdot \frac{N}{q^{f(n)}} \rceil$ (For $q = 2$ [20, Sec. III] showed $N - m(N) \leq 2\lceil N/q^{f(n)-1} \rceil$), and we shall use the overly zealous upper bound $N - m(N) \leq \frac{4N}{q^{f(n)}} + 2 \leq \frac{4a \log(n)}{q^{f(n)}} + \frac{4}{q} + 2 \leq \frac{4a \log(n)}{q^{f(n)}} + 4$.

Finally, we get that

$$\begin{aligned} \text{red}(\mathcal{C}_A(n)) &= K(L_{\min} - m(N)) + L_{\min} + (n \bmod L_{\min}) \\ &\leq \frac{n}{a \log(n)} \left(\log(n) + f(n) + \frac{\log(n)}{f(n) - 1} + \right. \\ &\quad \left. 9 + \frac{2}{f(n) - 1} + \frac{4a \log(n)}{q^{f(n)}} \right) + 2L_{\min} \\ &\leq \frac{n}{a} \left(1 + \frac{f(n)}{\log(n)} + \frac{1}{f(n) - 1} + \right. \\ &\quad \left. \frac{9 + 2/(f(n) - 1)}{\log(n)} + \frac{4a}{q^{f(n)}} + \frac{2a^2 + 2}{n} \right), \end{aligned}$$

which completes the proof of the first part. The second part follows by substitution of $f(n) = (1 + o(1))\sqrt{\log(n)}$ into the former. ■

By Theorem 18 and Corollary 2, efficient encoding and decoding is possible at asymptotically optimal rates. In comparison to Construction P (by Lemma 9), Construction A asymptotically achieves rate $1 - \frac{1}{a} - O\left(\frac{f(n)}{\log(n)} + \frac{1}{f(n)}\right)$ instead of $1 - \frac{1}{|a|} - O\left(\frac{1}{\log(n)}\right)$, for any channel parameter $a > 1$ (here, the integer value is used since Construction P must be operated at $m \triangleq \lceil a \rceil$ to produce an (L_{\min}, L_{\max}) -torn-paper code). For completeness, we also include specific construction parameters for several arbitrary choices of n, L_{\min} , and compare resulting rates, in Tables I to III (all for $q = 4$). It should however be stressed that, for Construction P, the choice of s, m optimizing the resulting rate $R(\mathcal{C}_{\text{pilot}}(n)) \geq (1 - \frac{1}{m}) \cdot R(\mathcal{O}_p)$ is not straightforward, even given the lower bounds of Lemmas 8 and 9; indeed, $R(\mathcal{O}_p)$ cannot easily be computed, for an optimal choice of p . We rely in our comparison on the lower-bounds of Lemmas 8 and 9 instead; note in particular that even for the same choice of n, m, s , i.e., for a specific code, these might provide distinct lower-bounds on the rate. As mentioned above, even then it is not immediately clear how to efficiently encode and decode $\mathcal{C}_{\text{pilot}}(n)$.

Next, we consider the case of $k > 1$ and $\log(k) = o(n)$. We know from Corollary 2 that if $\limsup \frac{L_{\min}}{nk} \leq 1$ then any family of (L_{\min}, L_{\max}) -multistrand torn-paper codes will only achieve vanishing asymptotic rate; hence we assume $L_{\min} = \lceil a \log(nk) \rceil$ for some $a > 1$. The following theorem summarizes our main results regarding (L_{\min}, L_{\max}) -multistrand torn-paper codes.

Theorem 19 *Take n, k such that $k > 1, \log(k) = o(n)$, and let $L_{\min} = \lceil a \log(nk) \rceil$, for $a > 1$. There exists a linear run-time (in the substrings length, i.e., nk) encoder-decoder pair for (L_{\min}, L_{\max}) -multistrand torn-paper codes achieving $1 - \frac{1}{a} - o_{nk}(1)$ asymptotic rate.*

Proof: Theorem 19 is justified by a simple amendment of Construction A. We encode $\mathbf{x} \in \Sigma^{kK^m}$ into $\{\{z^{(j)} : j \in [k]\}\}$,

where $|z^{(j)}| = n$ for all $j \in [k]$, as follows. We modify $I \triangleq \lceil \log(k \lceil n/L_{\min} \rceil) \rceil$ (recall, also, $\alpha \triangleq \lceil \frac{f(n)}{f(n)-1} (I+1) \rceil$) and $L_{\min} = \lceil a \log(nk) \rceil$. We then denote $\mathbf{x} = \mathbf{x}^{(0)} \circ \mathbf{x}^{(1)} \circ \dots \circ \mathbf{x}^{(k-1)}$, where $|x^{(j)}| = Km$ for all $j \in [k]$, and apply Algorithm 1 to $(x^{(j)})_{j \in [k]}$ in succession; observe that every operation requires only $\lceil n/L_{\min} \rceil$ distinct indices, and we utilize available indices in order throughout the k operations.

We observe that the proofs of Lemmas 15 and 16 hold without change, hence this amendment encodes into an (L_{\min}, L_{\max}) -multistrand torn-paper code, which we denote $\mathcal{C}_A(n, k) \in \mathcal{X}_{n,k}$. Finally, following the proof of Theorem 18 we have

$$\begin{aligned} \text{red}(\mathcal{C}_A(n, k)) &= k(K(L_{\min} - m(N)) + L_{\min} + \\ &\quad (n \bmod L_{\min})) \\ &\leq \frac{nk}{a} \left(1 + (1 + o(1)) \left(\frac{f(nk)}{\log(nk)} + \frac{1}{f(nk)} \right) \right); \end{aligned}$$

As in Theorem 18, for $f(n) = (1 + o(1))\sqrt{\log(nk)}$ we have

$$\text{red}(\mathcal{C}_A(n, k)) \leq \frac{nk}{a} \left(1 + \frac{2 + o(1)}{\sqrt{\log(nk)}} \right).$$

From Lemma 1 we have $\log|\mathcal{X}_{n,k}| \geq (n - \log(k))k$, concluding the proof. ■

Again, by Corollary 2 and Theorem 18 the rate of the construction is asymptotically optimal.

IV. ERROR-CORRECTING TORN-PAPER CODES

In this section, we extend the study of torn-paper codes to a noisy setup. We consider two models of noise. The first one assumes that the encoded string, before segmentation, suffers at most some t substitution errors. The second model corresponds to the case where some of the segments are deleted during segmentation.

A. Substitution-Correcting Torn-paper Codes

For a string \mathbf{x} , its t -error torn-paper ball, denoted by $\mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{x}; t)$, is defined as the set of all possible (L_{\min}, L_{\max}) -segmentations after introducing at most t errors to \mathbf{x} , that is,

$$\mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{x}; t) \triangleq \bigcup_{\mathbf{y} \in B_t(\mathbf{x})} \mathcal{T}_{L_{\min}}^{L_{\max}}(\mathbf{y}),$$

where $B_t(\mathbf{x}) = \{\mathbf{y} : d_H(\mathbf{x}, \mathbf{y}) \leq t\}$ is the radius- t Hamming ball centered at \mathbf{x} . A code \mathcal{C} is called a t -error single-strand torn-paper code if for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}, \mathbf{x}_1 \neq \mathbf{x}_2$, it holds that

$$\mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{x}_1; t) \cap \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{x}_2; t) = \emptyset.$$

Our goal in this section is to show how to adjust Construction A in order to produce t -error single-strand torn-paper codes. We first explain the main ideas of the required modifications. Let $\mathbf{z} = \text{Enc}_A(\mathbf{x}) \in \mathcal{C}_A(n)$ (encoded with Algorithm 1) and let $\mathcal{U} \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$ be an (L_{\min}, L_{\max}) -segmentation of some word \mathbf{z}' , where $d_H(\mathbf{z}, \mathbf{z}') \leq t$. The main task of the noiseless decoder of $\mathcal{C}_A(n)$ was to first calculate the index, and thus the location in \mathbf{z} , of every segment $\mathbf{u} \in \mathcal{U}$.

TABLE I
CONSTRUCTION P (LEMMA 8): SPECIFIC PARAMETERS $(m, s, R(\mathcal{C}_{\text{pilot}}(n)))$.

$L_{\min} \setminus n$	60	250	4000	60,000	400,000	6,000,000
10	2, 5, 0.379	n/a	n/a	n/a	n/a	n/a
50	15, 3, 0.856	10, 5, 0.844	5, 10, 0.798	3, 16, 0.667	2, 25, 0.5	2, 25, 0.5
100	n/a	10, 10, 0.9	10, 10, 0.9	6, 16, 0.833	5, 20, 0.8	4, 25, 0.75
300	n/a	n/a	32, 9, 0.968	25, 12, 0.96	20, 15, 0.95	15, 20, 0.933
1000	n/a	n/a	125, 8, 0.992	100, 10, 0.989	64, 15, 0.984	50, 20, 0.98

(Bold-face indicates Lemma 8 provides highest lower-bound on rate.)

TABLE II
CONSTRUCTION P (LEMMA 9): SPECIFIC PARAMETERS $(m, s, R(\mathcal{C}_{\text{pilot}}(n)))$.

$L_{\min} \setminus n$	60	250	4000	60,000	400,000	6,000,000
10	2, 5, 0.45	n/a	n/a	n/a	n/a	n/a
50	15, 3, 0.778	10, 5, 0.81	5, 10, 0.76	5, 10, 0.76	4, 12, 0.719	3, 16, 0.646
100	n/a	10, 10, 0.855	10, 10, 0.855	10, 10, 0.855	8, 12, 0.839	6, 16, 0.807
300	n/a	n/a	25, 12, 0.92	25, 12, 0.92	25, 12, 0.92	20, 15, 0.918
1000	n/a	n/a	50, 20, 0.956	50, 20, 0.956	50, 20, 0.956	50, 20, 0.956

(Bold-face indicates Lemma 9 provides highest lower-bound on rate. Background pattern indicates that the choice of m, s is only guaranteed by Lemma 9.)

TABLE III
CONSTRUCTION A (THEOREM 18): SPECIFIC PARAMETERS $(f, I, N, K, R(\mathcal{C}_A(n)))$.

$L_{\min} \setminus n$	60	250	4000	60,000	400,000	6,000,000
10	n/a	n/a	n/a	n/a	n/a	n/a
50	n/a	2, 2, 40, 4, 0.56	3, 4, 38, 79, 0.711	3, 6, 35, 1199, 0.659	4, 7, 34, 7999, 0.66	4, 9, 31, 119999, 0.6
100	n/a	2, 1, 92, 1, 0.32	3, 3, 89, 39, 0.839	3, 5, 86, 599, 0.829	4, 6, 85, 3999, 0.84	4, 8, 82, 59999, 0.81
300	n/a	n/a	3, 2, 291, 12, 0.843	3, 4, 288, 199, 0.925	4, 6, 9, 285, 1332, 0.939	4, 8, 282, 19999, 0.93
1000	n/a	n/a	3, 1, 992, 3, 0.721	3, 3, 989, 59, 0.942	4, 5, 986, 399, 0.976	4, 7, 984, 5999, 0.976

(Bold-face indicates Theorem 18 provides highest lower-bound on rate.)

However, in the presence of errors, calculating the index of a segment $\mathbf{u} \in \mathcal{U}$ based on the first (perhaps partial) occurrence of an encoded index within \mathbf{u} might result with the misplacement of all the (perhaps partial) information blocks \mathbf{y}_i that are contained in \mathbf{u} . Hence, a more careful approach is necessary for index decoding.

Before presenting our construction for t -error single-strand torn-paper codes, we introduce several additional required definitions. For a string \mathbf{u} , define $\mathcal{T}_{L_{\min}}^+(\mathbf{u})$ to be the multiset of non-overlapping L_{\min} -segments of \mathbf{u} , where the last segment is of length ℓ , $L_{\min} \leq \ell < 2L_{\min}$. A segment $\mathbf{w} \in \mathcal{T}_{L_{\min}}^+(\mathbf{u})$ is called *A-decodable* if, informally, Algorithm 2 returns (perhaps erroneous) output when given \mathbf{w} as input. More formally, if \mathbf{w} satisfies one of the following conditions.

- 1) \mathbf{w} either contains a unique complete occurrence of $10^{f(n)}1$, or it doesn't contain complete occurrences but contains a cyclic occurrence (if $L_{\min} < |\mathbf{w}| < 2L_{\min}$, require instead that either the L_{\min} -prefix or the L_{\min} -suffix of \mathbf{w} contain a cyclic occurrence).
- 2) \mathbf{w} contains precisely two complete occurrences of $10^{f(n)}1$, and there exist a unique pair of occurrences (either complete or complete-to-suffix/prefix) whose locations are at distance precisely L_{\min} . Recall that \mathbf{w} cannot contain more than two complete occurrences of $10^{f(n)}1$, except in the presence of errors, hence those cases can safely be discarded (see part 1 in the proof of Theorem 20 for a formal proof).

Let \mathbf{w} be an A-decodable segment. Then, by definition, there is at least one occurrence (perhaps cyclic) of $10^{f(n)}1$ within \mathbf{w} and, if there is more than a single occurrence, then

there is exactly one pair of occurrences such that the difference between their locations is L_{\min} . Consider the α -segments of \mathbf{w} preceding these occurrences as encoded indices; if the (first) occurrence of $10^{f(n)}1$ in \mathbf{w} is at location $\ell < \alpha$, concatenate the $(\alpha - \ell)$ -segment of \mathbf{w} at location $L_{\min} + \ell - \alpha$, to the ℓ -prefix of \mathbf{w} , and consider the resulting length- α string to be a *cyclic encoded index*.

An A-decodable segment \mathbf{w} is called *valid* if, informally, there appears at least one 'valid' encoded index in \mathbf{w} , and no conflicting pair of such indices. More formally, a valid segment \mathbf{w} is an A-decodable segment that satisfies one of the following conditions:

- 1) \mathbf{w} contains no complete encoded index, hence it contains only a cyclic encoded index.
- 2) \mathbf{w} contains a single complete encoded index, and its parity symbol is correct.
- 3) \mathbf{w} contains two complete encoded indices, and either exactly one of their parity symbols is correct, or both are correct and the indices are consecutive in the applied Gray code.

Construction B We construct a concatenated code, using Construction A as inner-code, and an arbitrary $(K, q^{m(N)M}, 2t + 1)_{q^{m(N)}}$ error-correcting code \mathcal{C}_{EC} , with an encoding algorithm $\text{Enc}_{\text{EC}}: (\Sigma^{m(N)})^M \rightarrow (\Sigma^{m(N)})^K$, as outer-code (here, K, N are the parameters of Construction A). The resulting t -error (L_{\min}, L_{\max}) -single-strand torn-paper code is denoted $\mathcal{C}_B(n)$, with the associated encoder $\text{Enc}_B: \Sigma^{m(N)M} \rightarrow \Sigma^n$. \square

We observe the following property of Construction B. Assume one retrieves a noisy version \mathbf{z}' of $\mathbf{z} = \text{Enc}_B(\mathbf{x})$, e.g., from any reconstruction algorithm; further assume that \mathbf{z}, \mathbf{z}' agree on all locations containing encoded indices \mathbf{c}_i'' or markers $10^{f(n)}1$ (as their locations in \mathbf{z} are known a priori and do not depend on the information \mathbf{x}). Thus, one extracts from \mathbf{z}' (perhaps erroneous) encoded information blocks, denoted \mathbf{y}'_i . Denote by e the number of encoded information blocks that were not recovered (e.g., due to conflicts in the reconstruction algorithm), and by s the number of encoded blocks that were recovered incorrectly (i.e., $\mathbf{y}'_i \neq \mathbf{y}_i$). Since the information string $(\mathbf{x}_i)_{i \in [M]} \in \Sigma^{m(N)M}$ is encoded using a $(K, q^{m(N)M}, 2t+1)_{q^{m(N)}}$ error-correcting code, it suffices that $2s + e \leq 2t$ to guarantee correct decoding.

In order to reconstruct a noisy version \mathbf{z}' of \mathbf{z} , we define a modification of Algorithm 2, as follows. First, given an (L_{\min}, L_{\max}) -segmentation $\mathcal{U} \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$ we apply the reconstruction algorithm not directly to \mathcal{U} , but rather to valid segments in

$$\mathcal{T}_{L_{\min}}^+(\mathcal{U}) \triangleq \{\{\mathcal{T}_{L_{\min}}^+(\mathbf{u}) : \mathbf{u} \in \mathcal{U}\}\}.$$

Secondly, in case a valid $\mathbf{w} \in \mathcal{T}_{L_{\min}}^+(\mathcal{U})$ contains multiple (perhaps cyclic) occurrences of an encoded index, the algorithm selects one to decode by prioritizing complete occurrences over cyclic ones, and in the case of complete occurrences, accepting the first containing a correct parity symbol (since \mathbf{w} is valid, such occurrence exists in this case). Decoding of the selected encoded index is then performed as described in Algorithm 2, and denoted by $\text{Ind}'(\mathbf{w})$.

For an (L_{\min}, L_{\max}) -segmentation $\mathcal{U} \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$, we define the set

$$\mathcal{Z}(\mathcal{U}) \triangleq \{(\text{Ind}'(\mathbf{w}), \mathbf{w}) : \mathbf{w} \in \mathcal{T}_{L_{\min}}^+(\mathcal{U}) \text{ is valid}\}.$$

If $(j, \mathbf{w}), (j, \mathbf{w}') \in \mathcal{Z}(\mathcal{U})$ for some j and $\mathbf{w} \neq \mathbf{w}'$, we define a restriction $\mathcal{Z}'(\mathcal{U})$ of $\mathcal{Z}(\mathcal{U})$ by including only the shortest, lexicographically-least, segment (i.e., $\mathcal{Z}'(\mathcal{U})$ defines a proper function).

Given the set $\mathcal{Z}'(\mathcal{U})$ we decode a string \mathbf{z}' as follows.

- 1) Fill the encoded indices and the markers in \mathbf{z}' in the correct locations as defined in Algorithm 1 (note again that these locations do not depend on the information).
- 2) Next, we iterate over any pair $(\text{Ind}'(\mathbf{w}), \mathbf{w}) \in \mathcal{Z}'(\mathcal{U})$ and update \mathbf{z}' with the symbols of the encoded blocks \mathbf{y}_i 's within \mathbf{w} ; If there is a collision of symbols in the same position within an encoded block \mathbf{y}_i for some $i, i \in [K]$, we erase \mathbf{y}_i completely from \mathbf{z}' .
- 3) If an encoded block \mathbf{y}_i is partially filled at the end of the process (i.e., there are missing symbols within \mathbf{y}_i) we erase the encoded block \mathbf{y}_i .

The output \mathbf{z}' of this decoding procedure over the segmentation $\mathcal{U} \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$ is denoted by $\text{Dec}_B(\mathcal{U}) \triangleq \mathbf{z}'$.

We now prove that $\mathcal{C}_B(n)$ is a t -error single-strand torn-paper code.

Theorem 20 *Let $\mathbf{z} = \text{Enc}_B(\mathbf{x})$, $\mathcal{U} \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$, and let $\mathbf{z}' = \text{Dec}_B(\mathcal{U})$ be the noisy version of \mathbf{z} reconstructed by the aforementioned algorithm. Then, it holds that $2s + e \leq 2t$,*

where e, s are defined as previously explained; i.e., any inner-channel error propagates as, at most, either one outer-channel error or two outer-channel erasures.

Proof: By definition, \mathcal{U} is obtained by first introducing up to t errors to \mathbf{z} , and then performing an (L_{\min}, L_{\max}) -segmentation to the obtained word. For the rest of the proof, we fix an arbitrary (L_{\min}, L_{\max}) -segmentation pattern, and for $\mathbf{z} \in \Sigma^n$ we denote $\mathcal{U} \in \mathcal{T}_{L_{\min}}^{L_{\max}}(\mathbf{z})$ obtained from this pattern by $\mathcal{U} = T(\mathbf{z})$. In particular, observe for $\|\mathbf{v}\| \leq t$ that $T(\mathbf{z} + \mathbf{v}) \in \mathcal{BT}_{L_{\min}}^{L_{\max}}(\mathbf{z}; t)$.

For convenience, we denote by $\mathbf{z}'_{\mathbf{v}} \triangleq \text{Dec}_B(T(\mathbf{z} + \mathbf{v}))$, and by $e_{\mathbf{v}}$ (respectively, $s_{\mathbf{v}}$) the number of encoded information blocks \mathbf{y}'_i in $\mathbf{z}'_{\mathbf{v}}$ which were not recovered (recovered erroneously). We shall prove the following proposition, which justifies the claim. Let $\mathbf{z} \triangleq \text{Enc}_B(\mathbf{x})$, $\mathbf{v} \in \Sigma^n$ such that $\|\mathbf{v}\| \leq t$. Then $e_{\mathbf{v}} + 2s_{\mathbf{v}} \leq 2\|\mathbf{v}\|$.

The proof is done by induction on $\|\mathbf{v}\|$. First observe by Lemma 16 that $e_{\mathbf{0}} = s_{\mathbf{0}} = 0$ (here, $\mathbf{0}$ is the all-zero string). For the induction step, assume that the claim holds for any $\mathbf{v}' \in \Sigma^n$, $\|\mathbf{v}'\| < t$. Let $\mathbf{v} \in \Sigma^n$, $\|\mathbf{v}\| = t$. Take any $\mathbf{u}' \in \mathcal{T}_{L_{\min}}^+(T(\mathbf{z} + \mathbf{v}))$ affected by $t' > 0$ errors. Decompose $\mathbf{v} = \mathbf{v}' + \mathbf{v}''$ such that $\|\mathbf{v}'\| = t'$, $\|\mathbf{v}''\| = t - t'$, and \mathbf{u}' contains the support of \mathbf{v}' . Consider the decoder output $\mathbf{z}'_{\mathbf{v}''}$; by the induction assumption,

$$e_{\mathbf{v}''} + 2s_{\mathbf{v}''} \leq 2(t - t').$$

We denote by \mathbf{u} the segment corresponding to \mathbf{u}' in $\mathbf{z}'_{\mathbf{v}''}$. Note that \mathbf{u} contains no errors and its index is recovered correctly by the decoder. Hence, each encoded block that intersects \mathbf{u} is either correct in $\mathbf{z}'_{\mathbf{v}''}$, or it is erased due to errors in other segments.

Denoting by δ the number of encoded information blocks intersecting \mathbf{u} , we let (i) ρ_1 be the number of those recovered correctly in $\mathbf{z}'_{\mathbf{v}''}$; (ii) ρ_2 be the number of those erased in $\mathbf{z}'_{\mathbf{v}''}$ due to collisions resulting from incorrect index-decoding in other segments; and (iii) ρ_3 be the number of those erased in $\mathbf{z}'_{\mathbf{v}''}$ due to erasures of other, intersecting, segments. Observe that $\delta = \rho_1 + \rho_2 + \rho_3 \in \{1, 2, 3\}$, depending on $|\mathbf{u}|$ and its location.

The rest of the proof is done by cases.

- 1) If \mathbf{u}' is not valid, then all encoded information blocks intersecting \mathbf{u}' are erased at the decoder. Hence, the ρ_1 correctly recovered blocks in $\mathbf{z}'_{\mathbf{v}''}$ which intersect \mathbf{u} are erased in $\mathbf{z}'_{\mathbf{v}}$.

In addition, each of the ρ_2 blocks corresponding to blocks intersecting \mathbf{u} which are erased due to collisions, might instead cause incorrect recovery of encoded information blocks in \mathbf{u}' . Hence, $e_{\mathbf{v}} \leq e_{\mathbf{v}''} + \rho_1 - \rho_2$ and $s_{\mathbf{v}} \leq s_{\mathbf{v}''} + \rho_2$, and we note

$$\begin{aligned} e_{\mathbf{v}} + 2s_{\mathbf{v}} &\leq (e_{\mathbf{v}''} + 2s_{\mathbf{v}''}) + \rho_1 + \rho_2 \\ &\leq 2(t - t') + \delta = 2\|\mathbf{u}\| - (2t' - \delta). \end{aligned}$$

Since $t' \geq 1$, we have $e_{\mathbf{v}} + 2s_{\mathbf{v}} \leq 2\|\mathbf{u}\|$ unless $\delta = 3$; however, in that case \mathbf{u} contains two complete instances of $10^{f(n)}1$ whose locations are at distance L_{\min} , both preceded by complete occurrences of encoded indices,

and since \mathbf{u}' is not valid we have $t' > 1$, which also concludes the proof.

- 2) If \mathbf{u}' is valid but its index is incorrectly decoded, then the ρ_1 encoded information blocks that are recovered correctly in $\mathbf{z}'_{v''}$ are erased in \mathbf{z}'_v , and ρ_2 encoded information blocks, corresponding to those intersecting \mathbf{u} which are erased in $\mathbf{z}'_{v''}$ due to collisions, might be recovered incorrectly in \mathbf{z}'_v .

Furthermore, the placement of \mathbf{u}' at an incorrectly decoded location causes δ additional encoded information blocks to be either erased (due to collisions) or incorrectly recovered (where the correct blocks appear in invalid segments, i.e., are erased in $\mathbf{z}'_{v''}$). Denoting the number of blocks of the former type by δ_1 , and the latter δ_2 , we then have $e_v = e_{v''} + \rho_1 - \rho_2 + \delta_1 - \delta_2$ and $s_v \leq s_{v''} + \rho_2 + \delta_2$. Hence,

$$\begin{aligned} e_v + 2s_v &\leq (e_{v''} + 2s_{v''}) + \rho_1 + \rho_2 + \delta_1 + \delta_2 \\ &\leq 2(t - t') + 2\delta = 2\|\mathbf{v}\| - 2(t' - \delta). \end{aligned}$$

To conclude, we require $t' \geq \delta$. Indeed, observe that if $\delta = 2$ then \mathbf{u} contains a complete occurrence of an encoded index followed by $10^{f(n)}1$, requiring $t' \geq 2$ for incorrect recovery. Likewise, if $\delta = 3$ then \mathbf{u} contains two complete occurrences of encoded indices whose locations are at distance L_{\min} , each followed by $10^{f(n)}1$; incorrect recovery of the index therefore requires at least two errors in one of them in addition to further errors in the other index or $10^{f(n)}1$ marker, or an error in each $10^{f(n)}1$ marker in addition to further errors to generate such a marker at an alternative location, hence $t' \geq 3$ as well.

- 3) Finally, if the index of \mathbf{u}' is decoded correctly (and, in particular, \mathbf{u}' is valid), then recalling that the index of \mathbf{u} is also decoded correctly, we clearly have $e_v = e_{v''}$. Since any error in \mathbf{u}' can cause an error in at most a single encoded information block, we have that $s_v \leq s_{v''} + t'$. Hence,

$$\begin{aligned} e_v + 2s_v &\leq e_{v''} + 2(s_{v''} + t') \\ &= (e_{v''} + 2s_{v''}) + 2t' \leq 2t = 2\|\mathbf{v}\|. \quad \blacksquare \end{aligned}$$

Theorem 21 Denote the redundancy of the outer-code \mathcal{C}_{EC} used in Construction B by $\rho_{EC} \triangleq K - M$. Then, operating Enc_A as in Theorem 18, with $f(n) = (1 + o(1))\sqrt{\log(n)}$, we have

$$\begin{aligned} \text{red}(\mathcal{C}_B(n)) &\leq \frac{n}{a} \left(1 + \frac{f(n)}{\log(n)} + \frac{1}{f(n) - 1} + \frac{9 + 2/(f(n) - 1)}{\log(n)} + \frac{4a}{q^{f(n)}} + \frac{2a^2 + 2}{n} \right) + \\ &\rho_{EC} \left((a - 1) \log(n) - 2\sqrt{\log(n)} - \right. \\ &\quad \left. 11 - \frac{3}{\sqrt{\log(n)} - 1} - \frac{4a \log(n)}{q\sqrt{\log(n)}} \right) \\ &= \frac{n}{a} \left(1 + \frac{2 + o(1)}{\sqrt{\log(n)}} \right) + \\ &\rho_{EC} \left((a - 1) \log(n) - (2 + o(1))\sqrt{\log(n)} \right). \end{aligned}$$

Furthermore, when $a > 2$ then the outer-code \mathcal{C}_{EC} can be an MDS code and hence $\rho_{EC} = 2t$.

Proof: By Construction B, $\text{red}(\mathcal{C}_B(n)) = \text{red}(\mathcal{C}_A(n)) + \rho_{EC} \cdot m(N)$.

We recall from the proof of Theorem 18 that for $f(n) = \lceil \sqrt{\log(n)} \rceil$ it holds that

$$\begin{aligned} m(N) &\geq L_{\min} - \log(n) - f(n) - \frac{\log(n)}{f(n) - 1} \\ &\quad 9 - \frac{2}{f(n) - 1} - \frac{4a \log(n)}{q^{f(n)}} \\ &\geq (a - 1) \log(n) - 2\sqrt{\log(n)} - \\ &\quad 11 - \frac{3}{\sqrt{\log(n)} - 1} - \frac{4a \log(n)}{q\sqrt{\log(n)}}, \end{aligned}$$

satisfying the former part of claim.

Next, for $a > 2$ we observe that $m(N) > \log(n) - \log \log(n) + O_n(1) = \log(K)$, implying that an RS code may be used in Construction B, satisfying the latter part. \blacksquare

Before concluding the section, we outline an extension of Construction B to the case $k > 1$, i.e., to t -error multi-strand torn-paper codes.

Corollary 22 Take n, k such that $k > 1$, $\log(k) = o(n)$; let $L_{\min} = \lceil a \log(nk) \rceil$, for $a > 1$, and take some $L_{\max} \geq L_{\min}$. Amend Construction B as was done in Theorem 19 to Construction A, using a $(kK, q^{m(N)M}, 2t + 1)_{q^{m(N)}}$ error-correcting code \mathcal{C}_{EC} , with redundancy $\rho_{EC} \triangleq kK - M$. Then the resulting code $\mathcal{C}_B(n, k)$ is a t -error (L_{\min}, L_{\max}) -multistrand torn-paper code, satisfying

$$\begin{aligned} \text{red}(\mathcal{C}_B(n, k)) &\leq \frac{nk}{a} \left(1 + \frac{2 + o(1)}{\sqrt{\log(nk)}} \right) + \\ &\rho_{EC} \left((a - 1) \log(nk) - \right. \\ &\quad \left. (2 + o(1))\sqrt{\log(nk)} \right). \end{aligned}$$

Proof: The proof of Theorem 20 applies without change. As in Theorem 19, we have

$m(N) = (a - 1) \log(nk) - (1 + o(1)) \left(f(nk) + \frac{\log(nk)}{f(nk)} \right)$, and following the steps of Theorem 21, we have the claimed upper bound on redundancy, for $f(n) = (1 + o(1))\sqrt{\log(nk)}$. \blacksquare

B. Deletion-Correcting Torn-paper Codes

For a string \mathbf{x} , its t -deletion torn-paper ball, $\mathcal{DT}_{L_{\min}}^{L_{\max}}(\mathbf{x}; t)$, is defined as all the subsets with at most t missing segments of all the possible (L_{\min}, L_{\max}) -segmentations of \mathbf{x} , that is,

$$\mathcal{DT}_{L_{\min}}^{L_{\max}}(\mathbf{x}; t) \triangleq \bigcup_{S \in \mathcal{T}_{L_{\min}}^{L_{\max}}(\mathbf{x})} \{S' \subseteq S : |S| - |S'| \leq t\}.$$

A code \mathcal{C} is called a t -deletion torn-paper code if for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}$ it holds that $\mathcal{DT}_{L_{\min}}^{L_{\max}}(\mathbf{x}_1; t) \cap \mathcal{DT}_{L_{\min}}^{L_{\max}}(\mathbf{x}_2; t) = \emptyset$.

In this section, we utilize *burst-erasure-correcting (BEC) codes* in our constructions, which are defined next. For a string

\mathbf{x} , its t -burst L -erasures ball, denoted by $\mathcal{B}_{\text{BE}}^L(\mathbf{x}; t)$, is defined as the set of all strings that can be obtained from \mathbf{x} by at most t burst of erasures, each of length at most L . A code \mathcal{C} is called a t -burst L -erasure correcting code if for all $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{C}$, $\mathcal{B}_{\text{BE}}^L(\mathbf{x}_1; t) \cap \mathcal{B}_{\text{BE}}^L(\mathbf{x}_2; t) = \emptyset$.

Next, we present a generic construction of t -deletion torn-paper codes. Let $\widehat{L}_{\text{max}} \triangleq L_{\text{max}} - \lceil \frac{L_{\text{max}}}{L_{\text{min}}} \rceil (\alpha + f(n) + 2)$. This construction is based on Construction A and assumes the existence of a systematic linear t -burst \widehat{L}_{max} -erasure correcting code, denoted by \mathcal{C}_{BEC} .

Construction C Let $\rho > 0$ be an integer that is determined next. This construction uses the following family of codes:

Systematic BEC encoding. Let $\text{ENC}_{\text{BEC}}: \Sigma^{(K-\rho)N} \rightarrow \Sigma^{\rho \text{BEC}}$ denote the systematic encoder of the code \mathcal{C}_{BEC} , such that for any string $\mathbf{v} \in \Sigma^{(K-\rho)N}$, $\mathbf{v} \circ \text{ENC}_{\text{BEC}}(\mathbf{v}) \in \mathcal{C}_{\text{BEC}}$ (for convenience we assume that $\text{ENC}_{\text{BEC}}(\mathbf{v})$ returns only the encoded systematic redundancy symbols). The redundancy of this encoder is denoted by ρ_{BEC} . The parameter ρ is defined $\rho \triangleq \left\lceil \frac{1}{N} \left\lceil \rho_{\text{BEC}} \cdot \frac{f(n)}{f(n)-1} \right\rceil \right\rceil$.

Next, we utilize a generalized concatenated coding approach, where Construction A is used as inner-code for $K - \rho$ information blocks, and with a slight adjustment also for the ρ redundant blocks, as follows:

- 1) *The length of the input string \mathbf{x} .* The input of this construction is $\mathbf{x} \in \Sigma^{(K-\rho)m(N)}$. That is, this construction has additional redundancy of $\rho m(N)$ symbols compared to Construction A. The input string is divided to $K - \rho$ information blocks each of length $m(N)$, denoted by $\mathbf{x}_0, \dots, \mathbf{x}_{K-\rho-1}$.
- 2) *The generation of the encoded blocks \mathbf{y}_i 's.* The first $K - \rho$ blocks are generated from the corresponding \mathbf{x}_i 's using the RLL encoder E_m^{RLL} similarly to Construction A. Let $\mathbf{y}^* \triangleq \mathbf{y}_0 \circ \dots \circ \mathbf{y}_{K-\rho-1} \in \Sigma^{(K-\rho)N}$ denote their concatenation. Next, we apply ENC_{BEC} to obtain $\mathbf{w} \triangleq \text{ENC}_{\text{BEC}}(\mathbf{y}^*)$, and denote by \mathbf{w}^* the result of inserting '1's into \mathbf{w} at every location divisible by $f(n)$ (in particular, $\mathbf{y}^* \circ \mathbf{w}^*$ does not contain a length- $f(n)$ zero-run). Then, \mathbf{w}^* is divided to the remaining segments $\mathbf{y}_{K-\rho}, \dots, \mathbf{y}_{K-1} \in \Sigma^N$ (if $|\mathbf{w}^*|$ is not a multiple of N , \mathbf{y}_{K-1} is padded with 1's to length N). Note that the parameter ρ satisfies $\rho N \geq \left\lceil \rho_{\text{BEC}} \cdot \frac{f(n)}{f(n)-1} \right\rceil = |\mathbf{w}^*|$, hence one may continue to follow the steps of Construction A without change.

We now indeed continue identically to Construction A. That is, an index and a marker are appended to the beginning of each encoded block \mathbf{y}_i to construct a segment \mathbf{z}_i of length L_{min} . Then, $\mathbf{z}_0, \dots, \mathbf{z}_{K-1}$ are concatenated along with $\mathbf{z}_K \circ 0^{n \bmod L_{\text{min}}} = \mathbf{c}_K'' \circ 10^{f(n)} 10^{N+(n \bmod L_{\text{min}})}$ to obtain the encoded output string $\mathbf{z} \in \Sigma^n$. \square

Let $\mathcal{C}_{\text{del}}(n)$ denote the constructed code. The correctness of Construction C and redundancy calculation are proved in the next theorem.

Theorem 23 $\mathcal{C}_{\text{del}}(n)$ is a t -deletion torn-paper code. Furthermore, it holds that

$$\text{red}(\mathcal{C}_{\text{del}}(n)) = \text{red}(\mathcal{C}_A(n)) + m(N) \left\lceil \frac{1}{N} \left\lceil \rho_{\text{BEC}} \cdot \frac{f(n)}{f(n)-1} \right\rceil \right\rceil.$$

Proof: Let $\mathbf{z} \in \mathcal{C}_{\text{del}}(n)$ be the encoded codeword of the input string \mathbf{x} , and take $\mathcal{U} \in \mathcal{DT}_{L_{\text{min}}}^{L_{\text{max}}}(\mathbf{z}; t)$. We shall prove that one can uniquely decode \mathbf{x} .

From Lemma 16, for every $\mathbf{u} \in \mathcal{U}$ which is not a substring of the suffix of length $(n \bmod L_{\text{min}}) + N + f(n)$ of \mathbf{z} , its index $\text{Ind}(\mathbf{u})$ can be decoded using Algorithm 2. The string \mathbf{z} can then be reconstructed by the locations of each received segment, with some segments erased (at identifiable locations). Let $\mathbf{z}' \in (\Sigma \cup \{?\})^n$ denote this partially reconstructed string, where '?' stands for erased symbols.

From the definition of $\mathcal{DT}_{L_{\text{min}}}^{L_{\text{max}}}(\mathbf{z}; t)$, at most t segments of \mathbf{z} are missing from \mathcal{U} . Therefore, $\mathbf{z}' \in \mathcal{B}_{\text{BE}}^{L_{\text{max}}}(\mathbf{z}; t)$. By removing coordinates of \mathbf{z}' corresponding to indices or markers (including '?' symbols), a string $\mathbf{y}' \in \mathcal{B}_{\text{BE}}^{\widehat{L}_{\text{max}}}(\mathbf{y}_0 \circ \dots \circ \mathbf{y}_{K-1}; t)$ is obtained, since there are at most $L_{\text{max}} - \widehat{L}_{\text{max}} = \lceil \frac{L_{\text{max}}}{L_{\text{min}}} \rceil (\alpha + f(n) + 2)$ symbols in any L_{max} -segment of \mathbf{z} belonging to either index or marker.

Finally, we remove from \mathbf{y}' coordinates corresponding to '1's inserted into \mathbf{w}^* ; thus, we obtain a string $\widehat{\mathbf{y}} \in \mathcal{B}_{\text{BE}}^{\widehat{L}_{\text{max}}}(\mathbf{y}^* \circ \mathbf{w}; t)$. A decoder for \mathcal{C}_{BEC} may be invoked on $\widehat{\mathbf{y}}$ to retrieve $\mathbf{y}^* = \mathbf{y}_0 \circ \dots \circ \mathbf{y}_{K-\rho-1}$, and consequently \mathbf{x} is obtained by applying the RLL decoder to each \mathbf{y}_i , $i \in [K - \rho]$.

To conclude the proof we observe that the asserted redundancy follows by definition, as precisely $\rho m(N)$ less information symbols are input at the encoder, in comparison to Construction A. \blacksquare

Next, we note that an extension to the case $k > 1$, i.e., to t -deletion multi-strand torn-paper codes, is again straightforward.

Corollary 24 Amending Construction C, one constructs a t -deletion multi-strand torn-paper code $\mathcal{C}_{\text{del}}(n, k)$ with redundancy

$$\text{red}(\mathcal{C}_{\text{del}}(n, k)) = \text{red}(\mathcal{C}_A(n, k)) + m(N) \left\lceil \frac{1}{N} \left\lceil \rho_{\text{BEC}} \cdot \frac{f(n)}{f(n)-1} \right\rceil \right\rceil.$$

Proof: Here, an information string $\mathbf{x} \in \Sigma^{(kK-\rho)m(N)}$ is encoded with E_m^{RLL} into \mathbf{y}^* , and \mathbf{w}^* is obtained utilizing a systematic BEC encoder on strings in $\Sigma^{(kK-\rho)N}$. It is segmented into $\mathbf{y}_{kK-\rho}, \dots, \mathbf{y}_{kK} \in \Sigma^N$; again, observing $\rho N \geq \rho_{\text{BEC}}$ assures that this is possible. Then, each K segments \mathbf{y}_j are encoded, in order, with the remaining steps of Algorithm 1, where again $I \triangleq \lceil \log(k \lceil n/L_{\text{min}} \rceil) \rceil$ and $L_{\text{min}} = \lceil a \log(nk) \rceil$, and indices are utilized by each operation in succession. It is straightforward that the proof of Theorem 23 can be followed to show that $\mathcal{C}_{\text{del}}(n, k)$ is a t -deletion multi-strand torn-paper code, with the above redundancy. \blacksquare

Before concluding the section, we discuss the cases of $t \in \{1, 2\}$, in which more is known on the construction of BEC codes.

For $t = 1$, we use a systematic interleaving parity BEC code as the code \mathcal{C}_{BEC} . Namely, the redundancy string $\mathbf{w} = \text{ENC}_{\text{BEC}}(\mathbf{y}^*)$ is of length $\rho_{\text{BEC}} = \widehat{L}_{\text{max}}$, and

$$w_i \triangleq \sum_{k \in \left\lceil \left\lfloor \frac{(K-\rho)N-i}{L_{\text{max}}} \right\rfloor \right\rceil} y_{i+k\widehat{L}_{\text{max}}}^*$$

for all $i \in [\widehat{L}_{\text{max}}]$, i.e., w_i is a single parity symbol for $(y_i^*, y_{i+\widehat{L}_{\text{max}}}^*, \dots)$. Denote this code by $\mathcal{C}_{\text{del},1}$.

For $t = 2$, we state for completeness the following basic proposition which draws the connection between burst-error-correcting codes and burst-erasure-correcting codes. We note that this fact has been mentioned before in [7], for a single burst of errors.

Lemma 25 *For $0 < \ell \leq n$ and $\mathbf{x}, \mathbf{y} \in \Sigma^n$, it holds that \mathbf{x}, \mathbf{y} are confusable under t bursts of errors of lengths at most ℓ if and only if they are confusable under $2t$ bursts of erasures of lengths at most ℓ .*

Proof: Denote $\mathbf{x} = (x_j)_{j \in [n]}$, $\mathbf{y} = (y_j)_{j \in [n]}$, and $I_i \triangleq \bigcup_{j \in [t]} (k_j^{(i)} + [\ell])$, for $i = 0, 1$ and some $\{k_j^{(i)}\}_{j \in [t]} \subseteq [n]$. Assume there exist $\mathbf{e}^{(0)}, \mathbf{e}^{(1)} \in \Sigma^n$ such that $\mathbf{x} + \mathbf{e}^{(0)} = \mathbf{y} + \mathbf{e}^{(1)}$, and $\text{supp}(\mathbf{e}^{(i)}) \subseteq I_i$, $i = 1, 2$. Then, one observes that $\mathbf{x}_{[n] \setminus (I_0 \cup I_1)} = \mathbf{y}_{[n] \setminus (I_0 \cup I_1)}$.

Conversely, assume $\mathbf{x}_{[n] \setminus I} = \mathbf{y}_{[n] \setminus I}$, where $I \subseteq \bigcup_{j \in [2t]} (k_j + [\ell])$ for some $\{k_j\}_{j \in [2t]} \subseteq [n]$, and without loss of generality $\{k_j\}_{j \in [2t]}$ are increasing, and $k_j \leq k_{t+1} - \ell$ for all $j \leq t$. Let $I_i \triangleq \bigcup_{j \in (it + [t])} (k_j + [\ell])$ for $i = 1, 2$, and observe that $I_0 \cup I_1 = I$, $I_0 \cap I_1 = \emptyset$. For $i = 1, 2$ and $j \in [n]$, let

$$e_j^{(i)} \triangleq \begin{cases} (-1)^i (y_j - x_j), & j \in I_i; \\ 0, & \text{otherwise.} \end{cases}$$

Then, denoting $\mathbf{e}^{(i)} = (e_j^{(i)})_{j \in [n]}$ for $i = 1, 2$, we have $\mathbf{x} + \mathbf{e}^{(0)} = \mathbf{y} + \mathbf{e}^{(1)}$, which completes the proof. ■

A construction of 2-deletion torn-paper codes is derived from Construction C, using a BEC code for $t = 2$. Hence, by Lemma 25 one may use an \widehat{L}_{max} -burst error-correcting code. Observe that Construction C requires a systematic encoder, which is guaranteed by several prior works with redundancy at most $\log((K - \rho)N) + \widehat{L}_{\text{max}}$; see, e.g. [1], [2]. These constructions require the alphabet Σ to be a field, and are linear and cyclic, which ensures the existence of a systematic encoder. For simplicity of derivation we bound this redundancy (from above) by $\widehat{L}_{\text{max}} + \log(n)$. Let $\mathcal{C}_{\text{del},2}$ denote this code.

The next corollary summarizes these results. For convenience, denote the difference

$$\Delta \text{red}(\mathcal{C}(n)) \triangleq \text{red}(\mathcal{C}(n)) - \text{red}(\mathcal{C}_A(n)),$$

for a t -deletion torn-paper code $\mathcal{C}(n) \subseteq \Sigma^n$.

Corollary 26 *For a prime power q and all admissible values of n and $f(n)$ in Construction A, where $f(n) = \omega(1)$, $f(n) =$*

$o(\log(n))$ and with the RLL encoders of [20], [37], it holds that

$$\Delta \text{red}(\mathcal{C}_{\text{del},1}(n)) \leq \widehat{L}_{\text{max}} \cdot \frac{f(n)}{f(n) - 1},$$

$$\Delta \text{red}(\mathcal{C}_{\text{del},2}(n)) \leq (\widehat{L}_{\text{max}} + \log(n)) \cdot \frac{f(n)}{f(n) - 1}.$$

In particular, for $f(n) = (1 + o(1))\sqrt{\log(n)}$,

$$\Delta \text{red}(\mathcal{C}_{\text{del},1}(n)) \leq \widehat{L}_{\text{max}} \left(1 + \frac{1 - o(1)}{\sqrt{\log(n)}} \right),$$

$$\Delta \text{red}(\mathcal{C}_{\text{del},2}(n)) \leq (\widehat{L}_{\text{max}} + \log(n)) \left(1 + \frac{1 - o(1)}{\sqrt{\log(n)}} \right).$$

Note that if $L_{\text{max}} = o(n)$ the asymptotic rate of $\mathcal{C}_{\text{del},1}(n)$ and $\mathcal{C}_{\text{del},2}(n)$ are asymptotically equal to the rate of $\mathcal{C}_A(n)$. Thus, efficient encoding and decoding of t -deletion torn-paper codes, $t = 1, 2$, is possible at rates arbitrarily close to the optimum.

V. CONCLUSION

In this paper, we study the adversarial torn-paper channel, for which we present fundamental bounds and code constructions. We further study several extensions of this model, including multi-strand storage, substitution errors, or incomplete coverage. Importantly, our proposed constructions have linear-run-time encoders and decoders, and the resulting codes achieve asymptotically optimal rates.

We mention again that the adversarial model we assume in this work is chosen to enable analysis in the worst-case. More realistically, an adversarial channel where the average of the received segments' lengths is bounded from below might be analyzed; unfortunately, this channel turns out to be hard to analyze in the worst-case, and such analysis is left for future works. It will be remarked that by the same methods of Lemma 1, it can be shown that the capacity of such an adversarial channel is bounded from above by $1 - \frac{1}{a}$, where the lower bound on the average segment length is chosen to be $a \log(n)$. Coding for this channel appears to be more challenging; we point to the fact that an adversary is able to segment a non-vanishing fraction of the channel input into short substrings as a likely reason for that difficulty.

A naive solution, where the lower bound on the average segment length is $a \log(n)$ and $a > \frac{q}{q-1}$, is to apply Construction C with parameter a' satisfying $1 < a' < (1 - \frac{1}{q})a$; the decoder then discards any received segment shorter than $a' \log(n)$, creating at most $\frac{n}{a \log(n)}$ bursts of erasures of lengths at most $(a' - 1) \log(n)$ (in the reconstructed information sequence). To recover the information, a BEC code \mathcal{C}_{BEC} is used; since $\frac{1}{Km(N)}(a' - 1) \log(n) \frac{n}{a \log(n)} = \frac{(a' - 1)/a}{Km(N)/n} = \frac{a'}{a} + o(1) < 1 - \frac{1}{q}$, a positive-rate BEC code exists for all a' in the permissible range (since positive-rate $\frac{a'-1}{a}n$ -erasure-correcting codes exist in $\Sigma^{Km(N)}$); hence a' can be optimized according to Theorem 23, i.e., to maximize the achieved rate of $(1 + o_n(1))(1 - \frac{1}{a'}) \cdot R(\mathcal{C}_{\text{BEC}})$. (Naively, one might utilize codes correcting $\frac{n}{a \log(n)} \cdot (a' - 1) \log(n)$

erasures; by the GV bound, the achievable rate of this construction is at least $(1 + o_n(1))(1 - \frac{1}{a'}) (1 - H_q(\frac{a'}{a}))$. Alternatively, if any integer a' falls within the given range, Construction P can also be used with parameter a' , where again segments shorter than $(a' + o(1)) \log(n)$ are discarded at the decoder, and reconstructed based on a BEC code \mathcal{C}_{BEC} correcting up to $\frac{n}{a \log(n)}$ bursts of erasures of lengths at most $(1 + o(1)) \log(n)$ in $\Sigma^{n/a'}$; however, the achieved rate of this construction is similarly $(1 - \frac{1}{a'}) \cdot R(\mathcal{C}_{\text{BEC}})$, and applicable BEC codes are equivalent (i.e., they correct the same number of bursts, of length $(1 + o(1)) \log(n)$ in $\Sigma^{n/a'}$ instead of length $(a' - 1) \log(n)$ in $\Sigma^{K^m(N)} = \Sigma^{(1-1/a'+o(1))n}$).

Finally, for future research, we believe that applying our methods to a generalized channel, including multiple sources of noise concurrently, one may achieve similar results. Studying the channel under edit-errors, including insertions/deletions in addition to substitutions, is also of great interest for applications to DNA data storage.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the valuable insight and advice offered to us by the two anonymous reviewers and associate editor, which were instrumental in streamlining the presentation of this paper.

REFERENCES

- [1] K. A. S. Abdel-Ghaffar, "On the existence of optimum cyclic burst correcting codes over $\text{GF}(q)$," *IEEE Trans. on Inform. Theory*, vol. 34, no. 2, pp. 329–332, Mar. 1988.
- [2] K. A. S. Abdel-Ghaffar, R. J. McEliece, A. M. Odlyzko, and H. C. A. van Tilborg, "On the existence of optimum cyclic burst-correcting codes," *IEEE Trans. on Inform. Theory*, vol. 32, no. 6, pp. 768–775, Nov. 1986.
- [3] J. Acharya, H. Das, O. Milenkovic, A. Orlitsky, and S. Pan, "String reconstruction from substring compositions," *SIAM J. Discrete Math.*, vol. 29, no. 3, pp. 1340–1371, 2015.
- [4] F. Balado, "Capacity of DNA data embedding under substitution mutations," *IEEE Trans. on Inform. Theory*, vol. 59, no. 2, pp. 928–941, Feb. 2013.
- [5] J. Bornholt, R. Lopez, D. M. Carmean, L. Ceze, G. Seelig, and K. Strauss, "A DNA-based archival storage system," *ACM SIGPLAN Notices*, vol. 51, no. 4, pp. 637–649, Mar. 2016.
- [6] G. Bresler, M. Bresler, and D. Tse, "Optimal assembly for high throughput shotgun sequencing," *BMC Bioinformatics*, vol. 14, no. 5, p. S18, Jul. 2013.
- [7] R. T. Chien, L. R. Bahl, and D. Tang, "Correction of two erasure bursts (corresp.)," *IEEE Trans. on Inform. Theory*, vol. 15, no. 1, pp. 186–187, Jan. 1969.
- [8] C.-S. Chin, D. H. Alexander, P. Marks, A. A. Klammer, J. Drake, C. Heiner, A. Clum, A. Copeland, J. Huddleston, E. E. Eichler, S. W. Turner, and J. Korlach, "Nonhybrid, finished microbial genome assemblies from long-read SMRT sequencing data," *Nature Methods*, vol. 10, no. 6, pp. 563–569, Jun. 2013.
- [9] G. M. Church, Y. Gao, and S. Kosuri, "Next-generation digital information storage in DNA," *Science*, vol. 337, no. 6102, pp. 1628–1628, 2012.
- [10] Contributing Members, "Preserving our digital legacy: an introduction to DNA data storage," The DNA Data Storage Alliance, White Paper, Jun. 2021. [Online]. Available: <https://dnastoragealliance.org/dev/wp-content/uploads/2021/06/DNA-Data-Storage-Alliance-An-Introduction-to-DNA-Data-Storage.pdf>
- [11] N. G. de Bruijn and T. van Aardenne-Ehrenfest, "Circuits and trees in oriented linear graphs," *Simon Stevin*, vol. 28, pp. 203–217, 1951.
- [12] O. Elishco, R. Gabrys, M. Médard, and E. Yaakobi, "Repeat-free codes," *IEEE Trans. on Inform. Theory*, vol. 67, no. 9, pp. 5749–5764, Sep. 2021.
- [13] Y. Erlich and D. Zielinski, "DNA fountain enables a robust and efficient storage architecture," *Science*, vol. 355, no. 6328, pp. 950–954, Mar. 2017.
- [14] R. Gabrys and O. Milenkovic, "Unique reconstruction of coded strings from multiset substring spectra," *IEEE Trans. on Inform. Theory*, vol. 65, no. 12, pp. 7682–7696, Dec. 2019.
- [15] S. Ganguly, E. Mossel, and M. Racz, "Sequence assembly from corrupted shotgun reads," in *Proceedings of the 2016 IEEE International Symposium on Information Theory (ISIT), Barcelona, Spain*, Jul. 2016, pp. 265–269.
- [16] N. Goldman, P. Bertone, S. Chen, C. Dessimoz, E. M. LeProust, B. Sipos, and E. Birney, "Towards practical, high-capacity, low-maintenance information storage in synthesized DNA," *Nature*, vol. 494, no. 7435, pp. 77–80, Feb. 2013.
- [17] R. N. Grass, R. Heckel, M. Puddu, D. Paunescu, and W. J. Stark, "Robust chemical preservation of digital information on DNA in silica with error-correcting codes," *Angewandte Chemie International Edition*, vol. 54, no. 8, pp. 2552–2555, 2015.
- [18] R. Heckel, I. Shomorony, K. Ramchandran, and D. N. C. Tse, "Fundamental limits of DNA storage systems," in *Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT), Aachen, Germany*, Jun. 2017, pp. 3130–3134.
- [19] A. Lenz, P. H. Siegel, A. Wachter-Zeh, and E. Yaakobi, "An upper bound on the capacity of the DNA storage channel," in *Proceedings of the 2019 IEEE Information Theory Workshop (ITW), Visby, Sweden*, Aug. 2019.
- [20] M. Levy and E. Yaakobi, "Mutually uncorrelated codes for DNA storage," *IEEE Trans. on Inform. Theory*, vol. 65, no. 6, pp. 3671–3691, Jun. 2019.
- [21] N. J. Loman, J. Quick, and J. T. Simpson, "A complete bacterial genome assembled de novo using only nanopore sequencing data," *Nature Methods*, vol. 12, no. 8, pp. 733–735, Aug. 2015.
- [22] A. Motahari, K. Ramchandran, D. Tse, and N. Ma, "Optimal DNA shotgun sequencing: Noisy reads are as good as noiseless reads," in *Proceedings of the 2013 IEEE International Symposium on Information Theory (ISIT), Istanbul, Turkey*, Jul. 2013, pp. 1640–1644.
- [23] A. S. Motahari, G. Bresler, and D. N. C. Tse, "Information theory of DNA shotgun sequencing," *IEEE Trans. on Inform. Theory*, vol. 59, no. 10, pp. 6273–6289, Oct. 2013.
- [24] S. Nassirpour, I. Shomorony, and A. Vahid, "Reassembly codes for the chop-and-shuffle channel," *arXiv preprint arXiv:2201.03590*, 2022.
- [25] L. Organick, S. D. Ang, Y.-J. Chen et al., "Random access in large-scale DNA data storage," *Nature Biotechnology*, vol. 36, no. 3, pp. 242–248, Mar. 2018.
- [26] A. N. Ravi, A. Vahid, and I. Shomorony, "Capacity of the torn paper channel with lost pieces," in *Proceedings of the 2021 IEEE International Symposium on Information Theory (ISIT), Melbourne, Victoria, Australia*, Jul. 2021, pp. 1937–1942.
- [27] —, "Coded shotgun sequencing," *IEEE Journal on Selected Areas in Information Theory*, vol. 3, no. 1, pp. 147–159, Mar. 2022.
- [28] S. L. Salzberg, "Mind the gaps," *Nature Methods*, vol. 7, no. 2, pp. 105–106, Feb. 2010.
- [29] I. Shomorony, T. Courtade, and D. Tse, "Do read errors matter for genome assembly?" in *Proceedings of the 2015 IEEE International Symposium on Information Theory (ISIT), Hong Kong, China*, Jun. 2015, pp. 919–923.
- [30] I. Shomorony and R. Heckel, "Capacity results for the noisy shuffling channel," in *Proceedings of the 2019 IEEE International Symposium on Information Theory (ISIT), Paris, France*, Jul. 2019, pp. 762–766.
- [31] I. Shomorony, G. M. Kamath, F. Xia, T. A. Courtade, and D. N. Tse, "Partial DNA assembly: A rate-distortion perspective," in *Proceedings of the 2016 IEEE International Symposium on Information Theory (ISIT), Barcelona, Spain*, Jul. 2016, pp. 1799–1803.
- [32] I. Shomorony and A. Vahid, "Torn-paper coding," *IEEE Trans. on Inform. Theory*, vol. 67, no. 12, pp. 7904–7913, Dec. 2021.
- [33] J. Spencer, "Asymptotic lower bounds for Ramsey functions," *Discrete Mathematics*, vol. 20, pp. 69–76, 1977.
- [34] R. R. Varshamov and G. M. Tenengolts, "Code correcting single asymmetric errors (in Russian)," *Automatika i Telemekhanika*, vol. 26, no. 2, pp. 288–292, 1965.
- [35] N. Weinberger and N. Merhav, "The DNA storage channel: Capacity and error probability bounds," *IEEE Trans. on Inform. Theory*, vol. 68, no. 9, pp. 5657–5700, Sep. 2022.
- [36] P. C. Wong, K.-k. Wong, and H. Foote, "Organic data memory using the DNA approach," *Communications of the ACM*, vol. 46, no. 1, pp. 95–98, Jan. 2003.

- [37] Y. Yehezkeally, D. Bar-Lev, S. Marcovich, and E. Yaakobi, “Generalized unique reconstruction from substrings,” *IEEE Trans. on Inform. Theory*, 2023.
- [38] Y. Yehezkeally and N. Polyanski, “On codes for the noisy substring channel,” in *Proceedings of the 2021 IEEE International Symposium on Information Theory (ISIT), Melbourne, Victoria, Australia*, Jul. 2021, pp. 1700–1705.

Daniella Bar-Lev (S’20) received the B.Sc. degree in computer science and the B.Sc. degree in mathematics in 2019, and the M.Sc. degree in computer science in 2021, from the Technion—Israel Institute of Technology, Haifa, Israel, where she is currently pursuing the Ph.D. degree with the Computer Science Department. Her research interests include algorithms, discrete mathematics, coding theory, and DNA storage.

Sagi Marcovich (S’20) received the B.Sc. degree in software engineering and the M.Sc. degree in computer science from the Technion—Israel Institute of Technology, Haifa, Israel, in 2016 and 2021, respectively, where he is currently pursuing the Ph.D. degree with the Computer Science Department. His research interests include algorithms, information theory, and coding theory with applications to DNA based storage.

Eitan Yaakobi (S’07–M’12–SM’17) received the B.A. degree in computer science and the B.A. degree in mathematics in 2005, and the M.Sc. degree in computer science in 2007, from the Technion—Israel Institute of Technology, Haifa, Israel, and the Ph.D. degree in electrical engineering in 2011, from the University of California at San Diego, San Diego, CA, USA. From 2011 to 2013, he was a Post-Doctoral Researcher with the Department of Electrical Engineering, California Institute of Technology, and the Center for Memory and Recording Research, University of California at San Diego. Since 2016, he has been with the Center for Memory and Recording Research, University of California at San Diego. Since 2018, he has been with the Institute of Advanced Studies, Technical University of Munich, where he holds a four-year Hans Fischer Fellowship, funded by the German Excellence Initiative and the EU 7th Framework Program. He is currently an Associate Professor with the Computer Science Department, Technion—Israel Institute of Technology. He also holds a courtesy appointment with the Electrical and Computer Engineering (ECE) Department, Technion—Israel Institute of Technology. His research interests include information and coding theory with applications to non-volatile memories, associative memories, DNA storage, data storage and retrieval, and private information retrieval. He was a recipient of several grants, including the ERC Consolidator Grant. He received the Marconi Society Young Scholar in 2009 and the Intel Ph.D. Fellowship during 2010-2011. Since 2020, he has been serving as an Associate Editor for Coding and Decoding for the IEEE TRANSACTIONS ON INFORMATION THEORY.

Yonatan Yehezkeally (S’12–M’20) received the B.Sc. degree (*cum laude*) in mathematics and the M.Sc. (*summa cum laude*) and Ph.D. degrees in electrical and computer engineering from Ben-Gurion University of the Negev, Beer-Sheva, Israel, in 2013, 2017, and 2020 respectively. He is currently a Carl Friedrich von Siemens Post-Doctoral Research Fellow of the Alexander von Humboldt Foundation, with the Associate Professorship of Coding and Cryptography (Prof. Wachter-Zeh), School of Computation, Information and Technology, Technical University of Munich. His research interests include coding for novel storage media, with a focus on DNA-based storage and nascent sequencing technologies, as well as combinatorial structures and finite group theory.