

Capacity Results for the Noisy Shuffling Channel

Ilan Shomorony
 University of Illinois at Urbana-Champaign
 ilans@illinois.edu

Reinhard Heckel
 Rice University
 rh43@rice.edu

Abstract—Motivated by DNA-based storage, we study the *noisy shuffling channel*, which can be seen as the concatenation of a standard noisy channel (such as the BSC) and a shuffling channel, which breaks the data block into small pieces and shuffles them. This channel models a DNA storage system, by capturing two of its key aspects: (1) the data is written onto many short DNA molecules that are stored in an unordered way and (2) the molecules are corrupted by noise at synthesis, sequencing, and during storage. For the BSC-shuffling channel we characterize the capacity exactly (for a large set of parameters), and show that a simple index-based coding scheme is optimal.

I. INTRODUCTION

Due to its longevity and enormous information density, and thanks to rapid advances in technologies for writing (synthesis) and reading (sequencing), DNA is on track to become an attractive medium for archival data storage. Computer scientists and engineers have dreamed of harnessing DNA’s storage capabilities already in the 60s [1, 2], and in recent years this idea has developed into an active field of research: In 2012 and 2013 groups lead by Church [3] and Goldman [4] independently stored about a megabyte of data in DNA. In 2015 Grass et al. [5] demonstrated that millenia long storage times are possible by protecting the data both physically and information-theoretically, and designed a robust DNA data storage scheme using modern error correcting codes. Later, in the same year, Yazdi et al [6] showed how to selectively access files, and in 2017, Erlich and Zielinski [7] demonstrated that practical DNA storage can achieve very high information densities. In 2018, Organick et al. [8] scaled up these techniques and stored about 200 megabytes of data.

DNA is a long molecule made up of four nucleotides (Adenine, Cytosine, Guanine, and Thymine) and, for storage purposes, can be viewed as a string over a four-letter alphabet. However, due to technological constraints, it is difficult and inefficient to synthesize long strands of DNA. Thus, data is stored on short DNA molecules which are preserved in a pool of DNA and cannot be spatially ordered. Indeed, all the aforementioned works did store information on molecules of no longer than a few hundred nucleotides.

Given these constraints, the simplest mathematical model for a DNA storage system is a *shuffling channel*: data is encoded into M DNA molecules (i.e., strings), each of length L , and the output of the channel consists of a random shuffle of the M sequences.

In our previous work [9], we studied the capacity of such a shuffling channel with one additional feature: random sampling of the shuffled sequences. In a DNA storage system, the information is accessed via shotgun sequencing, which

in essence corresponds to randomly sampling and reading molecules from the DNA pool. In order to formally study the capacity of this system, we considered an asymptotic regime in which M is the number of sequences/molecules stored, $N = cM$ sequences are sampled, the length of each sequence is small (specifically $L = \beta \log M$ for a fixed $\beta > 0$, and throughout \log is with respect to base 2), and $M \rightarrow \infty$. For this setting, and assuming an alphabet of size two for simplicity, the storage capacity, defined as the maximum number of bits that can be reliably stored per symbol (the total number of symbols is ML) was shown to be

$$(1 - e^{-c})(1 - 1/\beta), \quad (1)$$

provided that $\beta > 1$. If $\beta < 1$, no positive rate is achievable.

Surprisingly, the capacity of this shuffle-and-sampling channel is achieved by a simple scheme where each sequence is given a unique index (or header), and an erasure outer code is used to protect against missing molecules in the sampling step. The factor $1 - e^{-c}$ can be understood as the loss due to unseen molecules, and the factor $1 - 1/\beta$ corresponds to the loss due to shuffling or the lack of ordering in the molecules.

In practice, however, the DNA sequences are additionally corrupted by substitutions, deletions, and insertion errors (which occur both at synthesis, sequencing, and during the storage period). For that reason, several recent works have studied the design of error-correcting codes tailored to the error profiles of DNA storage systems [6, 10–13]. However, generalizing the capacity expression in (1) to a setting in which, in addition to random shuffling and random sampling, the molecules are corrupted by noise is quite nontrivial.

In this work we make progress on this front by studying the capacity of a noisy shuffling channel. We drop the random sampling constraint considered in [9], and focus on a setting in which the M sequences are first corrupted by noise and then




Fig. 1. Noisy (BSC) shuffling channel.

randomly shuffled. Since in all recently proposed systems [3–8], substitution errors are more frequent than insertion and deletion errors [14], we focus on substitution noise.

Contributions: We study the capacity of the noisy shuffling channel illustrated in Figure 1, where the input/output sequences are binary, and the noisy channel is a Binary Symmetric Channel (BSC) with crossover probability p . Our main result establishes that, for a large set of the parameters β and p , the capacity of the BSC shuffling channel is

$$C = 1 - H(p) - 1/\beta. \quad (2)$$

Since the capacity of a BSC channel is $C_{\text{BSC}} = 1 - H(p)$, this result implies that the shuffling reduces the BSC capacity by $1/\beta$ bits.

The capacity-achieving scheme consists of encoding each length- L string as a length- L codeword from a BSC capacity-achieving code, and using $\log M$ out of its LC_{BSC} information bits (prior to encoding) to store an index. Hence, at least asymptotically, there is no rate gain in coding across the M strings, and index-based approaches, which were used in several practical implementations [5–8] are optimal.

Outline: The paper is organized as follows. After formalizing the problem setting in Section II, we state our main result and describe the achievability in Section III. Section IV is dedicated to the converse. We conclude by discussing extensions and future directions in Section V.

Related literature: Motivated by DNA-based storage, a few recent works have considered the problem of coding across an unordered set of strings [15–18]. The setting studied in all these works bears similarities with the one in this paper, but they focus on providing explicit code constructions, as opposed to characterizing the channel capacity, as we do here.

II. PROBLEM SETTING

The BSC shuffling channel, illustrated in Figure 1, can be seen as the concatenation of a BSC with crossover probability p and a shuffling channel, which randomly reorders the binary strings. Formally, we will view the input to the channel as a length- ML binary string

$$X^{ML} = [X_1^L, X_2^L, \dots, X_M^L]$$

(i.e., M strings of length L concatenated to form a single string of length ML). Similarly,

$$Y^{ML} = [Y_1^L, Y_2^L, \dots, Y_M^L]$$

is the corresponding output. We let $Z^{ML} = [Z_1^L, \dots, Z_M^L]$ be the random binary error pattern created by the BSC on X^{ML} and $S^M \in \{1, \dots, M\}^M$ be a vector encoding the random shuffle induced by the channel, so that

$$Y_k^L = X_{S(k)}^L \oplus Z_{S(k)}^L,$$

for $k = 1, \dots, M$, where \oplus indicates elementwise modulo 2 addition. Similar to our previous work [9], we let $L = \beta \log M$, and consider the asymptotic regime $M \rightarrow \infty$. As shown in [9], this is the proper regime for positive data rates to be achievable.

A rate R is said to be achievable if we have a sequence of codes, each with 2^{MLR} codewords, and whose error probability goes to zero as $M \rightarrow \infty$. The shuffling channel capacity C is the supremum over all achievable rates.

III. MAIN RESULT

Our main result establishes that, provided that β is large enough, treating each length- L sequence as the input to a separate BSC and encoding a unique index into each sequence is capacity optimal. More precisely, consider code for a BSC with codewords of length L and rate $R_{\text{BSC}} = 1 - H(p) - \epsilon$, for some small ϵ . Using this code, we can store LR_{BSC} information bits in each length- L string. We utilize the first $\log M$ of those bits to encode a distinct index for each of the M strings. Hence, in each string we can effectively store $LR_{\text{BSC}} - \log M$ data bits, and the rate of this scheme is

$$\frac{M(LR_{\text{BSC}} - \log M)}{ML} = R_{\text{BSC}} - 1/\beta. \quad (3)$$

Since $\epsilon > 0$ can be chosen arbitrarily small, this scheme can have a rate arbitrarily close to

$$R_{\text{index}} = 1 - H(p) - 1/\beta, \quad (4)$$

if $1 - H(p) - 1/\beta > 0$.

Strictly speaking, the simple index-based scheme described above needs to be modified using an outer code to guarantee that the rate in (3) is actually achieved. That is because there is a vanishing, but positive, probability that a given string is decoded in error, and since $M \rightarrow \infty$, the probability that at least one of the strings is decoded in error may not go to zero. When an error occurs, the index is also read in error, which may cause a ‘‘collision’’ with a properly decoded string. Nevertheless, since the error probability for each given string vanishes as $M \rightarrow \infty$, this can be fixed with an outer code across the data symbols of all the M strings, and it is straightforward to see that one can achieve arbitrarily close to the rate in (3). Hence, (4) is a lower bound to the capacity C .

On the other hand, a simple upper bound is given by

$$C \leq \min[1 - H(p), 1 - 1/\beta], \quad (5)$$

since both the capacity of the BSC, $C_{\text{BSC}} = 1 - H(p)$, and the capacity of the shuffling channel, $1 - 1/\beta$ [9], must be upper bounds to C .

Our main result improves on the upper bound in (5), and establishes that for parameters (p, β) in a certain regime, the lower bound in equation (4) is actually the capacity.

Theorem 1. *If $p \leq 0.1$ and $\beta \geq 6.4$,*

$$C = 1 - H(p) - 1/\beta. \quad (6)$$

Moreover, if $\beta \leq 1$, $C = 0$.

As we will see, the converse argument presented in the next section and the capacity expression in (6) hold for a set of parameters (p, β) larger than those specified in Theorem 1. In fact, for all (p, β) in the blue region of Figure 2, the capacity is given by (6). Hence, for $p < 0.01$ for instance, we need $\beta \geq 2.35$ for (6) to hold.




Fig. 2. Parameter regions for which we characterized the capacity. The capacity in the blue region is given by $C = 1 - H(p) - 1/\beta$, and the capacity in the red region (i.e., for $\beta < 1$) is 0. In the gray region, it is still unknown.

IV. CONVERSE

Consider a sequence of codes for the noisy shuffling channel with rate R and vanishing error probability. Let

$$X^{ML} = [X_1^L, X_2^L, \dots, X_M^L]$$

be the input to the channel when we choose one of the 2^{MLR} codewords from one such code uniformly at random, and

$$Y^{ML} = [Y_1^L, Y_2^L, \dots, Y_M^L]$$

the corresponding output. We first observe that

$$MLR = H(X^{ML}) = I(X^{ML}; Y^{ML}) - ML\epsilon_M,$$

where $\epsilon_M \rightarrow 0$ as $M \rightarrow \infty$ by Fano's inequality. Then,

$$\begin{aligned} ML(R - \epsilon_M) &= H(Y^{ML}) - H(Y^{ML}|X^{ML}) \\ &= H(Y^{ML}) - H(S^M, Z^{ML}, Y^{ML}|X^{ML}) \\ &\quad + H(S^M, Z^{ML}|X^{ML}, Y^{ML}) \\ &= H(Y^{ML}) - H(S^M, Z^{ML}, Y^{ML}|X^{ML}) \\ &\quad + H(S^M|X^{ML}, Y^{ML}). \end{aligned} \quad (7)$$

The last equality follows by noticing that, given (S^M, X^{ML}, Y^{ML}) , one can compute

$$Z_{S(k)}^L = Y_k^L \oplus X_{S(k)}^L$$

for $k = 1, \dots, M$, and thus $H(Z^{ML}|X^{ML}, Y^{ML}, S^M) = 0$. The second term in (7) can be computed as

$$\begin{aligned} H(S^M, Z^{ML}, Y^{ML}|X^{ML}) &= H(S^M) + H(Z^{ML}) + H(Y^{ML}|X^{ML}, S^M, Z^{ML}) \\ &= \log M! + MLH(p) + 0 \\ &= M \log M + MLH(p) + o(ML). \end{aligned} \quad (8)$$

In order to finish the converse, we need to jointly bound the first and third terms in equation (7). This is the most challenging part of the proof and is summarized in the following lemma:

Lemma 1. *If β and p satisfy*

$$1 - H(2p) - \frac{2}{\beta} > 0 \quad (9)$$

then it holds that

$$H(Y^{ML}) + H(S^M|X^{ML}, Y^{ML}) \leq ML + o(ML).$$

The parameter regime (p, β) for which the inequality (9) holds is the regime in which our capacity expression holds. The regime is depicted in blue in Figure 2.

Combining (7), (8) and Lemma 1, we obtain

$$ML(R - \epsilon_M) \leq ML - MLH(p) - M \log M + o(ML).$$

Dividing by ML and letting $M \rightarrow \infty$ yields the converse.

A. Intuition for Lemma 1

If we trivially bound each entropy term separately, we obtain

$$H(Y^{ML}) + H(S^M|X^{ML}, Y^{ML}) \leq ML + M \log M.$$

However, intuitively, the bound $H(S^M|X^{ML}, Y^{ML}) \leq M \log M$ is too loose. Given $X^{ML} = x^{ML}$ and $Y^{ML} = y^{ML}$, one could estimate S^M by finding for each y_i^L the x_j^L that is closest to it and guessing that $S(i) = j$. This will be a good guess if there is no other x_k^L close to x_j^L . Hence, we can picture two opposite scenarios, illustrated in Figure 3. In the first case, the strings x_1^L, \dots, x_M^L are all




Fig. 3. Two opposite scenarios for estimating S^M from (X^{ML}, Y^{ML}) .

distant from each other (in the Hamming sense). Hence, the maximum likelihood estimate of S^M given $X^{ML} = x^{ML}$ and $Y^{ML} = y^{ML}$ will be “close” to the truth and we would expect $H(S^M|X^{ML} = x^{ML}, Y^{ML} = y^{ML})$ to be small. In the second case, illustrated in Fig. 3(b), several of x_1^L, \dots, x_M^L are close to each other. So we have less information about S^M , and $H(S^M|X^{ML} = x^{ML}, Y^{ML} = y^{ML})$ may be large.

The term $H(Y^{ML})$ is maximized by making the $\{X_i^L\}$ independent and taking values uniformly in $\{0, 1\}^L$. Hence, in order for $H(Y^{ML})$ to be large, we would expect the scenario in Fig. 3(a) to be common and the scenario in Fig. 3(b) to be rare. This creates a tension between the terms $H(Y^{ML})$ and $H(S^M|X^{ML}, Y^{ML})$, which we exploit to prove Lemma 1.

B. Proof of Lemma 1

Let $\{Y_1^L, \dots, Y_M^L\}$ be the set of M strings observed at the output of the channel. In order to capture whether we are in the scenario of Figure 3(a) or (b), we let T be the largest subset of $[1 : M] \triangleq \{1, \dots, M\}$ so that, for any $i, j \in T$, $d_H(Y_i^L, Y_j^L) \geq \alpha L$, where d_H is the Hamming distance and $\alpha > 2p$. We assume that in case of ties, an arbitrary tie-breaking rule is used to define T (the actual choice will not be relevant for the proof).

We will prove that, given the conditions in Lemma 1, the following two bounds involving $\mathbf{E}|T|$ hold:

$$\begin{aligned} \text{(B1)} \quad & H(Y^{ML}) \leq M + L\mathbf{E}|T| \\ & + (M - \mathbf{E}|T|)(\log \mathbf{E}|T| + LH(\alpha)), \\ \text{(B2)} \quad & H(S^M|X^{ML}, Y^{ML}) \leq M \log M \\ & - \mathbf{E}|T| \log \mathbf{E}|T| + o(ML). \end{aligned}$$

For large $\mathbf{E}|T|$, we are intuitively more often in the scenario in Figure 3(a), while Figure 3(b) corresponds to the case where $\mathbf{E}|T|$ is small. The bounds above capture the tension between $H(Y^{ML})$ and $H(S^M|X^{ML}, Y^{ML})$ because (B2) is decreasing in $\mathbf{E}|T|$, while (B1) is increasing in $\mathbf{E}|T|$ (as long as $\beta(1 - H(\alpha)) \geq 1$). Combining (B1) and (B2),

$$\begin{aligned} & H(Y^{ML}) + H(S^M|X^{ML}, Y^{ML}) \\ & \leq L\mathbf{E}|T| + (M - \mathbf{E}|T|)(\log \mathbf{E}|T| + LH(\alpha)) \\ & \quad + M \log M - \mathbf{E}|T| \log \mathbf{E}|T| + o(ML) \\ & = \mathbf{E}|T|L(1 - H(\alpha)) + M \log \mathbf{E}|T| - 2\mathbf{E}|T| \log \mathbf{E}|T| \\ & \quad + MLH(\alpha) + M \log M + o(ML). \end{aligned} \quad (10)$$

Replacing $\mathbf{E}|T|$ with x and ignoring the terms in this upper bound that do not involve x , we have the expression

$$f(x) \triangleq x\beta(1 - H(\alpha)) \log M + M \log x - 2x \log x.$$

Let $\gamma = \beta(1 - H(\alpha))$. For $x > 0$, we have

$$\begin{aligned} f'(x) &= \frac{1}{\ln(2)} \left(\gamma \log M + \frac{M}{x} - 2 \log x - 2 \right) \\ &> \frac{1}{\ln(2)} (\gamma \log M - 2 \log x - 2) \\ &= \frac{2}{\ln(2)} \left(\log \frac{M^{\gamma/2}}{x} - 1 \right). \end{aligned}$$

Hence $f'(x) > 0$ if

$$x < \frac{1}{2} M^{\gamma/2}. \quad (11)$$

We see that, as long as $\gamma > 2$, the right-hand side of (11) is greater than M for M large enough. This means that $f(x)$ is increasing for $1 \leq x \leq M$, and must attain its maximum at $f(M)$. Notice that $x = \mathbf{E}|T| \leq M$. Therefore, (10) can be upper-bounded by setting $x = \mathbf{E}|T| = M$, which yields

$$H(Y^{ML}) + H(S^M|X^{ML}, Y^{ML}) \leq ML + o(ML).$$

Notice that this holds as long as

$$\beta(1 - H(\alpha)) > 2 \Leftrightarrow 1 - H(\alpha) - \frac{2}{\beta} > 0,$$

for some $\alpha > 2p$. From the continuity of $H(\cdot)$, such α can be found if (9) is satisfied, proving the lemma. Finally, we prove (B1) and (B2).

Proof of (B1): Since T is a deterministic function of Y^{ML} and can take at most 2^M values,

$$\begin{aligned} H(Y^{ML}) &= H(Y^{ML}, T) = H(T) + H(Y^{ML}|T) \\ &\leq M + \sum_{t \subseteq [1:M]} \Pr(T = t) H(Y^{ML}|T = t). \end{aligned} \quad (12)$$

Next we notice that, for a given t , we can write

$$\begin{aligned} H(Y^{ML}|T = t) &\leq H(\{Y_i^L : i \in t\}|T = t) \\ &\quad + H(\{Y_i^L : i \notin t\}|T = t, \{Y_i^L : i \in t\}). \end{aligned} \quad (13)$$

The first term in (13) is trivially bounded as

$$H(\{Y_i^L : i \in t\}|T = t) \leq |t|L.$$

Each of the remaining length- L strings Y_i^L with $i \notin t$ must be within a distance αL from one of the strings in $\{Y_i^L : i \in t\}$, from the definition of T . Hence, conditioned on $\{Y_i^L : i \in t\}$, each of them can only take at most $|t|B(\alpha L)$ values, where $B(\alpha L)$ is a Hamming ball of radius αL . Since $|B(\alpha L)| \leq 2^{LH(\alpha)}$ for $\alpha < 0.5$, we bound the second term in (13) as

$$\begin{aligned} & H(\{Y_i^L : i \notin t\}|T = t, \{Y_i^L : i \in t\}) \\ & \leq (M - |t|)(\log |t| + LH(\alpha)). \end{aligned}$$

Using these bounds back in (12), we obtain

$$\begin{aligned} H(Y^{ML}) &\leq M + \mathbf{E}[L|T] + (M - |T|)(\log |T| + LH(\alpha)) \\ &\leq M + L\mathbf{E}|T| + (M - \mathbf{E}|T|)(\log \mathbf{E}|T| + LH(\alpha)), \end{aligned} \quad (14)$$

where we used the fact that $(M - x) \log x$ is a concave function of x and Jensen's inequality. \square

Proof of (B2): Since T is a deterministic function of Y^{ML} ,

$$\begin{aligned} H(S^M|X^{ML}, Y^{ML}) &= H(S^M|X^{ML}, Y^{ML}, T) \\ &= \sum_{t \subseteq [1:M]} \Pr(T = t) H(S^M|X^{ML}, Y^{ML}, T = t) \\ &= \sum_{t \subseteq [1:M]} \Pr(T = t) \sum_i^M H(S(i)|X^{ML}, Y^{ML}, T = t). \end{aligned} \quad (15)$$

Next we notice that the probability that δL or more errors occur in a single length- L string, for $\delta > p$, is at most $2^{-LD(\delta||p)}$ by the Chernoff bound (where $D(\cdot||\cdot)$ is the binary KL divergence). If we let \mathcal{E}_i be the event that $d_H(X_{S(i)}^L, Y_i^L) \geq \delta L$, then we have

$$\Pr(\mathcal{E}_i) \leq 2^{-LD(\delta||p)} = M^{-\beta D(\delta||p)}.$$

The conditional entropy term in (15) is upper bounded by

$$\begin{aligned}
& H(S(i), \mathbf{1}_{\mathcal{E}_i} | X^{ML}, Y^{ML}, T = t) \\
& \leq H(\mathbf{1}_{\mathcal{E}_i}) + \Pr(\mathcal{E}_i | T = t) H(S(i) | X^{ML}, Y^{ML}, T = t, \mathcal{E}_i) \\
& \quad + \Pr(\bar{\mathcal{E}}_i | T = t) H(S(i) | X^{ML}, Y^{ML}, T = t, \bar{\mathcal{E}}_i) \\
& \leq 1 + \Pr(\mathcal{E}_i | T = t) \log M \\
& \quad + H(S(i) | X^{ML}, Y^{ML}, T = t, \bar{\mathcal{E}}_i) \tag{16}
\end{aligned}$$

The final step is to bound the conditional entropy term in (16), for the case where $i \in t$. Set $\delta = \alpha/2$. Conditioned on $\bar{\mathcal{E}}_i$, $d_H(X_{S(i)}^L, Y_i^L) < \alpha L/2$. Moreover, conditioned on $T = t$, for any $j \in t - \{i\}$, $d_H(Y_i^L, Y_j^L) \geq \alpha L$. Define the set

$$A_i = \{j : Y_j^L \text{ is the closest output string in } t \text{ to } X_j^L\}.$$

We claim that, if $i \in t$, $S(i)$ must be in A_i . To see this notice that, for any $k \in t$, $k \neq i$, we have

$$\begin{aligned}
\alpha L & \leq d_H(Y_i^L, Y_k^L) \\
& \leq d_H(X_{S(i)}^L, Y_i^L) + d_H(X_{S(i)}^L, Y_k^L) \\
& < \alpha L/2 + d_H(X_{S(i)}^L, Y_k^L),
\end{aligned}$$

implying that $d_H(X_{S(i)}^L, Y_k^L) > \alpha L/2 \geq d_H(X_{S(i)}^L, Y_i^L)$. Therefore, $S(i)$ for each output string Y_i^L with $i \in t$, can take at most $|A_i|$ values. Hence we have

$$\begin{aligned}
& \sum_i^M H(S(i) | X^{ML}, Y^{ML}, T = t, \bar{\mathcal{E}}_i) \\
& \leq \sum_{i \notin t} H(S(i)) + \sum_{i \in t} \log |A_i| \\
& \leq (M - |t|) \log M + \sum_{i \in t} \log |A_i| \\
& \leq (M - |t|) \log M + |t| \log(M/|t|) \\
& = M \log M - |t| \log |t|, \tag{17}
\end{aligned}$$

where the last inequality follows because $\sum_{i \in t} |A_i| = M$, and the sum is maximized by $|A_i| = M/|t|$. Combining (15), (16), and (17), we obtain

$$\begin{aligned}
& H(S^M | X^{ML}, Y^{ML}) = \\
& = \sum_i^M \sum_{t \subseteq [1:M]} \Pr(T = t) [1 + \Pr(\mathcal{E}_i | T = t) \log M] \\
& \quad + \sum_{t \subseteq [1:M]} \Pr(T = t) \sum_i^M H(S(i) | X^{ML}, Y^{ML}, T = t, \bar{\mathcal{E}}_i) \\
& = M + \log M \sum_{i=1}^M \Pr(\mathcal{E}_i) + M \log M - \mathbf{E}[|T| \log |T|] \\
& \leq M + M \log M (1 + M^{-\beta D(\delta|p)}) - \mathbf{E}[|T| \log \mathbf{E}|T|]
\end{aligned}$$

where we used Jensen's inequality in the last step. Since $M^{-\beta D(\delta|p)} \rightarrow 0$ as $M \rightarrow \infty$, $M^{-\beta D(\delta|p)} M \log M = o(ML)$, concluding the proof. \square

V. DISCUSSION AND EXTENSIONS

While the parameter regime of (p, β) in the blue region of Figure 2 is practically the most relevant one, an interesting question for future work is whether expression (6) is still the capacity of the BSC-shuffling channel if β and p do not satisfy (9) (i.e., the gray region in Figure 2). Notice that this is a high-noise, short-block regime, and it is reasonable to imagine that coding across the different sequences can be helpful and an index-based approach might not be optimal.

A. General symmetric channels

If we view the capacity expression in (6) as $C_{\text{BSC}} - 1/\beta$, it is natural to wonder whether for a different noisy channel with capacity C_{noisy} , the corresponding noisy shuffling channel would have capacity $C_{\text{noisy}} - 1/\beta$. As it turns out, the converse proof in Section IV can be extended to the class of *symmetric* discrete memoryless channels (as described in [19, ch. 7.2]).

Consider a noisy shuffling channel formed by a symmetric discrete memoryless channel (SDMC) with output alphabet \mathcal{Y} , followed by a shuffling channel. We focus on the same asymptotic regime from Section II. Then we have:

Theorem 2. *If β is large enough, the capacity of the SDMC-shuffling channel is given by*

$$C = C_{\text{SDMC}} - 1/\beta. \tag{18}$$

Moreover, if $\beta \leq \log |\mathcal{Y}|$, $C = 0$.

For symmetric channels, capacity is achieved by making the distribution of the output uniform, which allows an analogous result to Lemma 1 to be obtained. Notice that how large β needs to be depends on the specific channel transition matrix.

VI. CONCLUDING REMARKS

In this paper we took steps towards the understanding of the fundamental limits of DNA-based storage systems. We proposed a simple model capturing the fact that molecules are stored in an unordered fashion, are short, and are corrupted by individual base errors. Our results show that a simple index-based coding scheme is asymptotically optimal.

While the model captures (moderate) substitution errors which are the prevalent error source on a nucleotide level of current DNA storage systems, the current generation of systems relies on *low-error* synthesis and sequencing technologies that are relatively expensive and limited in speed. A key idea towards developing the next-generation of DNA storage systems is to employ *high-error*, but cheaper and faster synthesis and sequencing technologies such as light-directed maskless synthesis of DNA and nanopore sequencing. Such systems induce a significant amount of insertion and deletion errors. Thus, an important area of further investigation is to understand the capacity of channels which introduce deletions and insertions as well.

In addition, in practice, one can introduce physical redundancy (i.e., duplicate molecules) in storage and redundancy in the DNA sequencing. Thus, an interesting question is to understand how to efficiently code for channels where we have several unordered, independently perturbed copies of each sequence.

REFERENCES

- [1] M. S. Neiman. Some fundamental issues of microminiaturization. *Radiotekhnika*, 1(1):3–12, 1964.
- [2] E. B. Baum. Building an associative memory vastly larger than the brain. *Science*, 268(5210):583–585, 1995.
- [3] G. M. Church, Y. Gao, and S. Kosuri. Next-generation digital information storage in DNA. *Science*, 337(6102):1628–1628, 2012.
- [4] 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*, 494(7435):77–80, 2013.
- [5] R. 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*, 54(8):2552–2555, 2015.
- [6] H. T. Yazdi, Y. Yuan, J. Ma, H. Zhao, and O. Milenkovic. A rewritable, random-access DNA-based storage system. *Sci. Rep.*, 5:14138, 2015.
- [7] Y. Erlich and D. Zielinski. Dna fountain enables a robust and efficient storage architecture. *Science*, 2017.
- [8] Lee Organick, Siena Dumas Ang, Yuan-Jyue Chen, Randolph Lopez, Sergey Yekhanin, Konstantin Makarychev, Miklos Z. Racz, Govinda Kamath, Parikshit Gopalan, Bichlien Nguyen, and et al. Random access in large-scale DNA data storage. *Nature Biotechnology*, 2018.
- [9] R. Heckel, I. Shomorony, K. Ramchandran, and D. N. C. Tse. Fundamental limits of dna storage systems. In *IEEE International Symposium on Information Theory (ISIT)*, pages 3130–3134, 2017.
- [10] H. M. Kiah, G. J. Puleo, and O. Milenkovic. Codes for DNA sequence profiles. *IEEE Trans. on Information Theory*, 62(6):3125–3146, 2016.
- [11] R. Gabrys, H. M. Kiah, and O. Milenkovic. Asymmetric Lee distance codes: New bounds and constructions. In *2015 IEEE Information Theory Workshop (ITW)*, pages 1–5, 2015.
- [12] F. Sala, R. Gabrys, C. Schoeny, and L. Dolecek. Exact reconstruction from insertions in synchronization codes. *IEEE Transactions on Information Theory*, 2017.
- [13] Mladen Kovačević and Vincent YF Tan. Asymptotically optimal codes correcting fixed-length duplication errors in dna storage systems. *IEEE Communications Letters*, 22(11):2194–2197, 2018.
- [14] R. Heckel, G. Mikutis, and R. N. Grass. A Characterization of the DNA Data Storage Channel. *arXiv:1803.03322*, 2018.
- [15] Mladen Kovačević and Vincent YF Tan. Codes in the space of multisetscoding for permutation channels with impairments. *IEEE Transactions on Information Theory*, 64(7):5156–5169, 2018.
- [16] Wentu Song and Kui Cai. Sequence-subset distance and coding for error control in dna data storage. *arXiv preprint arXiv:1809.05821*, 2018.
- [17] Andreas Lenz, Paul H Siegel, Antonia Wachter-zeh, and Eitan Yaakobi. Anchor-Based Correction of Substitutions in Indexed Sets. (291763).
- [18] Andreas Lenz, Paul H Siegel, Antonia Wachter-Zeh, and Eitan Yaakobi. Coding over sets for dna storage. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 2411–2415. IEEE, 2018.
- [19] T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley Series in Telecommunications and Signal Processing, 2nd edition, 2006.