

Polar Codes for the Deletion Channel: Weak and Strong Polarization

Ido Tal, *Senior Member, IEEE*, Henry D. Pfister, *Senior Member, IEEE*,
Arman Fazeli, *Member, IEEE*, Alexander Vardy, *Fellow, IEEE*

Abstract—This paper presents the first proof of polarization for the deletion channel with a constant deletion rate and a regular hidden-Markov input distribution. A key part of this work involves representing the deletion channel using a trellis and describing the plus and minus polar-decoding operations on that trellis. In particular, the plus and minus operations can be seen as combining adjacent trellis stages to yield a new trellis with half as many stages. Using this viewpoint, we prove a weak polarization theorem for standard polar codes on the deletion channel. To achieve strong polarization, we modify this scheme by adding guard bands of repeated zeros between various parts of the codeword. This gives a scheme whose rate approaches the mutual information and whose probability of error decays exponentially in the cube-root of the block length. We conclude by showing that this scheme can achieve capacity on the deletion channel by proving that the capacity of the deletion channel can be achieved by a sequence of regular hidden-Markov input distributions.

Index Terms—Polar codes, deletion channel, fast polarization, channels with memory, Markov processes

I. INTRODUCTION

IN many communications systems, symbol-timing errors may result in insertion and deletion errors. For example, a *deletion channel* with constant deletion rate maps a length- N input string to a substring using an i.i.d. process that deletes each input symbol with probability δ . These types of channels were first studied in the 1960s [1], [2] and modern coding techniques were first applied to them in [3]. Over the past 15 years, numerical bounds on the capacity of the deletion channel have been significantly improved but a closed-form expression for the capacity remains elusive [4]–[11]. Recently, polar codes were applied to the deletion channel in a series

of papers, but the question of polarization for non-vanishing deletion rates remained open [12]–[15]. In this work, we show that polar codes can be used to efficiently approach the mutual information rate between a regular (i.e., finite-state, irreducible, and aperiodic) hidden-Markov input process and the output of the deletion channel with constant deletion rate.

In [12], a polar code is designed for the binary erasure channel (BEC) and evaluated on a BEC that also introduces a single deletion. An inner cyclic-redundancy check (CRC) code is used and decoding is performed by running the successive cancellation list (SCL) decoder [16] exhaustively over all possible deletion locations. The idea is to insert a single erasure symbol into the received sequence at all possible locations in order to convert a deletion at that location into an erasure. The results show one can recover a single deletion in this setting. An extension to d deletions is also discussed but the decoding complexity grows faster than N^{d+1} , where N is the code length and d is the number of deletions.

In [13], a low-complexity decoder is proposed for the case of d deletions. Its complexity, for a length- N polar code, is roughly $d^3 N \log N$ when d deletions occur¹. The paper also presents simulation results for polar codes with lengths ranging from 256 to 2048 on two deletion channels. The first channel has a fixed deletion rate of 0.002 and the second introduces exactly 4 deletions. Based on their results, the authors of [13] conjecture that polarization occurs when $N \rightarrow \infty$ while the total number of deletions, d , is fixed.

The final papers [14], [15] in this series extend the previous results by proving that weak polarization occurs when $N \rightarrow \infty$ and $d = o(N)$. While this result is quite interesting, its proof does not extend to the case of constant deletion rate. For the case where $N \rightarrow \infty$ with d fixed, these papers also show strong polarization for the deletion channel and weak polarization for the cascade of the deletion channel and a discrete memoryless channel (DMC).

In this paper, we combine the well-known trellis representation for channels with synchronization errors [3] with low-complexity successive-cancellation (SC) trellis decoding for channels with memory [17], [18]. In particular, [3] describes how the joint input-output probability of the deletion channel (and other synchronization-error channels) can be represented using a trellis. This is closely related to fast algorithms for the edit distance between strings based on dynamic programming [19]. The main advantage of the trellis perspective is that it naturally generalizes to other channels

This paper was presented in part in the at the International Symposium on Information Theory (ISIT'2019). DOI:10.1109/ISIT.2019.8849705

The work of A. Fazeli and A. Vardy was supported in part by the National Science Foundation (NSF) under Grants CCF-1405119 and CCF-1719139.

The work of I. Tal and A. Vardy was supported in part by the United-States Israel Binational Science Foundation (BSF) under Grant No. 2018218.

The work of H. D. Pfister was supported in part by the National Science Foundation (NSF) under Grant No. 1718494.

Ido Tal is with the Viterbi Department of Electrical and Computer Engineering, Technion, Haifa 320003, Israel (email: idotal@ee.technion.ac.il).

Henry D. Pfister is with the Department of Electrical and Computer Engineering and the Department of Mathematics at Duke University, Durham, NC 27708 USA (email: henry.pfister@duke.edu).

Arman Fazeli is with Apple Inc., Los Altos, CA 94022 USA (email: afazelic@ucsd.edu).

Alexander Vardy is with the Department of Electrical and Computer Engineering, the University of California at San Diego, San Diego CA 92093, USA (email: vardy@ece.ucsd.edu).

Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

¹In [13], this complexity is misstated as $O(d^2 N \log N)$.

with synchronization errors (e.g., with insertions, deletions, and errors). The papers [17], [18] describe how the plus and minus polar-decoding operations can be efficiently applied to a channel whose input-output mapping is represented by a trellis. Putting these ideas together defines a low-complexity SC decoder for polar codes on the deletion channel that is essentially equivalent to the decoder defined in [13].

Building on previous proofs of polarization for channels with memory [20], [21], this paper proves weak and strong polarization for the deletion channel. In order to prove strong polarization, guard bands of ‘0’ symbols are embedded in the codewords of Arkan’s standard polar codes. Effectively, these guard bands allow the decoder to work on independent blocks and enable our proof of strong polarization.

The primary results of this research are summarized in Theorem 1. Conceptually, it provides a polynomial-time method to achieve the mutual information rate between a fixed regular hidden-Markov input process and the binary deletion channel.

Theorem 1. *Fix a regular hidden-Markov input process and a parameter $\nu \in (0, 1/3]$. The rate of our coding scheme approaches the mutual information rate between the input process and the binary deletion channel output. The encoding and decoding complexities of our scheme are $O(\Lambda \log \Lambda)$ and $O(\Lambda^{1+3\nu})$, respectively, where Λ is the blocklength. For any $0 < \nu' < \nu$ and sufficiently large blocklength Λ , the probability of decoding error is at most $2^{-\Lambda^{\nu'}}$.*

The family of allowed input distributions is defined in Subsection II-D and the structure of the codeword is defined in Section VII-A. Its proof can be found in Section VII. As we will see, the parameter ν controls the number of polar transforms that are treated via trellises (the weak polarization stage) versus the number of polar transform in which no trellises are employed (the strong polarization stage). As can be seen above, a larger ν implies a lower probability of error, but a higher decoding complexity. While the theorem is stated for a fixed input process, we note that the encoding and decoding complexities scale cubically with the number of states in the input process.

Theorem 2 establishes a sequence of regular hidden-Markov input processes whose mutual information rates approach the deletion channel capacity.

Theorem 2. *Let C be the capacity of the binary deletion channel with deletion probability δ . For any $\epsilon > 0$, there is a regular hidden-Markov input process whose mutual information rate on the binary deletion channel output is at least $C - \epsilon$.*

Together, the two theorems imply that the first scheme can be used to achieve capacity on the binary deletion channel. We should note, however, that we do not provide an efficient method to optimize the input distribution or to bound its complexity in terms of the gap to capacity. Also, Theorem 2 is weaker than a recent result by Li and Tan which proves the capacity can be approached by a sequence of finite-order Markov input distributions that are both irreducible and aperiodic [22]. Both results are predated by an earlier proof of Dobrushin that shows that a sequence of periodic finite-

state Markov input distributions can approach capacity on the deletion channel [2].

Here is an outline of the structure of this paper. Section II sets up the basic notation and definitions used in this paper. Section III defines the concept of a trellis and shows how it can be used to compactly represent various deletion patterns and their corresponding probabilities. In Section IV, we describe how plus and minus polarization operations are applied to trellises to yield new trellises. This provides a more detailed description of the SC trellis decoding method introduced in [17]. It is our hope that all sections up to and including Section IV will be accessible to practitioners who are primarily interested in the implementation details. Section V discusses information rates and Section VI proves that, in our setting, weak polarization occurs. Section VII focuses on strong polarization. The practitioner is advised to read Section VII-A which defines the structure and operation of an encoder with guard bands. The proof of the main theorem is presented in Section VII.

II. BACKGROUND

A. Notation

The natural numbers are denoted by $\mathbb{N} \triangleq \{1, 2, \dots\}$. We also define $[m] \triangleq \{1, 2, \dots, m\}$ for $m \in \mathbb{N}$. Let \mathcal{X} denote a finite set (e.g., the input alphabet of a channel). In this paper, we fix $\mathcal{X} = \{0, 1\}$ as the binary alphabet. Extensions to non-binary alphabets are straightforward, see for example [23, Chapter 3] and [21, Appendix A]. Let $\mathbf{x} = (x_1, \dots, x_N) \in \mathcal{X}^N$ be a vector of length $N = 2^n$. We use $[statement]$ to denote the Iverson bracket which evaluates to 1 if $statement$ is true and 0 otherwise. The concatenation of vectors $\mathbf{y} \in \mathcal{X}^{N_1}$ and $\mathbf{y}' \in \mathcal{X}^{N_2}$ lives in $\mathcal{X}^{N_1+N_2}$ and is denoted by $\mathbf{y} \odot \mathbf{y}'$. The length of a vector \mathbf{y} is denoted by $|\mathbf{y}|$. Random variables will typically be denoted by uppercase letters.

In this paper, we use the standard Arkan transform presented in the seminal paper [24]. The Arkan transform of $\mathbf{x} \in \mathcal{X}^N$, $N = 2^n$, is defined recursively using length- $N/2$ binary vectors, $\mathbf{x}^{[0]}$ and $\mathbf{x}^{[1]}$:

$$\mathbf{x}^{[0]} \triangleq (x_1 \oplus x_2, x_3 \oplus x_4, \dots, x_{N-1} \oplus x_N), \quad (1)$$

$$\mathbf{x}^{[1]} \triangleq (x_2, x_4, \dots, x_N), \quad (2)$$

where \oplus denotes modulo-2 addition. Then, for any sequence $b_1, b_2, \dots, b_\lambda \in \{0, 1\}$ with $\lambda \leq n$, we extend this notation to define the vector $\mathbf{x}^{[b_1, b_2, \dots, b_\lambda]} \in \mathcal{X}^{2^{n-\lambda}}$ recursively via

$$\mathbf{x}^{[b_1, b_2, \dots, b_\lambda]} = \left(\mathbf{x}^{[b_1, b_2, \dots, b_{\lambda-1}]} \right)^{[b_\lambda]}. \quad (3)$$

Specifically, if $\lambda = n$, then the vector $\mathbf{x}^{[b_1, b_2, \dots, b_\lambda]}$ is a scalar. This scalar is denoted by $u_{i(\mathbf{b})}$, where \mathbf{b} defines the index

$$i(\mathbf{b}) \triangleq 1 + \sum_{j=1}^n b_j 2^{n-j}. \quad (4)$$

That is, for $\mathbf{b} = (b_1, b_2, \dots, b_n)$,

$$u_{i(\mathbf{b})} = \mathbf{x}^{[b_1, b_2, \dots, b_n]}. \quad (5)$$

The transformed length- N vector is given by

$$\mathbf{u} = (u_1, \dots, u_N) = \mathcal{A}_n(\mathbf{x}), \quad (6)$$

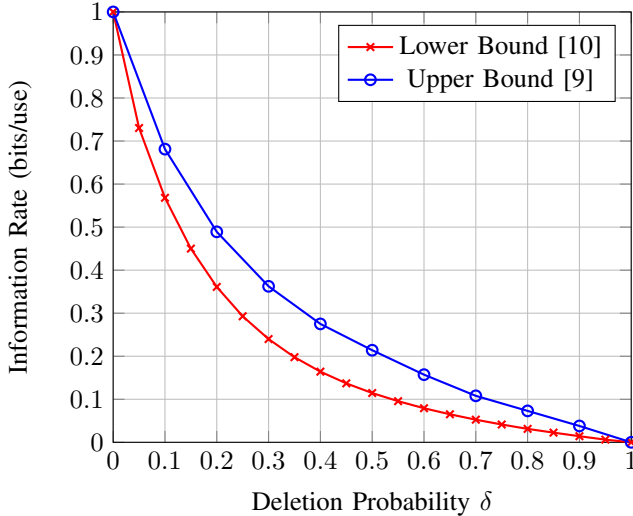


Fig. 1. Upper and lower bounds on the capacity of the binary deletion channel.

where $\mathcal{A}_n: \mathcal{X}^{2^n} \rightarrow \mathcal{X}^{2^n}$ is called the Arikan transform of order n . Its inverse is denoted \mathcal{A}_n^{-1} and satisfies $\mathcal{A}_n^{-1} = \mathcal{A}_n$.

Let $\mathbf{b} = (b_1, b_2, \dots, b_n)$ and $\mathbf{x} \in \mathcal{X}^N$ be given, where $N = 2^n$ and $i = i(\mathbf{b})$. As before, let $\mathbf{u} = \mathcal{A}(\mathbf{x})$. The vector $u_1^{i-1} = (u_1, u_2, \dots, u_{i-1})$ will play an important role later on. Thus, we now state explicitly that when $i = 0$, the vector u_1^{i-1} is taken as the empty vector.

B. Deletion Channel

Let $W(\mathbf{y}|\mathbf{x})$ denote the transition probability for N uses of the deletion channel with constant deletion rate δ . The input is denoted by $\mathbf{x} \in \mathcal{X}^N$ and the output \mathbf{y} has a random length $M = |\mathbf{y}|$ supported on $\{0, 1, \dots, N\}$. This channel is equivalent to a BEC with erasure probability δ followed by a device that removes all erasures from the output. Thus, $W(\mathbf{y}|\mathbf{x})$ equals the probability that $N - M$ deletions have occurred, which is $(1 - \delta)^M \cdot \delta^{N-M}$, multiplied by the number of distinct deletion patterns that produce \mathbf{y} from \mathbf{x} , see [4, Section 2].

The study of this channel was initiated in the 1960s [1], [2]. Modern coding techniques were first applied to this channel in [3]. Over the past 15 years, numerical bounds on the capacity of the deletion channel have been significantly improved [4]–[11]. Figure 1 shows the current best bounds on its capacity as a function of δ .

We will also consider a trimmed deletion channel whose output is given by removing all leading and trailing zeros from the output of the standard deletion channel. See Section VII for details.

C. Trellis Definition

An N -segment *trellis* \mathcal{T} is a labeled weighted directed graph $(\mathcal{V}, \mathcal{E})$. We assume that \mathcal{V} can be partitioned into $\mathcal{V}_0, \dots, \mathcal{V}_N$ so that \mathcal{V} is the union of $N + 1$ disjoint sets:

$$\mathcal{V} = \mathcal{V}_0 \cup \mathcal{V}_1 \cup \dots \cup \mathcal{V}_{N-1} \cup \mathcal{V}_N,$$

where \cup denotes a disjoint union. For channels with memory, \mathcal{V}_j represents the set of possible channel states after j channel

inputs. Similarly, the edge set \mathcal{E} is arranged into a sequence of N disjoint sets:

$$\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2 \cup \dots \cup \mathcal{E}_{N-1} \cup \mathcal{E}_N.$$

An edge in \mathcal{E}_j connects a vertex in \mathcal{V}_{j-1} to a vertex in \mathcal{V}_j . We define $\sigma(e)$ and $\tau(e)$ to be the starting and terminating vertices of edge e . Thus, for $e = u \rightarrow v$, we have $\sigma(e) = u$ and $\tau(e) = v$. Then,

$$e \in \mathcal{E}_j \text{ implies } \sigma(e) \in \mathcal{V}_{j-1} \text{ and } \tau(e) \in \mathcal{V}_j.$$

A *trellis section* comprises two adjacent sets of vertices along with the edges that connect them. That is, for $1 \leq j \leq N$, section j comprises vertex sets \mathcal{V}_{j-1} and \mathcal{V}_j , as well as edge set \mathcal{E}_j . See Fig. 2 for an example of a trellis with 4 sections.

Each edge $e \in \mathcal{E}$ has a weight $w(e) \in [0, 1]$ and a label $\ell(e) \in \mathcal{X}$. We also assume that \mathcal{V}_0 and \mathcal{V}_N have weight functions,

$$q: \mathcal{V}_0 \rightarrow [0, 1] \quad \text{and} \quad r: \mathcal{V}_N \rightarrow [0, 1],$$

that are associated with the initial and final states.

A path through a trellis is a sequence of N edges, e_1, e_2, \dots, e_N , which starts at a vertex in \mathcal{V}_0 and ends at a vertex in \mathcal{V}_N . Namely, $\sigma(e_1) \in \mathcal{V}_0$, $\tau(e_N) \in \mathcal{V}_N$, and for each $1 \leq j \leq N - 1$, we have $\tau(e_j) = \sigma(e_{j+1})$. The weight of a path through the trellis is defined as the product of the weights on each edge in the path times the weights of the initial and final vertices. Namely, the weight of the above path is

$$q(\sigma(e_1)) \cdot r(\tau(e_N)) \times \prod_{j=1}^N w(e_j).$$

Thus, an N -section trellis naturally defines a *path-sum* function $T: \mathcal{X}^N \rightarrow \mathbb{R}$, where $T(\mathbf{x})$ equals the sum of the path weights over all paths whose length- N label sequences match \mathbf{x} . That is,

$$\begin{aligned} T(\mathbf{x}) \triangleq & \sum_{\substack{e_1 \in \mathcal{E}_1, \\ \ell(e_1) = x_1}} \sum_{\substack{e_2 \in \mathcal{E}_2, \\ \ell(e_2) = x_2}} \dots \sum_{\substack{e_N \in \mathcal{E}_N, \\ \ell(e_N) = x_N}} q(\sigma(e_1)) r(\tau(e_N)) \\ & \times \prod_{j=1}^N w(e_j) \times \prod_{j=1}^{N-1} [\tau(e_j) = \sigma(e_{j+1})]. \quad (7) \end{aligned}$$

D. FAIM processes

In latter parts of this paper, for simplicity, we will often introduce key ideas by first framing them in the context of the uniform input distribution. That is, by first considering the case in which the input distribution is i.i.d. Bernoulli $1/2$. However, the uniform input distribution, or indeed any i.i.d. input distribution, is known to generally be sub-optimal with respect to the information rate between input and output, when transmitting over a deletion channel [4], [9]–[11]. Thus, we stand to benefit by considering a larger class of input distributions.

To this end, let \mathcal{S} be a given finite set. Each element of \mathcal{S} is a state of an input process. In the following² definition, we have for all $j \in \mathbb{Z}$ that $S_j \in \mathcal{S}$ and $X_j \in \mathcal{X}$.

Definition 1 (FAIM process). *A strictly stationary process (S_j, X_j) , $j \in \mathbb{Z}$ is called a finite-state, aperiodic, irreducible, Markov (FAIM) process if, for all j ,*

$$P_{S_j, X_j | S_{-\infty}^{j-1}, X_{-\infty}^{j-1}} = P_{S_j, X_j | S_{j-1}}, \quad (8)$$

is independent of j and the sequence (S_j) , $j \in \mathbb{Z}$ is a finite-state Markov chain that is stationary, irreducible, and aperiodic.

Note that by (8) we mean that for all $s_j, x_j, s_{j-1}, x_{j-1}, \dots$,

$$\begin{aligned} P(S_j = s_j, X_j = x_j | S_{j-1} = s_{j-1}, X_{j-1} = x_{j-1} \\ , S_{j-2} = s_{j-2}, X_{j-2} = x_{j-2}, \dots) \\ = P(S_j = s_j, X_j = x_j | S_{j-1} = s_{j-1}). \end{aligned}$$

For a FAIM process, consider the sequence X_j , for $j \in \mathbb{Z}$. In principle, the distribution of this sequence can be computed by marginalizing the states of the FAIM process (S_j, X_j) . Such a sequence is typically called a *hidden-Markov process*. In this paper, we sometimes add the term *regular* to emphasize that the hidden state process is a regular Markov chain.

Let us now connect the concept of a FAIM process to that of a trellis. Let a FAIM process (X_j, S_j) be given, and fix $N \geq 1$. We now define the corresponding trellis, having N stages. The vertex set is $\mathcal{V} = \mathcal{V}_0 \cup \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N$, where we define

$$\mathcal{V}_j = \{s_j : s \in \mathcal{S}\}$$

for $0 \leq j \leq N$ so that each \mathcal{V}_j contains a distinct copy of \mathcal{S} . For each $x \in \mathcal{X}$, $1 \leq j \leq N$, $\alpha_{j-1} \in \mathcal{V}_{j-1}$, and $\beta_j \in \mathcal{V}_j$, define an edge e from α_{j-1} to β_j with label $\ell(e) = x$ and weight $w(e) = P_{S_j, X_j | S_{j-1}}(\beta, x | \alpha)$. Lastly, for all $\alpha_0 \in \mathcal{V}_0$ define $q(\alpha_0) = \pi(\alpha)$, where $\pi(\alpha)$ is the stationary probability of state α in the Markov process $(S_j)_{j \in \mathbb{R}}$, and define $r(\beta_N) = 1$ for all $\beta_N \in \mathcal{V}_N$. It follows that the probability of $(X_1, X_2, \dots, X_N) = (x_1, x_2, \dots, x_N) = \mathbf{x}$ equals $T(\mathbf{x})$, where T was defined in (7).

III. TRELLIS REPRESENTATION OF JOINT PROBABILITY

We have just seen that a trellis can be used to compactly represent a hidden-Markov input distribution. In fact, a trellis can also represent more complicated distributions. Namely, we will now show how a trellis can be used to represent the *joint* distribution of a hidden-Markov input process and the channel output.

A. Trellis for uniform input

This trellis representation for the deletion channel can also be found in [3].

As previously explained, it is generally beneficial to use an input distribution with memory. However, for the sake of an easy exposition, we will first consider the simplest possible

²The definition of FAIM and FAIM-derived processes here is a specialization of the definition given in [21]. Here, we are interested in FAIM-derived (i.e., hidden-Markov) input processes. However, the input-output process of a deletion channel is neither FAIM nor FAIM-derived.

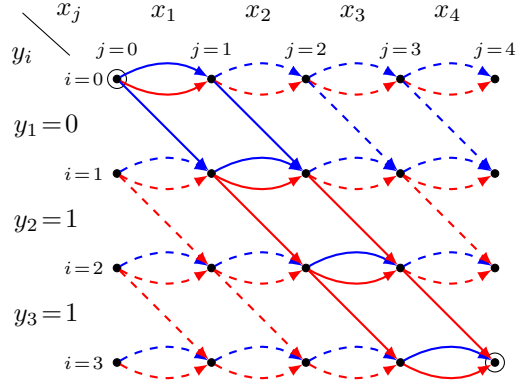


Fig. 2. A trellis for the binary deletion channel with uniform input, a codeword length of $N = 4$, and a received word $\mathbf{y} = (011)$ of length $M = 3$. Vertices are denoted $v_{i,j}$ with $0 \leq i \leq M$ and $0 \leq j \leq N$. All blue edges have label ‘0’ while all red edges have label ‘1’. The horizontal edges are weighted by the probability $\delta/2$. Diagonal edges are weighted by the probability $(1-\delta)/2$. The two circled vertices have $q(v_{0,0}) = r(v_{M,N}) = 1$, while all other vertices in \mathcal{V}_0 and \mathcal{V}_N have q and r values equal to 0, respectively. Edges that can be pruned without changing the function T in (7) are dashed.

input distribution, a uniform input distribution (i.e., i.i.d. and Bernoulli 1/2).

The trellis representation will be used on the decoder side. Thus, when building the trellis, we will have already received the output vector \mathbf{y} . Hence, the primary role of the trellis is to evaluate the probabilities associated with possible input vectors \mathbf{x} , of length N . That is, the trellis will be used to calculate the joint probability of \mathbf{x} and \mathbf{y} , denoted $P_{\mathbf{X}}(\mathbf{x}) \cdot W(\mathbf{y}|\mathbf{x})$, for \mathbf{y} fixed. Recall that $W(\mathbf{y}|\mathbf{x})$ is the deletion channel law, and in this subsection $P_{\mathbf{X}}$ is the uniform input distribution.

We will shortly define the concept of a valid path in the trellis. Each valid path will correspond to a specific transmitted \mathbf{x} and a specific deletion pattern that is compatible with the received \mathbf{y} (see Fig. 2). We term this trellis the *base trellis*, as we will ultimately construct other trellises derived from it.

Recalling our notation, we have \mathbf{x} as the unknown input vector, of known length N . The vector \mathbf{y} is the known output, having known length $M = |\mathbf{y}|$. The deletion probability is δ . The base trellis is defined as follows.

Definition 2 (Base Trellis for Uniform Input). *For N, δ, M , and $\mathbf{y} \in \mathcal{X}^M$:*

- 1) *The vertex set \mathcal{V} equals the disjoint union*

$$\mathcal{V} = \mathcal{V}_0 \cup \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N,$$

where, for $0 \leq j \leq N$,

$$\mathcal{V}_j = \{v_{i,j} : 0 \leq i \leq M\}. \quad (9)$$

- 2) *A path passing through vertex $v_{i,j}$ corresponds to the event where only i of the first j transmitted symbols were received. That is, from x_1, x_2, \dots, x_j , the channel has deleted $j - i$ symbols³.*
- 3) *Vertices $v_{i,j}$ with $0 \leq i \leq M$ and $0 \leq j < N$ each have up to three outgoing edges: two ‘horizontal’ edges, each*

³Note that we could have optimized our definition of \mathcal{V}_j . Namely, only i in the range $\max\{0, M - N + j\} \leq i \leq \min\{j, M\}$ are actually consistent with the described event (i.e., only the solid edges in Figure 2). We leave such optimization to the practitioner and settle for the simpler description in (9).

corresponding to a deletion, and one ‘diagonal’ edge, corresponding to a non-deletion.

- 4) For $0 \leq i \leq M$ and $0 \leq j < N$, there are two edges e, e' from $v_{i,j}$ to $v_{i,j+1}$. From 2) above, we deduce that these two ‘horizontal’ edges are associated with x_{j+1} being deleted by the channel. The first is associated with $x_{j+1} = 0$ and has $\ell(e) = 0$, while the second is associated with $x_{j+1} = 1$ and has $\ell(e') = 1$. Since the probability of deletion is δ , and in the uniform distribution $x_{j+1} = 0$ and $x_{j+1} = 1$ each occur with probability $1/2$, we set $w(e) = w(e') = \delta/2$.
- 5) For $0 \leq i < M$ and $0 \leq j < N$, there is a single edge e from $v_{i,j}$ to $v_{i+1,j+1}$. Recalling 2) above, we deduce that this ‘diagonal’ edge represents x_{j+1} not being deleted, and being observed as y_{i+1} . Thus, $\ell(e) = y_{i+1}$. Since the probability of sending x_{j+1} in the uniform case is $1/2$, regardless of its value, and the probability of a non-deletion is $1 - \delta$, we set $w(e) = (1 - \delta)/2$.
- 6) We set $q(v_{0,0}) = 1$. All other vertices $v \in \mathcal{V}_0$ have $q(v) = 0$. Thus, with respect to (7), we effectively force all paths to start at $v_{0,0}$. Namely, when starting a path, no symbols have yet been transmitted, and hence no symbols have yet been received.
- 7) We set $r(v_{M,N}) = 1$. All other vertices $v \in \mathcal{V}_N$ have $r(v) = 0$. Thus, with respect to (7), we effectively force all paths to end at $v_{M,N}$. That is, at the end of a path, N symbols have been transmitted, and of these, M have been received.

In line with the definitions above, let us call a path *valid* if it starts at $v_{0,0}$ and ends at $v_{M,N}$. For example, in Figure 2, valid paths are those that start at the circled vertex on the top left, end at the circled vertex on the bottom right, and hence contain only solid edges. Clearly, such a path is comprised of N edges, e_1, e_2, \dots, e_N . Denote by $\mathbf{x} = (x_1, x_2, \dots, x_N)$ the input vector corresponding to the above path, where $x_i = \ell(e_i)$. Each such \mathbf{x} is consistent with our received \mathbf{y} . Indeed, tracing the path, the type of the corresponding edge (horizontal or diagonal) shows exactly which of the x_i to delete and which to keep in order to arrive at \mathbf{y} . Also, the probability of the input sequence \mathbf{x} being transmitted and experiencing the above chain of deletion/no-deletion events is exactly equal to the product of the $w(e_i)$, times $q(v_{0,0}) \cdot r(v_{M,N}) = 1$.

From the above discussion, one has the following key lemma.

Lemma 3. *Let \mathcal{T} be a trellis as described in Definition 2. Then, for $\mathbf{x} \in \mathcal{X}^N$ and $T(\mathbf{x})$ as defined in (7), we have*

$$T(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x}) \cdot W(\mathbf{y}|\mathbf{x}),$$

where $P_{\mathbf{X}}$ is the uniform input distribution and W is the deletion channel law.

Proof: First, we observe that the weight of a trellis path equals the joint probability of (\mathbf{x}, \mathbf{y}) and the deletion pattern. Then, the claim follows from the fact that $T(\mathbf{x})$ sums the path weight over all paths through the trellis (i.e., all deletion patterns) consistent with the given (\mathbf{x}, \mathbf{y}) pair. ■

B. Trellises for hidden-Markov inputs

As explained earlier, a trellis is used on the decoding side, in order to capture the joint probability of \mathbf{x} and \mathbf{y} . We now show how such a trellis is built for the more general case in which \mathbf{x} is drawn from a regular hidden-Markov input process. Intuitively, this is done by simply ‘multiplying’ the trellis corresponding to the input distribution, as described at the end of Section II, with the trellis defined for the uniform case (with the correction that the edge weights $\delta/2$ and $(1 - \delta)/2$ are replaced by δ and $1 - \delta$, respectively). A formal definition follows.

Definition 3 (Base Trellis for Hidden-Markov Input). *For $N, \delta, M, \mathcal{S}, P_{S_j, X_j | S_{j-1}}, \pi$, and $\mathbf{y} \in \mathcal{X}^M$:*

- 1) *The vertex set \mathcal{V} equals the disjoint union*

$$\mathcal{V} = \mathcal{V}_0 \cup \mathcal{V}_1 \cup \dots \cup \mathcal{V}_N,$$

where, for $0 \leq j \leq N$,

$$\mathcal{V}_j = \{v_{s,i,j} : 0 \leq i \leq M, s \in \mathcal{S}\}. \quad (10)$$

Thus, $|\mathcal{V}_j| = (M + 1) \cdot |\mathcal{S}|$.

- 2) *A path passes through vertex $v_{s,i,j}$ if exactly i of the first j transmitted symbols are not deleted and the state of the input process is $s \in \mathcal{S}$ after the j -th input (i.e., $S_j = s$).*
- 3) *Vertices $v_{s,i,j}$ with $0 \leq i \leq M, 0 \leq j < N$, and $s \in \mathcal{S}$ each have up to $3 \cdot |\mathcal{S}|$ outgoing edges.*
- 4) *For $0 \leq i \leq M, 0 \leq j < N$, and $\alpha, \beta \in \mathcal{S}$, there are two edges e, e' from $v_{\alpha,i,j}$ to $v_{\beta,i,j+1}$. From item 2, we deduce that these two ‘horizontal’ edges are associated with x_{j+1} being deleted by the channel. The first is associated with $x_{j+1} = 0$ and has $\ell(e) = 0$, while the second is associated with $x_{j+1} = 1$ and has $\ell(e') = 1$. Recalling that by stationarity $P_{S_{j+1}, X_{j+1} | S_j} = P_{S_j, X_j | S_{j-1}}$, we set*

$$w(e) = \delta \cdot P_{S_j, X_j | S_{j-1}}(\beta, 0 | \alpha) \quad (11)$$

and

$$w(e') = \delta \cdot P_{S_j, X_j | S_{j-1}}(\beta, 1 | \alpha). \quad (12)$$

That is, the probability of a deletion, times the probability implied by the underlying FAIM distribution.

- 5) *For $0 \leq i < M, 0 \leq j < N$, and $\alpha, \beta \in \mathcal{S}$, there is a single edge e from $v_{\alpha,i,j}$ to $v_{\beta,i+1,j+1}$. Recalling item 2 above, we deduce that this ‘diagonal’ edge represents x_{j+1} being observed (i.e., not deleted) as y_{i+1} . Thus, $\ell(e) = y_{i+1}$. We set*

$$w(e) = (1 - \delta) \cdot P_{S_j, X_j | S_{j-1}}(\beta, y_{i+1} | \alpha).$$

That is, the probability of a non-deletion, times the probability implied by the underlying FAIM distribution⁴.

- 6) *For all $v_{s,0,0} \in \mathcal{V}_0$, where $s \in \mathcal{S}$, we set $q(v_{s,0,0}) = \pi(s)$. All other vertices $v \in \mathcal{V}_0$ have $q(v) = 0$. Thus, with respect to (7), we effectively force all paths to start at a vertex $v_{s,0,0}$, where $s \in \mathcal{S}$. Namely, when starting a path, no symbols have yet been transmitted,*

⁴As in the uniform case, one can reduce the trellis by restricting the index i in (10) to the range $\max\{0, M - N + j\} \leq i \leq \min\{j, M\}$ and removing any edge e with $w(e) = 0$.

and hence no symbols have yet been received. Moreover, the probability of starting the path at $v_{s,0,0}$ is $\pi(s)$, the stationary probability of s in the FAIM input process.

- 7) For all $v_{s,M,N} \in \mathcal{V}_N$, we set $r(v_{s,M,N}) = 1$. All other vertices $v \in \mathcal{V}_N$ have $r(v) = 0$. Thus, with respect to (7), we effectively force all paths to end at a vertex $v_{s,M,N}$. That is, at the end of a path, N symbols have been transmitted, and of these, M have been received.

As in the uniform case, we have the following lemma, which is easily proved.

Lemma 4. Let \mathcal{T} be a trellis as per Definition 3. Then, for $\mathbf{x} \in \mathcal{X}^N$ and $T(\mathbf{x})$ as defined in (7),

$$T(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x}) \cdot W(\mathbf{y}|\mathbf{x}),$$

where $P_{\mathbf{X}}$ is the hidden-Markov input distribution and W is the deletion channel law.

Proof: First, we observe that the weight of a trellis path equals the joint probability of (\mathbf{x}, \mathbf{y}) and the deletion pattern. Then, the claim follows from the fact that $T(\mathbf{x})$ sums the path weight over all paths through the trellis (i.e., all deletion patterns) consistent with the given (\mathbf{x}, \mathbf{y}) pair. ■

C. Trellis for the trimmed deletion channel

For reasons that will shortly become clear, we will now consider a slight variation of the deletion channel. Namely, we now define the trimmed deletion channel (TDC). A TDC is a deletion channel that, after the deletion process, trims its output of leading and trailing ‘0’ symbols. Thus, by definition, the output of a TDC is either an empty string, or a string that starts and ends with a ‘1’ symbol.

We now show how to alter Definition 3 in order to account for this variation. The change turns out to be minimal.

Definition 4 (Base Trellis for Hidden-Markov Input and TDC). For $N, \delta, M, \mathcal{S}, P_{S_j, X_j | S_{j-1}}, \pi$, and trimmed output $\mathbf{y}^* \in \mathcal{X}^M$, define the trellis \mathcal{T} as in Definition 3, but with the following changes.

- The probability of an edge e from $v_{\alpha,0,j}$ to $v_{\beta,0,j+1}$ with $\ell(e) = 0$ must be changed to $w(e) = P_{S_j, X_j | S_{j-1}}(\beta, 0 | \alpha)$. Namely, the δ factor in (11) is removed. In short, if the path is currently at vertex $v_{\alpha,0,j}$, then none of the j symbols x_1, x_2, \dots, x_j have made it to the output of the channel (they have either been deleted or trimmed). Thus, if $x_{j+1} = 0$, it will surely be either deleted, or else trimmed.
- The probability of an edge e from $v_{\alpha,M,j}$ to $v_{\beta,M,j+1}$ with $\ell(e) = 0$ must be changed to $w(e) = P_{S_j, X_j | S_{j-1}}(\beta, 0 | \alpha)$. Namely, the δ factor in (11) is removed. Note that the exact same reasoning from the previous point applies; the only difference is that now we are correcting for the trimming of the trailing ‘0’ symbols.

The result of the above altered trellis definition is the following lemma.

Lemma 5. Let \mathcal{T} be a trellis as described in Definition 4. Then, for $\mathbf{x} \in \mathcal{X}^N$ and $T(\mathbf{x})$ as defined in (7),

$$T(\mathbf{x}) = P_{\mathbf{X}}(\mathbf{x}) \cdot W^*(\mathbf{y}^*|\mathbf{x}),$$

where $P_{\mathbf{X}}$ is the hidden-Markov input distribution and W^* is the law of the TDC.

Proof: First, we observe that the weight of a trellis path equals the joint probability of $(\mathbf{x}, \mathbf{y}^*)$ and the deletion/trimming event associated with that path. Then, the claim follows from the fact that $T(\mathbf{x})$ sums the path weight over all paths through the trellis (i.e., all deletion/trimming events) consistent with the given $(\mathbf{x}, \mathbf{y}^*)$ pair. ■

IV. POLARIZATION OPERATIONS ON A TRELLIS

Polar plus and minus transforms for channels with memory were first presented in [17], [18]. Let an input distribution on \mathcal{X}^N be given, for N even. For this input distribution and a vector channel with input $\mathbf{x} \in \mathcal{X}^N$ and output \mathbf{y} , let \mathcal{T} be a trellis with N sections whose path-sum function satisfies

$$T(\mathbf{x}) = \Pr(\mathbf{Y} = \mathbf{y}, \mathbf{X} = \mathbf{x}). \quad (13)$$

A. Minus transform

For a given path-sum function $T(\mathbf{x})$, where $\mathbf{x} \in \mathcal{X}^N$, the polar *minus transform* defines a new path-sum function $T^{[0]}(\mathbf{z})$, $\mathbf{z} \in \mathcal{X}^{N/2}$. Specifically, $T^{[0]}(\mathbf{z})$ is the marginalization of $T(\mathbf{x})$ over all \mathbf{x} vectors satisfying

$$\mathbf{z} = \mathbf{x}^{[0]} = (x_1 \oplus x_2, \dots, x_{N-1} \oplus x_N).$$

That is,

$$\begin{aligned} T^{[0]}(\mathbf{z}) &\triangleq \sum_{\mathbf{x} \in \mathcal{X}^N : \mathbf{x}^{[0]} = \mathbf{z}} T(\mathbf{x}) \\ &= \sum_{\mathbf{x} \in \mathcal{X}^N} T(\mathbf{x}) \prod_{j=1}^{N/2} [x_{2j-1} \oplus x_{2j} = z_j] \\ &= \Pr(\mathbf{Y} = \mathbf{y}, \mathbf{X}^{[0]} = \mathbf{z}), \end{aligned} \quad (14)$$

where the last equality follows under the assumption of (13). Due to the local nature of this reparameterization, there is a modified trellis $\mathcal{T}^{[0]}$ with $N/2$ sections that represents the new path-sum function. Fig. 3 shows an example of the minus trellis.

Definition 5 (Minus Transform). Let $\mathcal{T} = \mathcal{T}(\mathcal{V}, \mathcal{E}, w, \ell, q, r)$ be a length- N trellis, where N is even. The trellis $\tilde{\mathcal{T}} = \tilde{\mathcal{T}}(\tilde{\mathcal{V}}, \tilde{\mathcal{E}}, \tilde{w}, \tilde{\ell}, \tilde{q}, \tilde{r}) = \mathcal{T}^{[0]}$ is defined as follows.

- The vertex set of $\tilde{\mathcal{T}}$ is

$$\tilde{\mathcal{V}} = \tilde{\mathcal{V}}_0 \cup \tilde{\mathcal{V}}_1 \cup \dots \cup \tilde{\mathcal{V}}_{N/2},$$

where

$$\tilde{\mathcal{V}}_j = \mathcal{V}_{2j}.$$

- We next define the edge set $\tilde{\mathcal{E}}$ implicitly. Consider an edge $\tilde{e} = \alpha \rightarrow \gamma \in \tilde{\mathcal{E}}$ in section j of $\tilde{\mathcal{T}}$ with label $\tilde{\ell}(\tilde{e}) = z$. Then,

$$\alpha \in \tilde{\mathcal{V}}_{j-1} = \mathcal{V}_{2j-2} \quad \text{and} \quad \gamma \in \tilde{\mathcal{V}}_j = \mathcal{V}_{2j}.$$

The weight $\tilde{w}(\tilde{e})$ of this edge equals the sum of the product of the edge weights along each two-step path $\alpha \xrightarrow{e_1} \beta \xrightarrow{e_2} \gamma$ in \mathcal{T} with $\ell(e_1) \oplus \ell(e_2) = z$. That is,

$$\tilde{w}(\tilde{e}) = \sum_{\substack{e_1 \in \mathcal{E}_{2j-1}: \\ \sigma(e_1) = \alpha}} \sum_{\substack{e_2 \in \mathcal{E}_{2j}: \\ \tau(e_2) = \gamma}} w(e_1) w(e_2)$$

$$\times [\tau(e_1) = \sigma(e_2)] \cdot [\ell(e_1) \oplus \ell(e_2) = z].$$

Edges with weight 0 may be removed from $\tilde{\mathcal{T}}$.

- The minus operation does not affect initial and final vertices and this implies that $\tilde{q}(v) = q(v)$ and $\tilde{r}(v) = r(v)$.

The following lemma states that applying a minus transform to a trellis indeed results in a trellis whose corresponding path-sum function is the minus transform of the path-sum function of the initial trellis.

Lemma 6. Let \mathcal{T} be a trellis with N sections, where N is even. Denote the minus transform of \mathcal{T} by $\mathcal{T}' = \mathcal{T}^{[0]}$ per Definition 5. Let T and T' be the path-sum functions corresponding to \mathcal{T} and \mathcal{T}' , respectively, as defined in (7). Then, T' equals $T^{[0]}$ as defined in (14).

Proof: This follows from the fact that the minus trellis is constructed by merging adjacent trellis stages and then combining paths according to their $\mathbf{x}^{[0]}$ values. Finally, the new paths are relabeled by their $\mathbf{x}^{[0]}$ values. ■

B. Plus transform

For a given path-sum function $T(\mathbf{x})$, where $\mathbf{x} \in \mathcal{X}^N$, the polar plus transform defines a new path-sum function $T^{[1]}(\mathbf{z}')$, $\mathbf{z}' \in \mathcal{X}^{N/2}$. This definition is always with respect to a vector $\mathbf{z} \in \mathcal{X}^{N/2}$, which is assumed to be fixed. Specifically, $T^{[1]}(\mathbf{z}')$ equals $T(\mathbf{x})$, where \mathbf{x} is the unique vector satisfying

$$\begin{aligned} \mathbf{z} &= \mathbf{x}^{[0]} = (x_1 \oplus x_2, \dots, x_{N-1} \oplus x_N) \quad \text{and} \\ \mathbf{z}' &= \mathbf{x}^{[1]} = (x_2, x_4, \dots, x_N). \end{aligned}$$

That is,

$$\begin{aligned} T^{[1]}(\mathbf{z}') &\triangleq T(\mathbf{x}) \Big|_{\mathbf{x}: \mathbf{x}^{[0]} = \mathbf{z}, \mathbf{x}^{[1]} = \mathbf{z}'} \quad (15) \\ &= \sum_{\mathbf{x} \in \mathcal{X}^N} T(\mathbf{x}) \prod_{j=1}^{N/2} [x_{2j-1} \oplus x_{2j} = z_j] \cdot [x_{2j} = z'_j] \\ &= \Pr(\mathbf{Y} = \mathbf{y}, \mathbf{X}^{[0]} = \mathbf{z}, \mathbf{X}^{[1]} = \mathbf{z}'), \end{aligned}$$

where the last equality follows under the assumption of (13).

As with the minus transform, there is a corresponding operation one can apply to the underlying trellis, which we now detail. Note that the plus-transform of a trellis is defined with respect to a fixed vector \mathbf{z} , which may not be specified explicitly when it is clear from the context. Fig. 3 shows an example of the plus trellis.

Definition 6 (Plus Transform). Let $\mathcal{T} = \mathcal{T}(\mathcal{V}, \mathcal{E}, w, \ell, q, r)$ be a length- N trellis, where N is even and let $\mathbf{z} \in \mathcal{X}^{N/2}$ be given. The trellis $\tilde{\mathcal{T}} = \tilde{\mathcal{T}}(\tilde{\mathcal{V}}, \tilde{\mathcal{E}}, \tilde{w}, \tilde{\ell}, \tilde{q}, \tilde{r}) = \mathcal{T}^{[1]}$ is defined as follows.

- The vertex set of $\tilde{\mathcal{T}}$ is the same as the minus trellis $\mathcal{T}^{[0]}$. This is also the case for the functions \tilde{q} and \tilde{r} .
- We next define the edge set $\tilde{\mathcal{E}}$ implicitly. Consider an edge $\tilde{e} = \alpha \rightarrow \gamma \in \tilde{\mathcal{E}}$ in section j of $\tilde{\mathcal{T}}$ with label $\tilde{\ell}(\tilde{e}) = z'$. Then,

$$\alpha \in \tilde{\mathcal{V}}_{j-1} = \mathcal{V}_{2j-2} \quad \text{and} \quad \gamma \in \tilde{\mathcal{V}}_j = \mathcal{V}_{2j}.$$

The weight $\tilde{w}(\tilde{e})$ of this edge equals the sum of the product of the edge weights along each two-step path $\alpha \xrightarrow{e_1} \beta \xrightarrow{e_2} \gamma$ in \mathcal{T} with $\ell(e_1) \oplus \ell(e_2) = z_j$ and $\ell(e_2) = z'$. That is,

$$\begin{aligned} \tilde{w}(\tilde{e}) &= \sum_{\substack{e_1 \in \mathcal{E}_{2j-1}: \\ \sigma(e_1) = \alpha}} \sum_{\substack{e_2 \in \mathcal{E}_{2j}: \\ \tau(e_2) = \gamma}} w(e_1) w(e_2) \\ &\times [\tau(e_1) = \sigma(e_2)] \cdot [\ell(e_1) \oplus z' = z_j] \cdot [\ell(e_2) = z']. \end{aligned}$$

Edges with weight 0 may be removed from $\tilde{\mathcal{T}}$.

This lemma states the key property of the plus transform.

Lemma 7. Let \mathcal{T} be a trellis with N sections where N is even, and let $\mathbf{z} \in \mathcal{X}^{N/2}$ be given. Denote the plus transform of \mathcal{T} by $\mathcal{T}' = \mathcal{T}^{[1]}$ per Definition 6. Let T and T' be the path-sum functions corresponding to \mathcal{T} and \mathcal{T}' , respectively, as defined in (7). Then, T' equals $T^{[1]}$ as defined in (15).

Proof: This follows from the fact that the plus trellis is constructed by merging adjacent trellis stages and then pruning paths that do not satisfy $\mathbf{x}^{[0]} = \mathbf{z}$. Finally, the remaining paths are relabeled with $\mathbf{x}^{[1]}$ values. ■

C. Successive cancellation decoding

As in Arıkan's seminal paper [24], the transform defined above leads to an SC decoding algorithm. In brief, given \mathbf{y} we first construct a base trellis \mathcal{T} . Then, there is a recursive decoder that, given $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda]}$, constructs $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda, 0]}$ and calls itself with that argument. When this returns the decoded $\mathbf{x}^{[b_1, b_2, \dots, b_\lambda, 0]}$, it then builds $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda, 1]}$ with respect to those hard decisions and calls itself to decode $\mathbf{x}^{[b_1, b_2, \dots, b_\lambda, 1]}$. Then, the two decoded vectors are combined to form $\mathbf{x}^{[b_1, b_2, \dots, b_\lambda]}$ and the function returns. The following lemma makes this precise.

Lemma 8. Let \mathcal{T} be a base trellis with $N = 2^n$ sections corresponding to a received word \mathbf{y} such that (13) holds for the corresponding path-sum function. For each $i \in [N]$ in order, let \hat{u}_1^{i-1} be a vector of past decisions and $b_1, b_2, \dots, b_n \in \{0, 1\}$ satisfy $i(\mathbf{b}) = i$. Construct $\mathcal{T}^{[b_1, b_2, \dots, b_n]}$ iteratively as follows. For $\lambda = 1, 2, \dots, n$, let us define

$$\mathcal{T}^{[b_1, b_2, \dots, b_\lambda]} \triangleq \begin{cases} (\mathcal{T}^{[b_1, b_2, \dots, b_{\lambda-1}]})^{[b_\lambda]} & \text{if } \lambda \geq 2, \\ \mathcal{T}^{[b_1]} & \text{if } \lambda = 1. \end{cases}$$

If $b_\lambda = 1$, then we apply the plus transform with respect to the fixed vector

$$\mathbf{z} = \mathcal{A}_{n-\lambda}^{-1}(\hat{u}_\tau^\theta), \quad (16)$$

where $\hat{u}_\tau^\theta \triangleq (\hat{u}_\tau, \hat{u}_{\tau+1}, \dots, \hat{u}_\theta)$ and

$$\theta = \sum_{j=1}^{\lambda} b_j 2^{n-j}, \quad \tau = \theta - 2^{n-\lambda} + 1. \quad (17)$$

Then, for $\mathbf{U} = \mathcal{A}_n(\mathbf{X}) \in \mathcal{X}^N$, we have

$$T^{[b_1, b_2, \dots, b_n]}(u) = \Pr(U_i = u, U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y} = \mathbf{y}).$$

Proof: To facilitate a proof by induction, we actually prove a stronger claim. Namely, let $0 \leq \lambda \leq n$ be given. Define \mathbf{b}_λ

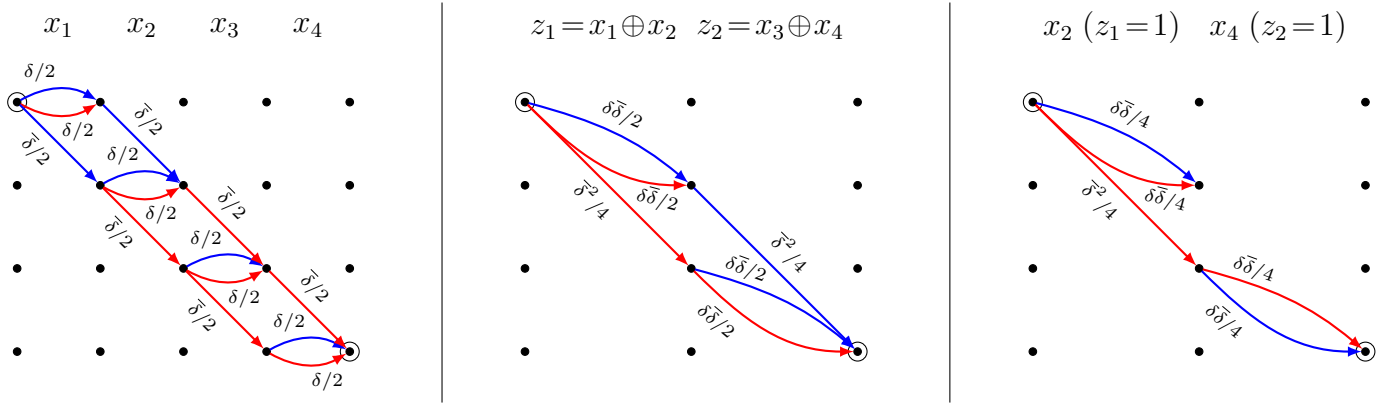


Fig. 3. The left panel continues the example from Fig. 2 by showing the same trellis except that edge weights are shown and edges with zero weight have been pruned. Like before, we color the edge blue if the associated input is 0 and red if it is 1. The middle panel shows the minus trellis for this example where adjacent stages have been merged and the input labels correspond to the modulo-2 sum of adjacent inputs. Here the edge colors are for the merged inputs $z_1 = x_1 \oplus x_2$ and $z_2 = x_3 \oplus x_4$. The right panel shows the plus trellis where the edges depend on a previous decision (z_1 or z_2) and are labeled by one input (x_2 or x_4). The edge colors indicate the values of the inputs x_2 and x_4 .

as the vector in $\{0, 1\}^n$ whose first λ entries equal those of \mathbf{b} , while the remaining entries are all-zero. That is,

$$\mathbf{b}_\lambda = (b_1, b_2, \dots, b_\lambda, 0, 0, \dots, 0). \quad (18)$$

Recalling the notation in (1)–(5), we will prove that for all $\boldsymbol{\mu} \in \mathcal{X}^{2^{n-\lambda}}$,

$$\begin{aligned} T^{[b_1, b_2, \dots, b_\lambda]}(\boldsymbol{\mu}) \\ = P(\mathbf{X}^{[b_1, b_2, \dots, b_\lambda]} = \boldsymbol{\mu}, U_1^{i(\mathbf{b}_\lambda)-1} = \hat{u}_1^{i(\mathbf{b}_\lambda)-1}, \mathbf{Y} = \mathbf{y}). \end{aligned} \quad (19)$$

For $\lambda = n$ this reduces to the claimed lemma, by (5).

The proof of (19) proceeds by induction on λ . For the base case, take $\lambda = 0$, and note that (19) holds by assumption: the LHS is by definition $T(\boldsymbol{\mu})$ while the RHS is simply $P(\mathbf{X} = \boldsymbol{\mu}, \mathbf{Y} = \mathbf{y})$, and the two are equal by (13).

For the induction step, we assume that (19) is true for λ , and prove it to be true for $\lambda + 1$. Assume first that $b_{\lambda+1} = 0$. In this case, $\mathbf{b}_\lambda = \mathbf{b}_{\lambda+1}$. Recall that since $b_{\lambda+1} = 0$, we get the trellis $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]}$ by applying a minus transform (Definition 5) on $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda]}$. We must prove that (19) holds with $\lambda + 1$ in place of λ , and this is indeed the case by Lemma 6. Indeed, recall that by our recursive definition, $\mathbf{X}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]} = (\mathbf{X}^{[b_1, b_2, \dots, b_\lambda]})^{[0]}$, and apply Lemma 6, where in (13) and (14) we replace \mathbf{X} , \mathbf{Y} , and \mathbf{y} with $\mathbf{X}^{[b_1, b_2, \dots, b_\lambda]}$, $(\mathbf{Y}, U_1^{i(\mathbf{b}_\lambda)-1})$, and $(\mathbf{y}, \hat{u}_1^{i(\mathbf{b}_\lambda)-1})$, respectively.

Now, let us assume that $b_{\lambda+1} = 1$. Because of this, note that $\mathbf{b}_\lambda \neq \mathbf{b}_{\lambda+1}$. As before, we assume that (19) is true for λ , and prove it to be true for $\lambda + 1$. By definition, we get the trellis $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]}$ by applying a plus transform (Definition 6) on $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda]}$, with respect to the vector \mathbf{z} defined in (16) and (17), with λ replaced by $\lambda + 1$. Thus, if we denote by T the probability function associated with $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda]}$, we get by Lemma 7 that the probability function associated with $\mathcal{T}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]}$, which we denote by T' , satisfies

$$\begin{aligned} T'(\mathbf{z}') &= T(\boldsymbol{\mu}) \\ &= P(\mathbf{X}^{[b_1, b_2, \dots, b_\lambda]} = \boldsymbol{\mu}, U_1^{i(\mathbf{b}_\lambda)-1} = \hat{u}_1^{i(\mathbf{b}_\lambda)-1}, \mathbf{Y} = \mathbf{y}), \end{aligned}$$

where $\boldsymbol{\mu}$ is the unique vector for which $\boldsymbol{\mu}^{[0]} = \mathbf{z}$ and $\boldsymbol{\mu}^{[1]} = \mathbf{z}'$. For this definition of $\boldsymbol{\mu}$, note that the condition $\mathbf{X}^{[b_1, b_2, \dots, b_\lambda]} = \boldsymbol{\mu}$ is equivalent to the pair of conditions

$$\mathbf{X}^{[b_1, b_2, \dots, b_\lambda, 0]} = \mathbf{z} \quad \text{and} \quad \mathbf{X}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]} = \mathbf{z}'.$$

We will shortly prove that the pair of conditions

$$U_1^{i(\mathbf{b}_\lambda)-1} = \hat{u}_1^{i(\mathbf{b}_\lambda)-1} \quad \text{and} \quad \mathbf{X}^{[b_1, b_2, \dots, b_\lambda, 0]} = \mathbf{z} \quad (20)$$

can be simplified to

$$U_1^{i(\mathbf{b}_{\lambda+1})-1} = \hat{u}_1^{i(\mathbf{b}_{\lambda+1})-1}. \quad (21)$$

Once this is proved, the lemma follows, since the above implies that

$$\begin{aligned} T'(\mathbf{z}') &= \\ &= P(\mathbf{X}^{[b_1, b_2, \dots, b_\lambda, b_{\lambda+1}]} = \mathbf{z}', U_1^{i(\mathbf{b}_{\lambda+1})-1} = \hat{u}_1^{i(\mathbf{b}_{\lambda+1})-1}, \mathbf{Y} = \mathbf{y}). \end{aligned}$$

Let us now show that (20) is equivalent to (21). To get from the vector $U_1^{i(\mathbf{b}_\lambda)-1}$ in (20) to the vector $U_1^{i(\mathbf{b}_{\lambda+1})-1}$ in (21) we need to append to the former the vector $U_i^{i(\mathbf{b}_{\lambda+1})-1}$. That is, by (4) and (5), and recalling that $b_{\lambda+1} = 1$, we need to append all the length- n polar transforms of \mathbf{X} with prefix $(b_1, b_2, \dots, b_\lambda, 0)$. Thus, recalling (6), all that remains is to show that $\hat{u}_i^{i(\mathbf{b}_{\lambda+1})-1} = \mathcal{A}_{n-(\lambda+1)}(\mathbf{z})$. To see that this indeed is the case, we observe that \mathbf{z} is defined by (16) and (17) with $\lambda + 1$ in place of λ . Recalling (4) and (18), and keeping in mind that in (17) we replace λ by $\lambda + 1$, we see that $\theta = i(\mathbf{b}_{\lambda+1}) - 1$ while $\tau = i(\mathbf{b}_\lambda)$. \blacksquare

Actually, the above lemma is not unique to the deletion channel and it applies to any base trellis for which (13) holds. The above lemma also gives an efficient method for deciding⁵ the value of \hat{u}_i at stage i , since

⁵Note that the $\arg \max_{u \in \{0, 1\}}$ of $\Pr(U_i = u | U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y} = \mathbf{y})$ is equal to that of $\Pr(U_i = u, U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y} = \mathbf{y})$. Thus, if we are using SC decoding to decide on the value of \hat{u}_i , where U_i is non-frozen, then we can save some computation time by not carrying out the division on the RHS of (22). However, if the trellis is used to track probabilities on the encoding side, then we must use conditioning when employing the Honda-Yamamoto scheme.

$$\begin{aligned} \Pr(U_i = u | U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y} = \mathbf{y}) \\ = \frac{T^{[b_1, b_2, \dots, b_n]}(u)}{\sum_{u' \in \mathcal{X}} T^{[b_1, b_2, \dots, b_n]}(u')} \end{aligned} \quad (22)$$

when $\Pr(U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y} = \mathbf{y}) > 0$.

D. Complexity

In [17], SC trellis decoding is generalized to finite-state channels with memory. For a finite-state channel with A states, the decoding complexity of a length- N code is shown to be $O(A^3 N \log N)$. While there are some connections between finite-state channels and deletion channels [10], it is not clear if this complexity result can be applied directly to the deletion channel. Using a different formulation, an SC decoder for polar codes on the deletion channel is defined in [13]. Its complexity is $O(N^4 \log N)$ for a constant deletion rate and a uniform input distribution⁶.

In this section, we bound the complexity of computing the plus and minus transformations of a trellis. For a trellis \mathcal{T} with N sections, let $P_2(j)$ be the number of distinct 2-step paths from states in \mathcal{V}_{2j} to states in \mathcal{V}_{2j+2} and define

$$C(\mathcal{T}) \triangleq \sum_{j=0}^{N/2-1} P_2(j).$$

From Definition 5, one can verify that the minus transform requires $C(\mathcal{T})$ multiplies and adds to compute $\mathcal{T}^{[0]}$. Similarly, from Definition 6, it follows that the plus transform requires at most $C(\mathcal{T})$ multiplies and adds to compute $\mathcal{T}^{[1]}$.

Consider a trellis \mathcal{T}_λ at depth- λ in the decoding process. Such a trellis will have $2^{n-\lambda}$ sections each corresponding to 2^λ channel uses. For the deletion channel, we observe that each state in \mathcal{V}_{2j} has at most $2(2^\lambda + 1)|\mathcal{S}|$ outgoing edges. This is because each edge can be labeled by 0 or 1, the number of deletions (between 0 and 2^λ) determines the change in the channel state, and the input state can change to any of $|\mathcal{S}|$ possibilities. Combining these observations, and noting that the number of vertices in each segment is at most $2^n |\mathcal{S}|$, we see that

$$C(\mathcal{T}_\lambda) \leq 2^n |\mathcal{S}| \cdot (2(2^\lambda + 1)|\mathcal{S}|)^2 2^{n-\lambda} \leq 2^{2n+2} (2^\lambda + 3) |\mathcal{S}|^3.$$

Since the full decoder uses 2^λ plus and minus operations at depth λ , the overall decoding complexity is

$$\sum_{\lambda=0}^{n-1} 2^\lambda 2^{2n+2} (2^\lambda + 3) |\mathcal{S}|^3 = O(|\mathcal{S}|^3 N^4),$$

which is lower than previous methods by a $\log N$ factor. This occurs because the $\lambda = n - 1$ decoding step dominates the calculation and has $O(|\mathcal{S}|^3 N^4)$ complexity by itself.

The reader should happily note that the above quartic growth in N is *not* present in Theorem 1. The overall complexity of our scheme is much smaller because the guard bands allow the codeword to be separated into many smaller blocks whose trellises can be processed separately.

⁶As noted earlier, the complexity of the decoding algorithm in [13] is misstated as $O(d^2 N \log N)$ for d deletions but it is actually $O(d^3 N \log N)$.

V. INFORMATION RATES

In this section, we will introduce and analyze various information rates related to polar codes on the deletion channel. For a given regular hidden-Markov input distribution, let \mathbf{X} be an input vector of length N and let \mathbf{Y} be the corresponding output vector (i.e., the observation of \mathbf{X} through the deletion channel). The main goal of this paper is to show that our polar coding scheme achieves the information rate

$$\mathcal{I} = \lim_{N \rightarrow \infty} \frac{I(\mathbf{X}; \mathbf{Y})}{N}, \quad (23)$$

where \mathbf{X} and \mathbf{Y} depend implicitly on N . This existence of this limit is well-known [2] but we revisit it here because the same argument will be used later with slight variations.

Lemma 9. *Fix a hidden-Markov input distribution. For a given N , let $\mathbf{X} = (X_1, X_2, \dots, X_N)$ be a random vector with the above distribution. Let \mathbf{Y} be the result of passing \mathbf{X} through a deletion channel with deletion probability δ . Then, the following two limits exist,*

$$\lim_{N \rightarrow \infty} \frac{H(\mathbf{X})}{N} \quad \text{and} \quad \lim_{N \rightarrow \infty} \frac{H(\mathbf{X}|\mathbf{Y})}{N}. \quad (24)$$

Proof: The proof of this lemma is detailed below for uniform inputs in Section V-A and hidden-Markov inputs in Section V-B. ■

Once the limits in (24) are established, the limit in (23) follows because

$$\frac{I(\mathbf{X}; \mathbf{Y})}{N} = \frac{H(\mathbf{X})}{N} - \frac{H(\mathbf{X}|\mathbf{Y})}{N}.$$

A. Uniform input

In this subsection, we prove Lemma 9, for the restricted case in which the input distribution is i.i.d. and uniform.

Proof of Lemma 9 for Uniform Inputs: In such a setting, the first limit in (24) clearly exists and equals 1. To prove the second limit in (24), let us first define

$$\mathcal{H}_N = H(\mathbf{X}|\mathbf{Y}), \quad |\mathbf{X}| = N. \quad (25)$$

Our plan is to show that the sequence \mathcal{H}_N is superadditive, implying [25, Lemma 1.2.1, page 3] the existence of the second limit in (24). Indeed, let N_1 and N_2 be given, and let \mathbf{X} and \mathbf{X}' be distributed according to the input distribution, and having lengths N_1 and N_2 , respectively. Denote the outputs corresponding to \mathbf{X} and \mathbf{X}' by \mathbf{Y} and \mathbf{Y}' , respectively. We have

$$\begin{aligned} \mathcal{H}_{N_1+N_2} &= H(\mathbf{X} \circ \mathbf{X}' | \mathbf{Y} \circ \mathbf{Y}') \\ &\stackrel{(a)}{=} H(\mathbf{X}, \mathbf{X}' | \mathbf{Y} \circ \mathbf{Y}') \\ &\geq H(\mathbf{X}, \mathbf{X}' | \mathbf{Y} \circ \mathbf{Y}', \mathbf{Y}, \mathbf{Y}') \\ &\stackrel{(b)}{=} H(\mathbf{X}, \mathbf{X}' | \mathbf{Y}, \mathbf{Y}') \\ &\stackrel{(c)}{=} H(\mathbf{X} | \mathbf{Y}, \mathbf{Y}') + H(\mathbf{X}' | \mathbf{X}, \mathbf{Y}, \mathbf{Y}') \\ &\stackrel{(d)}{=} H(\mathbf{X} | \mathbf{Y}) + H(\mathbf{X}' | \mathbf{Y}') \\ &= \mathcal{H}_{N_1} + \mathcal{H}_{N_2}, \end{aligned}$$

where (a) holds because N_1 and N_2 , the lengths of \mathbf{X} and \mathbf{X}' , respectively, are constant parameters; (b) holds because

$\mathbf{Y} \odot \mathbf{Y}'$ is a function of \mathbf{Y} and \mathbf{Y}' ; (c) follows by the chain rule; (d) holds because, for the i.i.d. uniform input distribution, the pair (\mathbf{X}, \mathbf{Y}) is independent of the pair $(\mathbf{X}', \mathbf{Y}')$. Hence, the sequence \mathcal{H}_N is indeed superadditive. ■

B. Hidden-Markov input

We now prove Lemma 9 for the case where the input distribution is a regular hidden-Markov process. Since now \mathcal{H}_N is not generally superadditive, we will take an indirect route to prove Lemma 9. Indeed, the following lemma is proved by defining a related quantity, $\hat{\mathcal{H}}_N$, which is superadditive. Recall from (8) that in the following, S_0 is the state of the process before X_1 is output and S_N is the state before X_{N+1} is output.

Lemma 10. *Fix a regular hidden-Markov input distribution. For a given N , let $\mathbf{X} = (X_1, X_2, \dots, X_N)$ be a random vector with the above distribution. Let \mathbf{Y} be the result of passing \mathbf{X} through a deletion channel with deletion probability δ . Then, the following limit exists:*

$$\lim_{N \rightarrow \infty} \frac{H(\mathbf{X}|\mathbf{Y}, S_0, S_N)}{N}. \quad (26)$$

Proof: Define

$$\hat{\mathcal{H}}_N = H(\mathbf{X}|\mathbf{Y}, S_0, S_N), \quad |\mathbf{X}| = N. \quad (27)$$

To borrow the terminology of [21], the above defines the *boundary-state-aware entropy*. Note that S_0 and S_N are the states just before transmission has started, and just after transmission has ended, respectively.

We now show that $\hat{\mathcal{H}}_N$ is superadditive. Indeed, let \mathbf{X} and \mathbf{X}' be consecutive input vectors of length N_1 and N_2 , respectively. That is, $\mathbf{X} \odot \mathbf{X}'$ is a vector of length $N_1 + N_2$ drawn from the input distribution. Denote by \mathbf{Y} and \mathbf{Y}' the output vectors corresponding to \mathbf{X} and \mathbf{X}' , respectively. Then,

$$\begin{aligned} \hat{\mathcal{H}}_{N_1+N_2} &= H(\mathbf{X} \odot \mathbf{X}'|\mathbf{Y} \odot \mathbf{Y}', S_0, S_{N_1+N_2}) \\ &\stackrel{(a)}{=} H(\mathbf{X}, \mathbf{X}'|\mathbf{Y} \odot \mathbf{Y}', S_0, S_{N_1+N_2}) \\ &\stackrel{(b)}{\geq} H(\mathbf{X}, \mathbf{X}'|\mathbf{Y}, \mathbf{Y}', S_0, S_{N_1+N_2}) \\ &\geq H(\mathbf{X}, \mathbf{X}'|\mathbf{Y}, \mathbf{Y}', S_0, S_{N_1}, S_{N_1+N_2}) \\ &\stackrel{(c)}{=} H(\mathbf{X}|\mathbf{Y}, \mathbf{Y}', S_0, S_{N_1}, S_{N_1+N_2}) \\ &\quad + H(\mathbf{X}'|\mathbf{X}, \mathbf{Y}, \mathbf{Y}', S_0, S_{N_1}, S_{N_1+N_2}) \\ &\stackrel{(d)}{=} H(\mathbf{X}|\mathbf{Y}, S_0, S_{N_1}) + H(\mathbf{X}'|\mathbf{Y}', S_{N_1}, S_{N_1+N_2}) \\ &= \hat{\mathcal{H}}_{N_1} + \hat{\mathcal{H}}_{N_2}, \end{aligned}$$

where (a) holds because N_1 and N_2 , the lengths of \mathbf{X} and \mathbf{X}' , respectively, are constant parameters; (b) holds because $\mathbf{Y} \odot \mathbf{Y}'$ is a function of \mathbf{Y} and \mathbf{Y}' ; (c) follows by the chain rule; (d) holds because of conditional independence: given S_{N_1} , $(\mathbf{X}, \mathbf{Y}, S_0)$ is independent of $(\mathbf{X}', \mathbf{Y}', S_{N_1+N_2})$. Hence, the sequence $\hat{\mathcal{H}}_N$ is indeed superadditive, and the following limit exists by [25, Lemma 1.2.1, page 3],

$$\lim_{N \rightarrow \infty} \frac{\hat{\mathcal{H}}_N}{N}. \quad \blacksquare$$

All that remains now is to account for the difference in the entropies of \mathcal{H}_N and $\hat{\mathcal{H}}_N$, incurred by conditioning on S_0

and S_N . As will be made clear in the following proof, this difference can be bounded by a constant, and hence vanishes when we divide by N .

Proof of Lemma 9 for hidden-Markov inputs: We first note that the existence of the second limit in (24) implies the existence of the first limit. Indeed, taking the deletion probability δ equal to 1 makes the second limit equal the first. Hence, all that remains is to prove the existence of the second limit.

To show that the second limit in (24) exists, note that, for $|\mathbf{X}| = N$, we have on the one hand that

$$\begin{aligned} H(\mathbf{X}, S_0, S_N|\mathbf{Y}) &= H(\mathbf{X}|\mathbf{Y}) + H(S_0, S_N|\mathbf{X}, \mathbf{Y}) \\ &\geq H(\mathbf{X}|\mathbf{Y}) \\ &= \mathcal{H}_N, \end{aligned}$$

and on the other hand that

$$\begin{aligned} H(\mathbf{X}, S_0, S_N|\mathbf{Y}) &= H(S_0, S_N|\mathbf{Y}) + H(\mathbf{X}|\mathbf{Y}, S_0, S_N) \\ &\leq 2 \log_2 |\mathcal{S}| + H(\mathbf{X}|\mathbf{Y}, S_0, S_N) \\ &= 2 \log_2 |\mathcal{S}| + \hat{\mathcal{H}}_N. \end{aligned}$$

Thus,

$$\mathcal{H}_N \leq \hat{\mathcal{H}}_N + 2 \log_2 |\mathcal{S}|.$$

Since it is easily seen that $\hat{\mathcal{H}}_N \leq \mathcal{H}_N$, we have that

$$\frac{\hat{\mathcal{H}}_N}{N} \leq \frac{\mathcal{H}_N}{N} \leq \frac{\hat{\mathcal{H}}_N}{N} + \frac{2 \log_2 |\mathcal{S}|}{N}. \quad (28)$$

We have already proved that the limit of the LHS of (28) exists, in Lemma 10. Since the limit of $(2 \log_2 |\mathcal{S}|)/N$ is 0, the limit of the RHS of (28) exists and equals that of the LHS. By the sandwich property, the limit of the middle term exists as well, which is the desired result. ■

We finish by restating the last part of the proof as a lemma.

Lemma 11. *Fix a hidden-Markov input distribution. For a given N , let $\mathbf{X} = (X_1, X_2, \dots, X_N)$ be a random vector with the above distribution. Let \mathbf{Y} be the result of passing \mathbf{X} through a deletion channel with deletion probability δ . Then,*

$$\lim_{N \rightarrow \infty} \frac{H(\mathbf{X}|\mathbf{Y}, S_0, S_N)}{N} = \lim_{N \rightarrow \infty} \frac{H(\mathbf{X}|\mathbf{Y})}{N}. \quad (29)$$

VI. WEAK POLARIZATION

In this section, we prove weak polarization for both the deletion channel and the trimmed deletion channel, as defined in Subsection III-C. The term ‘‘weak polarization’’ alludes to ϵ being fixed in Theorems 12 and 20 below. As in [24], we will first prove that a certain process is submartingale, and then prove that it either converges to 0 or to 1.

As a first step, we will shortly define three entropies. These are defined with respect to an input \mathbf{X} of length $N = 2^n$, which has a regular hidden-Markov input distribution, and $\mathbf{U} = \mathcal{A}_n(\mathbf{X})$. The corresponding output is denoted \mathbf{Y} . Recall that S_0 and S_N are the (hidden) states of the input process, just before \mathbf{X} is transmitted and right after \mathbf{X} is transmitted, respectively. Lastly, denote by \mathbf{Y}^* the result of trimming all leading and trailing ‘0’ symbols from \mathbf{Y} . Then, for a given

n and $1 \leq i \leq N = 2^n$, define the following (deterministic) entropies:

$$h_i = H(U_i | U_1^{i-1}, \mathbf{Y}), \quad (30)$$

$$\hat{h}_i = H(U_i | U_1^{i-1}, S_0, S_N, \mathbf{Y}), \quad (31)$$

$$h_i^* = H(U_i | U_1^{i-1}, \mathbf{Y}^*). \quad (32)$$

Clearly,

$$h_i^* \geq h_i \geq \hat{h}_i.$$

Note that in the case of a uniform input distribution, there is only one state, and hence h_i and \hat{h}_i are equal.

Following [24], we show weak polarization by considering a sequence B_1, B_2, \dots of i.i.d. $\text{Ber}(1/2)$ random variables. For any $n \in \mathbb{N}$, let $J_n = i(B_1, B_2, \dots, B_n)$ be the random index defined by (4), with B_t in place of b_t . We will study the three related random processes defined for $n \in \mathbb{N}$ by

$$H_n = h_{J_n}, \quad (33)$$

$$\hat{H}_n = \hat{h}_{J_n}, \quad (34)$$

$$H_n^* = h_{J_n}^*. \quad (35)$$

The arguments below will show that \hat{H}_n is a bounded submartingale. Thus, the sequence \hat{H}_n converges and the limit is denoted by \hat{H}_∞ . One can also argue that $\hat{H}_\infty \in \{0, 1\}$. From this, we will infer that H_n and H_n^* must converge to either 0 or 1 as well. Though neither H_n nor H_n^* are necessarily submartingales.

Theorem 12. *The sequence \hat{H}_n converges (almost surely and in L^1) to a well-defined random variable $\hat{H}_\infty \in \{0, 1\}$ and, for any $\epsilon > 0$, it follows that*

$$\frac{1}{N} |\{i \in [N] \mid H(U_i | U_1^{i-1}, S_0, S_N, \mathbf{Y}) \in [\epsilon, 1 - \epsilon]\}| \rightarrow 0. \quad (36)$$

Proof: Lemma 13 below shows that $\hat{H}_1, \hat{H}_2, \hat{H}_3, \dots \in [0, 1]$ is a bounded submartingale with respect to J_n . This implies that the sequence \hat{H}_n converges (almost surely and in L^1) to a limit that is denoted by \hat{H}_∞ [26, p. 236]. Lemma 18 below shows that, for any $\epsilon > 0$, there is a $\Delta > 0$ such that $\hat{H}_n \in [\epsilon, 1 - \epsilon]$ implies $\hat{H}_{n+1} > \hat{H}_n + \Delta$ with probability $\frac{1}{2}$. Thus, the sequence \hat{H}_n cannot converge to any real number in the set $(0, 1)$ and hence $\hat{H}_\infty \in \{0, 1\}$.

From (31) and (34), we see that $\Pr(\hat{H}_n \in [\epsilon, 1 - \epsilon])$ equals

$$\frac{1}{N} |\{i \in [N] \mid H(U_i | U_1^{i-1}, S_0, S_N, \mathbf{Y}) \in [\epsilon, 1 - \epsilon]\}|.$$

Since \hat{H}_n converges almost surely to \hat{H}_∞ and $\epsilon, 1 - \epsilon$ are continuity points of $\Pr(\hat{H}_\infty \leq x)$ [26, Ch. 4], it follows that

$$\lim_{n \rightarrow \infty} \Pr(\hat{H}_n \in [\epsilon, 1 - \epsilon]) = \Pr(\hat{H}_\infty \in [\epsilon, 1 - \epsilon]) = 0.$$

This completes the proof. \blacksquare

Lemma 13. *For a hidden-Markov input distribution and a deletion channel with deletion probability δ , let \hat{H}_n and J_n be as defined above. Then, the sequence $\hat{H}_1, \hat{H}_2, \hat{H}_3, \dots$ is a bounded submartingale with respect to the J_1, J_2, J_3, \dots sequence.*

Proof: Since \hat{H}_n is clearly bounded between 0 and 1, it remains to show that $E(\hat{H}_{n+1} | J_1, J_2, \dots, J_n) \geq \hat{H}_n$. Let $\mathbf{X} \odot \mathbf{X}'$ be a length- $2N$ input to the channel. Denote by $\mathbf{Y} \odot \mathbf{Y}'$ the corresponding output, where \mathbf{Y} only contains inputs from \mathbf{X} and \mathbf{Y}' only contains inputs from \mathbf{X}' . Recall that $\mathbf{U} = \mathcal{A}_n(\mathbf{X})$ and define $\mathbf{V} = \mathcal{A}_n(\mathbf{X}')$ and

$$\mathbf{F} = (U_1 \oplus V_1, V_1, U_2 \oplus V_2, V_2, \dots, U_N \oplus V_N, V_N).$$

By (4), we have that $J_{n+1} = 2J_n - 1$ with probability $1/2$ and $J_{n+1} = 2J_n$ with probability $1/2$. Thus,

$$\begin{aligned} E(\hat{H}_{n+1} | J_1^n) &= E(H(F_{J_{n+1}} | F_1^{J_{n+1}-1}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) | J_1^n) \\ &= \frac{1}{2} H(F_{2J_n-1} | F_1^{2J_n-2}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &\quad + \frac{1}{2} H(F_{2J_n} | F_1^{2J_n-1}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &= \frac{1}{2} H(F_{2J_n-1}, F_{2J_n} | F_1^{2J_n-2}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &= \frac{1}{2} H(U_{J_n} \oplus V_{J_n}, V_{J_n} | F_1^{2J_n-2}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &= \frac{1}{2} H(U_{J_n}, V_{J_n} | U_1^{J_n-1}, V_1^{J_n-1}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &\stackrel{(a)}{\geq} \frac{1}{2} H(U_{J_n}, V_{J_n} | U_1^{J_n-1}, V_1^{J_n-1}, \mathbf{Y}, \mathbf{Y}', S_0, S_{2N}) \\ &\stackrel{(b)}{\geq} \frac{1}{2} H(U_{J_n}, V_{J_n} | U_1^{J_n-1}, V_1^{J_n-1}, \mathbf{Y}, \mathbf{Y}', S_0, S_N, S_{2N}) \\ &\stackrel{(c)}{=} \frac{1}{2} H(U_{J_n} | U_1^{J_n-1}, \mathbf{Y}, S_0, S_N) \\ &\quad + \frac{1}{2} H(V_{J_n} | V_1^{J_n-1}, \mathbf{Y}', S_N, S_{2N}) \\ &\stackrel{(d)}{=} \hat{H}_n. \end{aligned}$$

The inequality (a) follows from the fact that $\mathbf{Y} \odot \mathbf{Y}'$ is a deterministic function of \mathbf{Y}, \mathbf{Y}' . Inequality (b) follows since conditioning reduces entropy. Step (c) holds by the Markov property. Finally, (d) is due to stationarity: $\hat{H}_n = H(U_{J_n} | U_1^{J_n-1}, \mathbf{Y}, S_0, S_N) = H(V_{J_n} | V_1^{J_n-1}, \mathbf{Y}', S_N, S_{2N})$. \blacksquare

Since the sequence \hat{H}_n is a bounded submartingale, it converges almost surely and in L_1 to a random variable $\hat{H}_\infty \in [0, 1]$. To show that $\hat{H}_\infty \in \{0, 1\}$ with probability 1, one can show that, if $\epsilon \leq \hat{H}_n \leq 1 - \epsilon$, then there is a $\Delta = \Delta(\epsilon) > 0$ such that $\hat{H}_n^- - \hat{H}_n > \Delta(\epsilon)$, where

$$\hat{H}_n^- \triangleq H(U_{J_n} \oplus V_{J_n} | U_1^{J_n-1}, V_1^{J_n-1}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}). \quad (37)$$

That is, a ‘minus’ operation applied to non-polarized entropy changes the entropy by at least Δ . Such a result indeed establishes the above, since it dictates that \hat{H}_n cannot converge to anything other than either 0 or 1. As before, we first prove the above for the simple case of i.i.d. uniform input, and then generalize to a hidden-Markov input.

A. Uniform input

Lemma 14. *Let \mathbf{X} and \mathbf{X}' be independent vectors of length $N = 2^n$, both drawn from an i.i.d. uniform distribution. Let \hat{H}_n and \hat{H}_n^- be as defined in (31), (34), and (37), with S_0, S_N*

and S_{2N} being degenerate random variables always taking the value 1. Then, for every $\epsilon > 0$ there exists $\Delta(\epsilon) > 0$ such that if $\epsilon \leq \hat{H}_n \leq 1 - \epsilon$, then $\hat{H}_n^- - \hat{H}_n > \Delta(\epsilon)$.

Proof: Denote $i = J_n$, and assume a fixed ϵ for which $\epsilon \leq \hat{H}_n \leq 1 - \epsilon$. Then, since S_0 , S_N , and S_{2N} are degenerate, we observe that $(U_i, U_1^{i-1}, \mathbf{Y})$ is independent of $(V_i, V_1^{i-1}, \mathbf{Y}')$. It follows that

$$H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}')$$

is the entropy of the modulo-2 sum of the independent binary random variables U_i and V_i . Thus, Mrs. Gerber's Lemma [27, Lemma 2.2] implies that, for every $\epsilon > 0$, there is $\Delta > 0$ such that

$$H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}') - H(U_i | U_1^{i-1}, \mathbf{Y}) \geq \Delta.$$

Since

$$\begin{aligned} \hat{H}_{n+1}^- &= H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y} \odot \mathbf{Y}') \\ &\geq H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}'), \end{aligned}$$

the result follows. \blacksquare

B. Hidden-Markov input

The proof of Lemma 14 above relied on the mutual independence of $(U_i, U_1^{i-1}, \mathbf{Y})$ and $(V_i, V_1^{i-1}, \mathbf{Y}')$. To emulate⁷ this property in a FAIM setting, we note that for s_0 , s_N , and s_{2N} fixed, we indeed have that $(U_i, U_1^{i-1}, \mathbf{Y})$ and $(V_i, V_1^{i-1}, \mathbf{Y}')$ are independent, when conditioning on the event $S_0 = s_0$, $S_N = s_N$, $S_{2N} = s_{2N}$. Towards this end, for $s_0, s_N, s_{2N} \in \mathcal{S}$, we denote the probability of these three states occurring as

$$p(s_0, s_N, s_{2N}) = \Pr(S_0 = s_0, S_N = s_N, S_{2N} = s_{2N}). \quad (38)$$

In the remainder of this subsection, we will assume that N is large enough such that the above probability is always positive. This is indeed possible, by the following lemma.

Lemma 15. *For $s \in \mathcal{S}$, denote by $\pi(s)$ the stationary probability of s . That is, the probability that $S_0 = s$. Let*

$$\pi_{\min} = \min_{s \in \mathcal{S}} \pi(s),$$

Then, $\pi_{\min} > 0$, and there exists a ν such that for all $N \geq 2^\nu$ and all $s_0, s_N, s_{2N} \in \mathcal{S}$ we have

$$\Pr(S_0 = s_0, S_N = s_N, S_{2N} = s_{2N}) > \frac{(\pi_{\min})^3}{2}. \quad (39)$$

Proof: Since the underlying Markov chain is regular (i.e., finite-state, irreducible, and aperiodic), some power of the transition matrix must be strictly positive and this implies that $\pi_{\min} > 0$. Regularity further implies that S_0, S_N, S_{2N} become asymptotically independent as N increases. Thus, there must be an $N_0 = 2^{n_0}$ such that (39) holds for all $N \geq N_0$. \blacksquare

For (s_0, s_N, s_{2N}) , we define the quantities $\alpha(s_0, s_N, s_{2N})$ and $\beta(s_0, s_N, s_{2N})$ as follows.

$$\alpha(s_0, s_N, s_{2N}) \triangleq \quad (40)$$

⁷For independence, it is sufficient to condition on the event $S_N = s_N$. Conditioning on the more specific event $S_0 = s_0$, $S_N = s_N$, $S_{2N} = s_{2N}$ is needed for latter parts.

$$H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}', S_0 = s_0, S_N = s_N, S_{2N} = s_{2N})$$

and

$$\beta(s_0, s_N, s_{2N}) \triangleq \frac{\gamma(s_0, s_N) + \gamma(s_N, s_{2N})}{2}, \quad (41)$$

where

$$\gamma(s_0, s_N) \triangleq H(U_i | U_1^{i-1}, \mathbf{Y}, S_0 = s_0, S_N = s_N). \quad (42)$$

Note that by stationarity,

$$\gamma(s_N, s_{2N}) = H(V_i | V_1^{i-1}, \mathbf{Y}', S_N = s_N, S_{2N} = s_{2N}).$$

The following lemma states how α and β are related to our quantities of interest, \hat{H}_n and \hat{H}_n^- .

Lemma 16. *Let $N = 2^n > 2^\nu$, where ν was promised in Lemma 15. Then, for α and β as defined above, we have that*

$$\hat{H}_n^- \geq \sum_{s_0, s_N, s_{2N} \in \mathcal{S}} p(s_0, s_N, s_{2N}) \cdot \alpha(s_0, s_N, s_{2N}), \quad (43)$$

and

$$\hat{H}_n = \sum_{s_0, s_N, s_{2N} \in \mathcal{S}} p(s_0, s_N, s_{2N}) \cdot \beta(s_0, s_N, s_{2N}). \quad (44)$$

Furthermore, for all $s_0, s_N, s_{2N} \in \mathcal{S}$,

$$\alpha(s_0, s_N, s_{2N}) \geq \beta(s_0, s_N, s_{2N}). \quad (45)$$

Proof: Define $i = J_n$. To prove (43), we proceed similarly to the proof in Lemma 13 and deduce that

$$\begin{aligned} \hat{H}_n^- &= H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y} \odot \mathbf{Y}', S_0, S_{2N}) \\ &\geq H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}', S_0, S_{2N}) \\ &\geq H(U_i \oplus V_i | U_1^{i-1}, V_1^{i-1}, \mathbf{Y}, \mathbf{Y}', S_0, s_N, S_{2N}) \\ &= \sum_{s_0, s_N, s_{2N} \in \mathcal{S}} p(s_0, s_N, s_{2N}) \cdot \alpha(s_0, s_N, s_{2N}), \end{aligned}$$

The proof of (44) follows by stationarity. That is,

$$\begin{aligned} \hat{H}_n &= H(U_i | U_1^{i-1}, \mathbf{Y}, S_0, S_N) \\ &= \frac{H(U_i | U_1^{i-1}, \mathbf{Y}, S_0, S_N) + H(V_i | V_1^{i-1}, \mathbf{Y}', S_N, S_{2N})}{2} \\ &= \sum_{s_0, s_N, s_{2N} \in \mathcal{S}} p(s_0, s_N, s_{2N}) \cdot \frac{\gamma(s_0, s_N) + \gamma(s_N, s_{2N})}{2} \\ &= \sum_{s_0, s_N, s_{2N} \in \mathcal{S}} p(s_0, s_N, s_{2N}) \cdot \beta(s_0, s_N, s_{2N}). \end{aligned}$$

By (41), we deduce that (45) will follow from proving that

$$\alpha(s_0, s_N, s_{2N}) \geq \gamma(s_0, s_N) \quad (46)$$

and

$$\alpha(s_0, s_N, s_{2N}) \geq \gamma(s_N, s_{2N}) \quad (47)$$

W.l.o.g, we prove (46). Indeed, given that $S_N = s_N$, we have by the Markov property that $(S_0, U_1^{i-1}, U_i, \mathbf{Y})$ and $(V_1^{i-1}, V_i, \mathbf{Y}', S_{2N})$ are independent. Hence, for any s_{2N} we may also write γ , defined in (42), as

$$\begin{aligned} \gamma(s_0, s_N) &= H(U_i | U_1^{i-1}, V_1^{i-1}, V_i, \mathbf{Y}, \mathbf{Y}', \\ &\quad S_0 = s_0, S_N = s_N, S_{2N} = s_{2N}). \end{aligned}$$

Lastly, note that in the above expression for γ , since we condition on V_i , we could have written $U_i \oplus V_i$ in place of

U_i . This would give us the expression for α in (40), up to a further conditioning on V_i . Since conditioning reduces entropy, (46) follows. As noted, the proof of (47) is similar. Hence, we deduce (45). ■

In light of Lemma 16, our plan is to show the existence of a triplet (s_0, s_N, s_{2N}) for which $\alpha(s_0, s_N, s_{2N})$ is substantially greater than $\beta(s_0, s_N, s_{2N})$. The next lemma assures us such a triplet indeed exists.

Lemma 17. *For every $\epsilon > 0$ there exists a $\Delta' = \Delta'(\epsilon)$ for which the following holds. Let $N = 2^n > 2^\nu$, where ν was promised in Lemma 15. Then, if $\epsilon \leq \hat{H}_n \leq 1 - \epsilon$, then there exists a triplet s_0, s_N, s_{2N} such that*

$$\alpha(s_0, s_N, s_{2N}) > \beta(s_0, s_N, s_{2N}) + \Delta'. \quad (48)$$

Proof: By definition of γ in (42), we have that

$$\begin{aligned} \hat{H}_n &= \sum_{s_0, s_N \in \mathcal{S}} \Pr(S_0 = s_0, S_N = s_N) \cdot \gamma(s_0, s_N) \\ &= \sum_{s_N, s_{2N} \in \mathcal{S}} \Pr(S_N = s_N, S_{2N} = s_{2N}) \cdot \gamma(s_N, s_{2N}), \end{aligned} \quad (49)$$

where the second equality follows by stationarity. A crucial point will be to show the existence of a triplet (s_0, s_N, s_{2N}) for which $(\hat{H}_n - \gamma(s_0, s_N)) \cdot (\hat{H}_n - \gamma(s_N, s_{2N})) \leq 0$. In other words, either

$$\gamma(s_0, s_N) \leq \hat{H}_n \quad \text{and} \quad \gamma(s_N, s_{2N}) \geq \hat{H}_n, \quad (50)$$

or

$$\gamma(s_0, s_N) \geq \hat{H}_n \quad \text{and} \quad \gamma(s_N, s_{2N}) \leq \hat{H}_n. \quad (51)$$

To show this by contradiction, we start by supposing that this is not the case. Then, for all $s_0, s_N, s_{2N} \in \mathcal{S}$, it must be that

$$(\hat{H}_n - \gamma(s_0, s_N)) \cdot (\hat{H}_n - \gamma(s_N, s_{2N})) > 0. \quad (52)$$

Fix some arbitrary $a, b \in \mathcal{S}$. By specializing s_0 to a and s_N to b in (52), we deduce that $\hat{H}_n \neq \gamma(a, b)$. Assume w.l.o.g. that $\gamma(a, b) < \hat{H}_n$. We now claim that for all $c, d \in \mathcal{S}$,

$$\gamma(c, d) < \hat{H}_n. \quad (53)$$

Indeed, let $c, d \in \mathcal{S}$ be given. By setting $s_0 = a$, $s_N = b$, $s_{2N} = c$, we deduce from (52) that $\gamma(b, c) < \hat{H}_n$. Hence, if we set $s_0 = b$, $s_N = c$, $s_{2N} = d$ in (52), we deduce (53).

From the above paragraph, we conclude that for all $s_0, s_N \in \mathcal{S}$, we must have that $\gamma(s_0, s_N) < \hat{H}_n$. However, recalling from (49) that \hat{H}_n is a weighted average of such γ terms, we arrive at a contradiction. Hence, there exists a triplet (s_0, s_N, s_{2N}) for which either (50) or (51) holds. This is the triplet we are searching for. Indeed, since we have assumed that $\epsilon \leq \hat{H}_n \leq 1 - \epsilon$, the above triplet satisfies

$$\min\{\gamma(s_0, s_N), \gamma(s_N, s_{2N})\} \leq 1 - \epsilon$$

and

$$\max\{\gamma(s_0, s_N), \gamma(s_N, s_{2N})\} \geq \epsilon.$$

Our result now follows by combining part (i) of [27, Lemma 2.2] with⁸ [20, Lemma 11]. ■

⁸The first two strict inequalities in the statement of [20, Lemma 11] are essentially typos: they should both be replaced by weak inequalities, as is evident from reading the beginning of the proof.

Combining Lemmas 16 and 17 gives the following key result.

Lemma 18. *For every $\epsilon > 0$ there exists $\Delta = \Delta(\epsilon)$ for which the following holds. Let $N = 2^n > 2^\nu$, where ν was promised in Lemma 15. Then, if $\epsilon < \hat{H}_n \leq 1 - \epsilon$, then*

$$\hat{H}_n^- - \hat{H}_n > \Delta(\epsilon)$$

Proof: Take

$$\Delta = \frac{\Delta' \cdot (\pi_{\min})^3}{2},$$

where Δ' is as defined in Lemma 17. Now, simply combine (39), (43), (44), (45) and the existence of triplet s_0, s_N, s_{2N} for which (48) holds, to yield the claim. ■

The following lemma will be useful.

Lemma 19. *For $n \in \mathbb{N}$, let A_n and B_n be real random variables defined on a common probability space. Suppose B_n converges in L^1 to B_∞ and $E(A_n)$ converges to $E(B_\infty)$. If $A_n \geq B_n$ for all $n \in \mathbb{N}$, then A_n converges in L^1 to B_∞ .*

Proof: By definition, B_n converges to B_∞ in L^1 if and only if $E(|B_n - B_\infty|) \rightarrow 0$. Thus, by the triangle inequality,

$$\begin{aligned} E(|A_n - B_\infty|) &\leq E(|A_n - B_n|) + E(|B_n - B_\infty|) \\ &= E(A_n - B_n) + E(|B_n - B_\infty|) \\ &= E(A_n) - E(B_n) + E(|B_n - B_\infty|). \end{aligned}$$

In the limit, the first two terms converge to $E(B_\infty)$ and the last term converges to 0. Thus, $E(|A_n - B_\infty|) \rightarrow 0$. ■

The following theorem claims weak polarization for the three cases discussed earlier.

Theorem 20. *Fix $\epsilon \in (0, 1)$ and let $N = 2^n$. For a given hidden-Markov input distribution, let $\mathbf{X} = (X_1, X_2, \dots, X_N)$ be a random vector with the above distribution. Let \mathbf{Y} be the result of passing \mathbf{X} through a deletion channel with deletion probability δ . Denote $\mathbf{U} = \mathcal{A}(\mathbf{X})$. Let S_0 and S_N be as in Definition 1. Then,*

$$\lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}, S_0, S_N) < \epsilon\}|}{N} \quad (54a)$$

$$= \lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}) < \epsilon\}|}{N} \quad (54b)$$

$$= \lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}^*) < \epsilon\}|}{N} \quad (54c)$$

$$= 1 - \lim_{n \rightarrow \infty} \frac{H(\mathbf{X} | \mathbf{Y})}{N} \quad (54d)$$

and

$$\lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}, S_0, S_N) > 1 - \epsilon\}|}{N} \quad (55a)$$

$$= \lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}) > 1 - \epsilon\}|}{N} \quad (55b)$$

$$= \lim_{n \rightarrow \infty} \frac{|\{i : H(U_i | U_1^{i-1}, \mathbf{Y}^*) > 1 - \epsilon\}|}{N} \quad (55c)$$

$$= \lim_{n \rightarrow \infty} \frac{H(\mathbf{X} | \mathbf{Y})}{N}. \quad (55d)$$

Proof: For simplicity, the proof is split into 4 parts.

Part I: (54d) and (55d) are well defined: Recall from Lemma 9 that $\lim_{n \rightarrow \infty} H(\mathbf{X}|\mathbf{Y})/N$ exists. Thus, the right hand sides of both (54d) and (55d) are well defined.

Part II: (54a)=(54d) and (55a)=(55d): Since the Arkan transform is invertible, it follows that $\hat{\mathcal{H}}_N = H(\mathbf{X}|\mathbf{Y}, S_0, S_N) = H(\mathbf{U}|\mathbf{Y}, S_0, S_N)$, where $\hat{\mathcal{H}}_N$ is defined in (27). Thus, from the chain rule for entropy, we observe that

$$\begin{aligned} E(\hat{H}_n) &= \frac{1}{N} \sum_{i=1}^N H(U_i|U_1^{i-1}, \mathbf{Y}, S_0, S_N) \\ &= \frac{1}{N} H(\mathbf{U}|\mathbf{Y}, S_0, S_N) \\ &= \frac{1}{N} \hat{\mathcal{H}}_N. \end{aligned}$$

From Theorem 12, we see that \hat{H}_n converges in L^1 to $\hat{H}_\infty \in \{0, 1\}$. This implies that $E(\hat{H}_\infty) = \lim_{n \rightarrow \infty} E(\hat{H}_n)$ which exists and equals $\lim_{N \rightarrow \infty} \hat{\mathcal{H}}_N/N$ by Lemma 10. Since $\hat{H}_\infty \in \{0, 1\}$, observing that $E(\hat{H}_\infty) = \Pr(\hat{H}_\infty = 1)$ shows that

$$(55a) = \lim_{n \rightarrow \infty} \Pr(\hat{H}_n > 1 - \epsilon) = \Pr(\hat{H}_\infty = 1) = \lim_{n \rightarrow \infty} \frac{1}{N} \hat{\mathcal{H}}_N,$$

where the second equality holds because convergence in L^1 implies convergence in distribution and $1 - \epsilon$ is a continuity point of $\Pr(\hat{H}_\infty \leq x)$ [26, Ch. 4]. Since Lemma 11 shows that $\lim_{N \rightarrow \infty} \hat{\mathcal{H}}_N/N$ equals (55d), it follows that (55a) equals (55d). The last step is observing that

$$(54a) = \lim_{n \rightarrow \infty} \Pr(\hat{H}_n < \epsilon) = \Pr(\hat{H}_\infty = 0) = 1 - \Pr(\hat{H}_\infty = 1)$$

holds because convergence in L^1 implies convergence in distribution and ϵ is a continuity point of $\Pr(\hat{H}_\infty \leq x)$. Thus, (54a) equals (54d).

Part III: (54c)=(54d) and (55c)=(55d): To prove these equalities, we will apply Lemma 19 to the sequences $A_n = H_n^*$ and $B_n = \hat{H}_n$. Theorem 12 shows that \hat{H}_n converges in L^1 to \hat{H}_∞ and we established in the previous part that $E(\hat{H}_\infty)$ equals (55d). From the definitions in (34) and (35), it follows that $H_n^* \geq \hat{H}_n$ for all $n \in \mathbb{N}$. The only other element required for Lemma 19 is that $E(H_n^*) \rightarrow E(\hat{H}_\infty)$ and this will be shown below. Assuming this for now, we observe Lemma 19 implies that H_n^* converges in L^1 to \hat{H}_∞ and gives the desired result

$$(54c) = \lim_{n \rightarrow \infty} \Pr(H_n^* < \epsilon) = \Pr(\hat{H}_\infty < \epsilon) = (54d)$$

$$(55c) = \lim_{n \rightarrow \infty} \Pr(H_n^* > 1 - \epsilon) = \Pr(\hat{H}_\infty > 1 - \epsilon) = (55d),$$

where the second equality on each line holds because convergence in L^1 implies convergence in distribution and $\epsilon, 1 - \epsilon$ are continuity points of $\Pr(\hat{H}_\infty \leq x)$ [26, Ch. 4].

To show that $E(H_n^*) \rightarrow E(\hat{H}_\infty)$, we will use the fact that

$$\begin{aligned} H(\mathbf{U}|\mathbf{Y}, S_0, S_N) &\leq H(\mathbf{U}|\mathbf{Y}^*) \leq \\ &H(\mathbf{U}|\mathbf{Y}, S_0, S_N) + 2 \log_2 |\mathcal{S}| + 2 \log_2(N+1). \end{aligned} \quad (56)$$

Indeed, the first inequality holds because \mathbf{Y}^* is a function of \mathbf{Y} . The second inequality follows from first noting that

$$H(\mathbf{U}|\mathbf{Y}^*) \leq H(\mathbf{Y}, S_0, S_N, \mathbf{U}|\mathbf{Y}^*).$$

And then observing that

$$\begin{aligned} H(\mathbf{Y}, S_0, S_N, \mathbf{U}|\mathbf{Y}^*) &= H(\mathbf{Y}|\mathbf{Y}^*) + H(S_0, S_N|\mathbf{Y}, \mathbf{Y}^*) + H(\mathbf{U}|\mathbf{Y}, \mathbf{Y}^*, S_0, S_N) \\ &\stackrel{(a)}{=} H(\mathbf{Y}|\mathbf{Y}^*) + H(S_0, S_N|\mathbf{Y}, \mathbf{Y}^*) + H(\mathbf{U}|\mathbf{Y}, S_0, S_N) \\ &\stackrel{(b)}{\leq} H(\mathbf{Y}|\mathbf{Y}^*) + 2 \log_2 |\mathcal{S}| + H(\mathbf{U}|\mathbf{Y}, S_0, S_N) \\ &\stackrel{(c)}{\leq} 2 \log_2(N+1) + 2 \log_2 |\mathcal{S}| + H(\mathbf{U}|\mathbf{Y}, S_0, S_N), \end{aligned}$$

where (a) follows from \mathbf{Y}^* being a function of \mathbf{Y} , (b) follows by S_0 and S_N each having a support of size $|\mathcal{S}|$, and (c) follows since in order to construct \mathbf{Y} from \mathbf{Y}^* , it suffices to be told how many ‘0’ symbols have been trimmed from each side of \mathbf{Y} , and both numbers are always between 0 and N . Combining the above two displayed equations yields the RHS of (56).

Finally, we divide both sides of (56) by N and take the limit as $N \rightarrow \infty$. Since the left-most and right-most terms converge to $E(\hat{H}_\infty)$, the sandwich property implies that the center term, $E(H_n^*)$ also converges to this quantity.

Part IV: (54a)=(54b)=(54c) and (55a)=(55b)=(55c): Note that, for $1 \leq i \leq N$, we have

$$H(U_i|U_1^{i-1}, \mathbf{Y}, S_0, S_N) \leq H(U_i|U_1^{i-1}, \mathbf{Y}) \leq H(U_i|U_1^{i-1}, \mathbf{Y}^*).$$

We have already proved that (54a)=(54c) and (55a)=(55c). Thus, by the sandwich property, (54a)=(54b)=(54c) and (55a)=(55b)=(55c). ■

Remark 1. We end this section by remarking in passing that the ideas used up to this point can be easily generalized to settings in which the channel introduces insertions and substitutions, in addition to deletions. That is, representing such outcomes via a trellis (recall Figure 2) is straightforward: to accommodate substitutions we only need to change the probabilities associated with the diagonal edges we have previously defined, and an insertion would correspond to an edge going one column across and two rows down. Theorem 20 would continue to hold for the more general case, as well as Theorem 2. This is because, in both, we have not made specific use of the fact that the channel only introduces deletions⁹. Indeed, the only essential complication we had to overcome was that of synchronization. That is, we were not told which part of the output corresponds to which part of the input. The next section, in which we prove strong polarization, is different in this regard. Proving strong polarization for the more general case requires different tools and ideas, and is done in the recent paper [28].

VII. STRONG POLARIZATION

To rigorously claim a coding scheme for the deletion channel, one must also show strong polarization. For this, Theorem 20 is not sufficient and, so far, we have been unable to prove strong polarization for the *standard* polar code construction. Thus, we will modify the standard coding scheme to proceed.

⁹For this statement to be correct, we must remove (54c) and (55c) from Theorem 20. However, the claim still holds if we do not remove these lines.

A. Overview of Coding Scheme

Fix a deletion probability δ and a regular hidden Markov input distribution. Recall that our goal is to achieve the information rate \mathcal{I} given in (23). For didactic reasons, we first consider a simplified setting in which this goal is easily attained. Specifically, let N_0 be a given parameter, and consider a block-TDC with block length N_0 and deletion probability δ . That is, for each input block $\mathbf{X}(\phi)$ of length N_0 , where $\phi = 1, 2, \dots$, the channel outputs $\mathbf{Y}^*(\phi)$, which is the result of passing $\mathbf{X}(\phi)$ through a TDC with deletion probability δ . The crucial point to note is that, contrary to a deletion channel, the output of a block-TDC *contains commas between segments*. That is, we know exactly which output segment corresponds to which input block.

How would one code for such a channel and achieve a rate approaching \mathcal{I} ? For this, we will assume that

$$N_0 = 2^{n_0}, \quad (57)$$

and that we can choose N_0 to be arbitrarily large. Let

$$\Phi = 2^{n_1} \quad (58)$$

be the number of blocks we will transmit through the channel. Consider the following input distribution: each block $\mathbf{X}(\phi)$ will be distributed according to the input distribution that we have fixed at the start of this subsection, and the input blocks $\mathbf{X}(1), \mathbf{X}(2), \dots, \mathbf{X}(\Phi)$ will be *i.i.d.* In a nutshell, this suffices to achieve a coding rate of \mathcal{I} with vanishing probability of error for the following two reasons. First, Theorem 20 shows weak polarization for each block and, in each block, we have the required fractions of high-entropy/low-entropy indices. Second, the independence between blocks implies that strong polarization will occur.

We now back the above claim with a few more details. We denote the output of the encoder — the concatenation of the above blocks — by

$$\mathbf{X} = \mathbf{X}(1) \odot \mathbf{X}(2) \odot \dots \odot \mathbf{X}(\Phi). \quad (59)$$

This output has length

$$N = N_0 \cdot \Phi = 2^{n_0+n_1} = 2^n. \quad (60)$$

We will use a sans-serif font to denote a vector whose elements are ‘blocks’. Thus, we will denote the partitioning of the above \mathbf{X} into blocks of length N_0 by

$$\mathbf{X} = (\mathbf{X}(1), \mathbf{X}(2), \dots, \mathbf{X}(\Phi)). \quad (61)$$

The corresponding output of the block-TDC is denoted

$$\mathbf{Y}^* = (\mathbf{Y}^*(1), \mathbf{Y}^*(2), \dots, \mathbf{Y}^*(\Phi)). \quad (62)$$

That is, \mathbf{Y}^* is comprised of Φ distinguishable blocks — it is *not* simply the concatenation of the $\mathbf{Y}^*(\phi)$. The superscript ‘*’ in \mathbf{Y}^* suggest that trimming operation is applied *blockwise*.

We first consider the polar transform of $\mathbf{X}(\phi)$, denoted¹⁰

$$\mathbf{V}(\phi) = \mathcal{A}(\mathbf{X}(\phi)), \quad (63)$$

¹⁰We reserve the letter U , commonly used to denote the result of a polar transform, for a related yet distinct definition that is yet to appear.

where $1 \leq \phi \leq \Phi$. Note that $\mathbf{V}(\phi)$ is a binary vector of length N_0 ,

$$\mathbf{V}(\phi) = (V_1(\phi), V_2(\phi), \dots, V_{N_0}(\phi)).$$

Recall that $\mathbf{Y}^*(\phi)$ is the output corresponding to $\mathbf{X}(\phi)$, and note that since we have assumed that the $\mathbf{X}(\phi)$ are *i.i.d.*, then this must also hold for triplets $(\mathbf{X}(\phi), \mathbf{V}(\phi), \mathbf{Y}^*(\phi))$, when ranging over $1 \leq \phi \leq \Phi$.

For a fixed $1 \leq \phi \leq \Phi$ and a given $1 \leq i_0 \leq N_0$, consider the pair of entropies

$$H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) \quad \text{and} \\ H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)). \quad (64)$$

We now make two important observations. First, since we have already established that the $(\mathbf{X}(\phi), \mathbf{V}(\phi), \mathbf{Y}^*(\phi))$ are *i.i.d.* over ϕ , we deduce that (64) is independent of ϕ . Second, both entropies in (64) exhibit slow polarization, in the sense of Theorem 20. That is, on one hand, we deduce that (54c)=(54d) and (55c)=(55d), if in both (54c) and (55c) we replace $U_i, U_1^{i-1}, \mathbf{Y}^*, n$ and N by $V_{i_0}(\phi), V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi), n_0$ and N_0 , respectively. These statements hold for all $\delta \in [0, 1]$. For the special case of $\delta = 1$, one gets a degenerate channel where $\mathbf{Y}^*(\phi)$ always equals the empty string. Thus, on the other hand, the same claim of (54c)=(54d) and (55c)=(55d), under the above substitutions continues to hold, with \mathbf{Y} and \mathbf{Y}^* removed from these equations.

Since the first entropy in (64) is always less than or equal to the second, we deduce from the above paragraph and the first half of Theorem 20 that for $\epsilon \in (0, 1)$ fixed, the fraction of indices i_0 for which

$$H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) < \epsilon \quad \text{and} \\ H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) \geq \epsilon$$

tends to

$$\left(1 - \lim_{n_0 \rightarrow \infty} \frac{H(\mathbf{X}(\phi)|\mathbf{Y}(\phi))}{N_0}\right) \\ - \left(1 - \lim_{n_0 \rightarrow \infty} \frac{H(\mathbf{X}(\phi))}{N_0}\right) = \mathcal{I},$$

as $n_0 \rightarrow \infty$. For simplicity of exposition, let us further restrict ϵ to $\epsilon \in (0, 1/2)$. By both halves of Theorem 20, we deduce that the fraction of indices i_0 for which

$$\epsilon \leq H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) \leq 1 - \epsilon$$

vanishes. The conclusion is stated as a lemma, for future reference.

Lemma 21. *For $\epsilon \in (0, 1/2)$ fixed, the fraction of indices $1 \leq i_0 \leq N_0$ for which*

$$H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) < \epsilon \quad \text{and} \\ H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) > 1 - \epsilon \quad (65)$$

tends to \mathcal{I} , as $n_0 \rightarrow \infty$, and is the same for every $1 \leq \phi \leq \Phi$.

We now note that for a given ϕ and i_0 , we have an efficient method of calculating the probabilities corresponding to (65). Namely, this is achieved by using the base trellis defined for

a TDC in Subsection III-C, applying a series of plus and minus polarization operations on it, according to the binary representation of $i_0 - 1$, and then invoking (22). That is, the only thing stopping us from applying the Honda-Yamamoto scheme [29] at this point is the fact that the above ϵ is fixed. That is, if we were to follow the Honda-Yamamoto scheme by placing information bits in the V_{i_0} for which (65) holds, then we would not be able to prove that the probability of error decays to zero as the codelength increases. For this, a stronger version of (65) is needed.

Informally, we overcome the above problem as follows. Take ϵ ‘small’ and n_0 as well as n_1 ‘large’. Consider a ‘good’ index i_0 . That is, an index i_0 for which (65) holds. This will be the case for a fraction of indices ‘very close’ to \mathcal{I} . Next, recall the definition of \mathbf{X} in (59), and denote its polar transform as

$$\mathbf{U} = \mathcal{A}(\mathbf{X}) .$$

Consider the subvector $U_{(i_0-1)\cdot\Phi+1}^{i_0\cdot\Phi}$. It is not hard to prove that

$$U_{(i_0-1)\cdot\Phi+1}^{i_0\cdot\Phi} = \mathcal{A}((V_{i_0}(1), V_{i_0}(2), \dots, V_{i_0}(\Phi))) . \quad (66)$$

That is, the LHS of (66) is gotten by applying the Arkan transform to the vector $(V_{i_0}(1), V_{i_0}(2), \dots, V_{i_0}(\Phi))$. Since each entry of this vector satisfies (65), ‘almost all’ indices i of \mathbf{U} , where $(i_0 - 1) \cdot \Phi + 1 \leq i \leq i_0 \cdot \Phi$ are strongly polarized. That is, satisfy

$$Z(U_i|U_1^{i-1}, \mathbf{Y}^*) < 2^{-n_1\beta} \quad \text{and} \quad K(U_i|U_1^{i-1}) < 2^{-n_1\beta} \quad (67)$$

where Z and K are the conditional Bhattacharyya parameter and the conditional total variation (see Definitions 7 and 8 in Appendix A), $\beta < 1/2$ is some fixed constant, and \mathbf{Y}^* is the block-TDC output vector defined in (62). That is, the overall fraction of useful indices $1 \leq i \leq N_0\Phi$ with respect to the Honda-Yamamoto scheme [29] will be ‘very close’ to \mathcal{I} , and the error of the scheme will approach 0 at a rate of roughly $2^{-\sqrt{\Phi}}$.

The reader may not be surprised to learn that the above informal statements can be made rigorous and proven¹¹. Indeed, this will be done as part of the proof of Theorem 1. However, one important point remains to be addressed. That is, the channel we will in fact be coding for is the deletion channel, and *not* the block-TDC. Hence, in the above description, we have implicitly assumed a genie which has manufactured the punctuated vector \mathbf{Y}^* for us. The purpose of the guard-bands, defined shortly, is to approximate such a genie in practice.

Our actual coding scheme will be as follows. For the encoding step, we will first use the Honda-Yamamoto scheme with respect to the block-TDC. Thus, the information bits will be placed in indices j of \mathbf{U} for which (67) holds. The resulting codeword will be \mathbf{X} . Then, we will add to \mathbf{X} runs of ‘0’ symbols in key locations, and transmit the resulting word (which will be longer than \mathbf{X}) on the deletion channel. On the decoder side, a preliminary step will be to deduce the

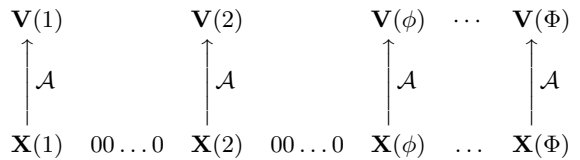


Fig. 4. The $\Phi = N/N_0$ blocks, denoted $\mathbf{X}(1), \mathbf{X}(2), \dots, \mathbf{X}(\Phi)$, have length $N_0 = 2^{n_0}$, are i.i.d., and each is distributed according to the regular hidden-Markov input distribution. Their polar transforms are $\mathbf{V}(1), \mathbf{V}(2), \dots, \mathbf{V}(\Phi)$. An additional $n - n_0$ polarization steps (not shown) will be applied to $\mathbf{V}(1), \mathbf{V}(2), \dots, \mathbf{V}(\Phi)$, resulting in \mathbf{U} . The transmitted codeword is gotten by separating consecutive $\mathbf{X}(\cdot)$ vectors by a ‘guard band’. That is, by a string of ‘0’ symbols. The length of the guard bands is not constant. For example, the middle guard band is always the longest, while the first and last guard bands are always the shortest.

punctuated vector \mathbf{Y}^* from the received vector \mathbf{Y} . That is, we will remove the guard bands (and trim the $\mathbf{Y}(\phi)$ into $\mathbf{Y}^*(\phi)$ in the process), thus producing \mathbf{Y}^* . Then, the decoder will be applied on \mathbf{Y}^* to yield \mathbf{U} , and thus the information bits.

B. Guard bands

In this subsection, we first describe how the guard bands are added to \mathbf{X} on the encoder side. We then explain how the decoder deduces the punctuated vector \mathbf{Y}^* from the received vector \mathbf{Y} .

We start by defining how guard bands are added between the blocks $\mathbf{X}(1), \mathbf{X}(2), \dots, \mathbf{X}(\Phi)$, see Figure 4. That is, we define how \mathbf{X} is transformed into $g(\mathbf{X})$. This is done in a simple recursive manner. Informally, let \mathbf{x} be a vector of length 2^n . If this length is greater than the designated block-length N_0 , we halve \mathbf{x} , add ℓ_n ‘0’ symbols in the middle, and then apply g recursively to each original half. Namely, for $\mathbf{x} = \mathbf{x}_I \odot \mathbf{x}_{II} \in \mathcal{X}^{2^n}$ with

$$\mathbf{x}_I = x_1^{2^{n-1}} \in \mathcal{X}^{2^{n-1}}, \quad \mathbf{x}_{II} = x_{2^{n-1}+1}^{2^n} \in \mathcal{X}^{2^{n-1}}$$

being the first and second halves of \mathbf{x} , respectively, we define

$$g(\mathbf{x}) \triangleq \begin{cases} \mathbf{x} & \text{if } n \leq n_0 \\ g(\mathbf{x}_I) \odot \overbrace{00 \dots 0}^{\ell_n} \odot g(\mathbf{x}_{II}) & \text{if } n > n_0, \end{cases} \quad (68)$$

and

$$\ell_n \triangleq \lfloor 2^{(1-\xi)(n-1)} \rfloor, \quad (69)$$

where $\xi \in (0, 1/2)$ is a yet-to-be-specified ‘small’ constant. The parameter ξ controls the rate penalty of adding guard bands, on one hand, and the probability of the decoder successfully removing the guard bands, on the other hand. We will require that $n_0 > 1$, so that the inequality

$$\ell_n > 2^{(n-1)(1-\xi)-1} \quad (70)$$

used later on will hold for all relevant n , i.e., for $n > n_0$. Note the above specifically implies that $\ell_n > 0$.

We now explain how the guard bands are removed, from the received word \mathbf{Y} , in order to produce the punctuated sequence \mathbf{Y}^* defined in (62). Equivalently, we now show a procedure with the following outcome: for each block index $1 \leq \phi < \Phi$, we will produce the trimmed vector $\mathbf{Y}^*(\phi)$ corresponding to the block $\mathbf{X}(\phi)$. Before explaining how this is done, we first

¹¹Such a proof is not a straightforward adaptation of the ideas in [24] and [30]. Namely, it requires the use of [31, Lemma 40], which we indeed invoke in the proof of Theorem 1.

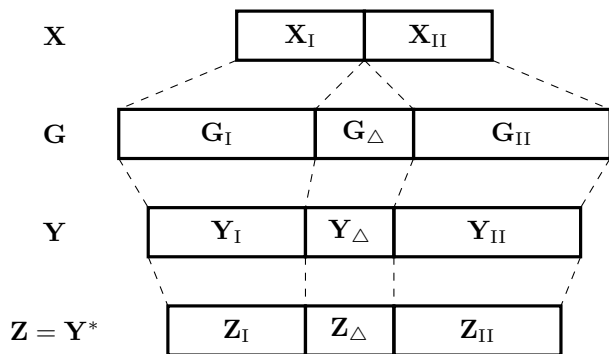


Fig. 5. The random variables X , G , Y , and Z .

mention that our method has a small yet non-zero probability of failing. That is, there is a non-zero probability that our method will fail to produce Y^* . This probability will be analyzed at a later stage.

Our procedure for producing Y^* will have a preliminary step, and will then involve a recursion. The preliminary step is simple: we trim the received vector Y of leading and trailing zeros to produce Y^* . We stress that, generally, Y^* does *not* equal the punctuated sequence Y^* defined in (62). In order to introduce notation required later on, let us now define the above operation more verbosely. Let X_I and X_{II} be the left and right halves of X , see Figure 5. Thus, the transmitted word is $g(X) = G_I \odot G_\Delta \odot G_{II}$, where $G_I = g(X_I)$, $G_{II} = g(X_{II})$, and G_Δ is the middle guard band of length ℓ_n , where n is \log_2 of the length of X . Clearly, G_I and G_{II} are of equal length. Denote the parts of Y corresponding to G_I , G_Δ and G_{II} by Y_I , Y_Δ , and Y_{II} , respectively. Note that at this stage, the decoder sees Y , but can only make an informed guess as to what parts of Y constitute Y_I , Y_Δ , and Y_{II} . We remove from the received word Y all leading and trailing ‘0’ symbols and denote the resulting vector $Z = Y^*$. We denote the parts of Z corresponding to Y_I , Y_Δ , and Y_{II} by Z_I , Z_Δ , and Z_{II} , respectively. In order to build up the reader’s intuition, we note that in a ‘typical case’, Z_I is Y_I after the leading zeros have been removed, Z_{II} is Y_{II} after the trailing zeros have been removed, and Z_Δ is simply Y_Δ . As explained, the production of Z from Y constitutes the preliminary step of our method.

We will now specify how the punctuated vector Y^* is recursively produced from Z . For the base case, note that if $\Phi = 1$, then Y^* is simply Z . Our procedure hinges on the assumption that the middle index of Z originated from a guard band symbol. Specifically, we will assume that the middle index of Z (rounding down) belongs to Z_Δ . As explained, there is a probability of this assumption being false, and this will be analyzed at a later stage. For now, consider the case in which the assumption holds. In this case, the crucial observation is that Y_I^* equals the first half of Z , trimmed, while Y_{II}^* equals the second half of Z , trimmed. Namely, if ζ is the length of Z , then

$$Y_I^* = (Z_1, Z_2, \dots, Z_{\lfloor \zeta/2 \rfloor})^*, \quad (71)$$

$$Y_{II}^* = (Z_{\lfloor \zeta/2 \rfloor + 1}, Z_{\lfloor \zeta/2 \rfloor + 2}, \dots, Z_\zeta)^*, \quad (72)$$

since the guard band Z_Δ has been ‘trimmed out’. Thus, we

have reduced our original problem of producing Y^* from Y^* into two equivalent problems, each half the size of the original: find the first half of Y^* , namely $Y^*(1), Y^*(2), \dots, Y^*(\Phi/2)$, from Y_I^* and the second half of Y^* from Y_{II}^* . Thus, we continue recursively: we apply our method first to the RHS (71) and then to the RHS of (72). If, during all these recursive invocations, our assumptions on the middle index being part of the middle guard band were indeed correct, then we will have succeeded in producing Y^* . Note that the recursion depth is $n - n_0$.

There are two points that must be addressed. First, recall that adding guard bands makes the transmitted word longer. We must show that this has a vanishingly small effect on the rate of our scheme. Second, we must show that our scheme of producing Y^* from Y has a vanishingly small probability of failing. Once this is done, the proof of Theorem 1 will follow easily.

C. Auxiliary lemmas

In this section, we state and prove a number of lemmas key to the proof of Theorem 1.

In the sequel, we will choose a fixed $\nu \in (0, \frac{1}{3}]$ and set $n_0 = \lfloor \nu n \rfloor$. The parameter ν will trade-off reliability and decoding complexity (e.g., see Theorem 1). Recall that both ξ , the parameter through which ℓ_n is defined in (69), and ν are positive and fixed (not a function of n). Thus, the following lemma ensures that the rate penalty of adding guard bands is negligible as $n \rightarrow \infty$.

Lemma 22. *Let x be a vector of length $|x| = 2^n$. Then,*

$$|x| \leq |g(x)| < \left(1 + \frac{2^{-(\xi \cdot n_0 + 1)}}{1 - 2^{-\xi}}\right) \cdot |x|. \quad (73)$$

Proof: From the definition of $g(x)$, induction shows

$$|g(x)| = \begin{cases} 2^n & \text{if } n \leq n_0 \\ 2^n + \sum_{t=n_0+1}^n 2^{n-t} \cdot \ell_t & \text{otherwise.} \end{cases} \quad (74)$$

Thus, the lower bound in (73) is trivial, since $|x| = 2^n$, and every term in the sum in (74) is non-negative, by (69). The upper bound in (73) is trivially true for $n \leq n_0$. For the case $n > n_0$, we have that

$$\begin{aligned} |g(x)|/|x| &\stackrel{(a)}{=} 1 + \sum_{t=n_0+1}^n 2^{-t} \cdot \ell_t \\ &\stackrel{(b)}{\leq} 1 + \sum_{t=n_0+1}^n 2^{-t} \cdot 2^{(1-\xi) \cdot (t-1)} \\ &= 1 + \sum_{t=n_0+1}^n 2^{-\xi \cdot (t-1) - 1} \\ &< 1 + \sum_{t=n_0+1}^{\infty} 2^{-\xi \cdot (t-1) - 1} \\ &\stackrel{(c)}{=} 1 + \frac{2^{-(\xi \cdot n_0 + 1)}}{1 - 2^{-\xi}}. \end{aligned}$$

where (a) follows from $|x| = 2^n$ and (74); (b) follows from (69); (c) is simply the sum of geometric series. ■

A key idea enabling the ‘genie’ described earlier is the recursive processing of each half of the received sequence. This processing will be successful if the middle symbol of the received sequence is a ‘0’ originating from the outermost guard band, as per the recursive definition in (68). The following lemma shows that this is indeed the case, with very high probability.

Lemma 23. *Let the guard-band length ℓ_n in (69) use a fixed $\xi \in (0, 1/2)$. Fix the channel deletion probability δ and a regular hidden-Markov input distribution. Let $n > n_0 > 1$ and let \mathbf{X} be a random vector of length $N = 2^n$ distributed according to the modified input distribution described above: i.i.d. blocks of length $N_0 = 2^{n_0}$, each distributed according to the specified input distribution. Denote by \mathbf{Y} the result of transmitting $g(\mathbf{X})$ through the deletion channel. Then, there exists a constant $\theta > 0$, dependent only on the input distribution and the deletion probability such that, for n_0 large enough, the probability that the middle symbol of \mathbf{Y}^* (rounding down) is not a ‘0’ from the outer guard band of length ℓ_n is at most $2^{-\theta \cdot 2^{(1-2\xi)n_0}}$.*

Proof: Let $\mathbf{G} = g(\mathbf{X})$ (see Fig. 5). Recall that we denote the first and second halves of \mathbf{X} by \mathbf{X}_I and \mathbf{X}_{II} , respectively. Let $\mathbf{G}_I = g(\mathbf{X}_I)$ and $\mathbf{G}_{II} = g(\mathbf{X}_{II})$, and denote by \mathbf{G}_Δ the guard band comprised of ℓ_n ‘0’ symbols between \mathbf{G}_I and \mathbf{G}_{II} . Hence, by (68),

$$\mathbf{G} = \mathbf{G}_I \odot \mathbf{G}_\Delta \odot \mathbf{G}_{II}.$$

Denote by \mathbf{Y} the (untrimmed) result of passing \mathbf{G} through the deletion channel. Let \mathbf{Y}_I , \mathbf{Y}_{II} , and \mathbf{Y}_Δ be the parts of \mathbf{Y} corresponding to \mathbf{G}_I , \mathbf{G}_{II} , and \mathbf{G}_Δ , respectively. Let $\mathbf{Z} = \mathbf{Y}^*$ be the trimmed \mathbf{Y} . Define \mathbf{Z}_I , \mathbf{Z}_{II} , and \mathbf{Z}_Δ , as the parts of \mathbf{Z} corresponding to \mathbf{G}_I , \mathbf{G}_{II} , and \mathbf{G}_Δ , respectively.

For $\mathbf{Z} = (Z_1, Z_2, \dots, Z_t)$ with $t \geq 1$, the middle index of \mathbf{Z} (rounding down) is $s = \lfloor (t+1)/2 \rfloor$. A sufficient condition for Z_s belonging to \mathbf{Z}_Δ is

$$|\mathbf{Z}_I| < |\mathbf{Z}_\Delta| + |\mathbf{Z}_{II}|, \quad |\mathbf{Z}_{II}| < |\mathbf{Z}_I| + |\mathbf{Z}_\Delta|. \quad (75)$$

To see that this is sufficient, we observe that $|\mathbf{Z}_I| < |\mathbf{Z}_\Delta| + |\mathbf{Z}_{II}|$ implies that the middle index does not fall in \mathbf{Z}_I because then

$$\begin{aligned} \lfloor (|\mathbf{Z}| + 1)/2 \rfloor &= \lfloor (|\mathbf{Z}_I| + |\mathbf{Z}_\Delta| + |\mathbf{Z}_{II}| + 1)/2 \rfloor \\ &\geq \lfloor (|\mathbf{Z}_I| + |\mathbf{Z}_I| + 2)/2 \rfloor = |\mathbf{Z}_I| + 1. \end{aligned}$$

Similarly, if $|\mathbf{Z}_{II}| < |\mathbf{Z}_I| + |\mathbf{Z}_\Delta|$, then the middle index does not fall in \mathbf{Z}_{II} because then

$$\begin{aligned} \lfloor (|\mathbf{Z}| + 1)/2 \rfloor &= \lfloor (|\mathbf{Z}_I| + |\mathbf{Z}_\Delta| + |\mathbf{Z}_{II}| + 1)/2 \rfloor \\ &\leq \lfloor (|\mathbf{Z}_I| + |\mathbf{Z}_\Delta| + |\mathbf{Z}_I| + |\mathbf{Z}_\Delta|)/2 \rfloor = |\mathbf{Z}_I| + |\mathbf{Z}_\Delta|. \end{aligned}$$

Now, we will analyze the probability of (75). Denote by α , β , and γ the following length differences between the three parts of \mathbf{G} and the three corresponding parts of \mathbf{Y} ,

$$\begin{aligned} \alpha &= |\mathbf{G}_I| - |\mathbf{Y}_I|, \\ \beta &= |\mathbf{G}_\Delta| - |\mathbf{Y}_\Delta|, \\ \gamma &= |\mathbf{G}_{II}| - |\mathbf{Y}_{II}|. \end{aligned}$$

Also, denote by α' , β' , and γ' the length differences resulting from trimming,

$$\begin{aligned} \alpha' &= |\mathbf{Y}_I| - |\mathbf{Z}_I|, \\ \beta' &= |\mathbf{Y}_\Delta| - |\mathbf{Z}_\Delta|, \\ \gamma' &= |\mathbf{Y}_{II}| - |\mathbf{Z}_{II}|. \end{aligned}$$

Suppose that the trimming on both sides stopped short of the guard band. In this case, $\beta' = 0$. Since $|\mathbf{G}_I| = |\mathbf{G}_{II}|$ and $|\mathbf{G}_\Delta| = \ell_n$, condition (75) would reduce to

$$\alpha + \alpha' < \gamma + \gamma' + \ell_n - \beta, \quad (76)$$

$$\gamma + \gamma' < \alpha + \alpha' + \ell_n - \beta. \quad (77)$$

Our aim is to show that, with very high probability, both (76) and (77) hold, as well as the assumption leading to their formulation.

Recall that δ is the channel deletion probability and let

$$\hat{\ell} = \ell_n \cdot (1 - \delta)/2. \quad (78)$$

We define the following ‘good’ events on the random variables α , α' , β , β' , γ , and γ' :

$$A : \delta|\mathbf{G}_I| - \hat{\ell}/4 < \alpha < \delta|\mathbf{G}_I| + \hat{\ell}/4 \quad (79)$$

$$A' : 0 \leq \alpha' < \hat{\ell}/4 \quad (80)$$

$$B : 0 \leq \beta < \delta \cdot \ell_n + \hat{\ell} \quad (81)$$

$$B' : \beta' = 0 \quad (82)$$

$$C : \delta|\mathbf{G}_{II}| - \hat{\ell}/4 < \gamma < \delta|\mathbf{G}_{II}| + \hat{\ell}/4 \quad (83)$$

$$C' : 0 \leq \gamma' < \hat{\ell}/4 \quad (84)$$

First, we note that the total number of symbols deleted or trimmed from \mathbf{G}_I is given by $|\mathbf{G}_I| - |\mathbf{Z}_I| = \alpha + \alpha'$. If A and A' hold, then this is bounded by

$$\begin{aligned} \alpha + \alpha' &< \delta|\mathbf{G}_I| + \hat{\ell}/4 + \hat{\ell}/4 \\ &= \delta|\mathbf{G}_I| + \hat{\ell}/2. \end{aligned} \quad (85)$$

By (73), $|\mathbf{G}_I| = 2^{n-1} + t$, where $t \geq 0$. We now show that if A and A' hold, then $\alpha + \alpha' < |\mathbf{G}_I|$. Indeed, by (69) and (78),

$$\begin{aligned} \delta|\mathbf{G}_I| + \hat{\ell}/2 &< \delta|\mathbf{G}_I| + \hat{\ell} \\ &= \delta(2^{n-1} + t) + 2^{-1}(1 - \delta)[2^{(1-\xi)(n-1)}] \\ &< \delta(2^{n-1} + t) + (1 - \delta)2^{n-2} \\ &= \delta 2^{n-1} + (1 - \delta)2^{n-2} + \delta t \\ &< 2^{n-1} + \delta t \\ &< 2^{n-1} + t = |\mathbf{G}_I|. \end{aligned}$$

The analogous claim also holds for C , C' , and \mathbf{G}_{II} . Thus, if A , A' , C , and C' hold, then some parts of \mathbf{G}_I and \mathbf{G}_{II} must remain in \mathbf{Z}_I and \mathbf{Z}_{II} after deletion and trimming. Hence, the trimming has stopped short of the guard band, which implies B' .

If, in addition, B occurs, then both (76) and (77) must also hold. To verify that (76) holds, note that

$$\begin{aligned} \gamma + \gamma' + \ell_n - \beta &\stackrel{(a)}{>} \delta|\mathbf{G}_{II}| - \hat{\ell}/4 + \ell_n - \delta \cdot \ell_n - \hat{\ell} \\ &= \delta|\mathbf{G}_{II}| - \hat{\ell}/4 + (1 - \delta)\ell_n - \hat{\ell} \\ &\stackrel{(b)}{=} \delta|\mathbf{G}_{II}| - \hat{\ell}/4 + 2\hat{\ell} - \hat{\ell} \end{aligned}$$

$$\begin{aligned}
&= \delta |\mathbf{G}_{\text{II}}| + 3\hat{\ell}/4 \\
&\stackrel{(c)}{>} \delta |\mathbf{G}_{\text{II}}| + \hat{\ell}/2,
\end{aligned}$$

where (a) follows from (82), (83), and (84); (b) follows from (78); (c) follows since ℓ_n is positive, by (70), and thus so is $\hat{\ell}$, by (78). Next, observe that $|\mathbf{G}_{\text{I}}| = |\mathbf{G}_{\text{II}}|$, and apply (85). The proof of (77) is the same except that the upper and lower bounds are swapped for $\alpha + \alpha'$ and $\gamma + \gamma'$.

To recap, the occurrence of all the ‘good’ events in (79)–(84) implies that the middle index falls inside \mathbf{Z}_{Δ} . Hence, the next step is to show that each of the above events occurs with very high probability, if n is large enough.

We now recall Hoeffding’s bound [32, Theorem 2] [33, proof of Lemma 4.13] and apply it to the deletion channel with deletion probability δ . Namely, let D be a random variable equal to the number of deletions after N channel uses. Hence, $E[D] = \delta N$, and for $t \geq 0$ we have by Hoeffding’s bound that

$$\Pr(D \geq \delta N + t) \leq e^{-2t^2/N}, \quad (86)$$

$$\Pr(D \leq \delta N - t) \leq e^{-2t^2/N}. \quad (87)$$

Recalling that $\xi > 0$, we now require that n_0 be large enough that the bracketed term in (73) is at most 2. That is, we assume that n_0 is large enough such that, for $n > n_0$, we have

$$|\mathbf{G}_{\text{I}}| \leq 2 \cdot 2^{n-1}. \quad (88)$$

Applying both (86) and (87), we deduce that, for $n > n_0$, we have

$$\begin{aligned}
1 - \Pr(A) &\leq 2e^{-2(\hat{\ell}/4)^2/|\mathbf{G}_{\text{I}}|} \\
&= 2e^{-2(\ell_n(1-\delta)/8)^2/|\mathbf{G}_{\text{I}}|} \\
&\stackrel{(a)}{<} 2e^{-2(2^{(n-1) \cdot (1-\xi)^{-1}}(1-\delta)/8)^2/|\mathbf{G}_{\text{I}}|} \\
&\stackrel{(b)}{\leq} 2e^{-2(2^{(n-1) \cdot (1-\xi)^{-1}}(1-\delta)/8)^2/(2 \cdot 2^{n-1})} \\
&= 2e^{-\left(\frac{(1-\delta)^2}{256}\right) \cdot 2^{(n-1)(1-2\xi)}} \\
&\stackrel{(c)}{\leq} 2e^{-\left(\frac{(1-\delta)^2}{256}\right) \cdot 2^{n_0 \cdot (1-2\xi)}}, \quad (89)
\end{aligned}$$

where (a) follows from (70); (b) holds by (88); and (c) follows from $n > n_0$. Exactly the same bound applies to $1 - \Pr(C)$. For $\Pr(B)$, we again use (86) to deduce that

$$\begin{aligned}
1 - \Pr(B) &\leq e^{-2\hat{\ell}^2/\ell_n} \\
&\stackrel{(a)}{=} e^{-2\left(\frac{\ell_n(1-\delta)}{2}\right)^2/\ell_n} \\
&= e^{-2\left(\frac{(1-\delta)}{2}\right)^2 \cdot \ell_n} \\
&\stackrel{(b)}{<} e^{-2\left(\frac{(1-\delta)}{2}\right)^2 \cdot 2^{(n-1)(1-\xi)^{-1}}} \\
&= e^{-\left(\frac{(1-\delta)^2}{4}\right) \cdot 2^{(n-1)(1-\xi)}} \\
&\stackrel{(c)}{\leq} e^{-\left(\frac{(1-\delta)^2}{4}\right) \cdot 2^{n_0 \cdot (1-\xi)}}, \quad (90)
\end{aligned}$$

where (a) follows from (78); (b) follows from (70); and (c) holds because $n > n_0$.

We now bound $1 - \Pr(A' \cap C')$ from above. Consider \mathbf{G}_{I} and \mathbf{Y}_{I} first. Next, recall that by the recursive definition of g in (68), the prefix of length $N_0 = 2^{n_0}$ of \mathbf{G}_{I} is distributed according to the underlying regular Markov input distribution (it does not contain a guard band). Denote this prefix as X_1, X_2, \dots, X_{N_0} ,

and denote the state of the process at time 0 as S_0 . Since our input distribution is not degenerate, there exists an integer $\tau > 0$ and a probability $0 < p < 1$ such that, for any $s \in \mathcal{S}$,

$$\Pr((X_1, X_2, \dots, X_{\tau}) = (0, 0, \dots, 0) | S_0 = s) < p. \quad (91)$$

Let

$$\tilde{\ell} = \ell_{n_0+1} \cdot (1 - \delta)/2. \quad (92)$$

Since $n > n_0$, we have by (69) and (78) that $\tilde{\ell} \leq \hat{\ell}$ and that

$$\tilde{\ell}/4 < 2^{n_0}.$$

Let

$$\rho = \tau \cdot \left\lfloor \frac{\tilde{\ell}/4}{\tau} \right\rfloor,$$

and partition $X_1, X_2, \dots, X_{\rho}$ into consecutive segments of length τ . Then, we define event A'' to occur if there exists a segment that is not an all-zero vector of length τ , and its first non-zero entry has not been deleted. We define C'' as the analogous event, with respect to \mathbf{G}_{II} and \mathbf{Y}_{II} , the only difference being that we are now considering the length ρ suffix of \mathbf{X}_{II} , and considering the last non-zero entry of a segment. By construction, if A'' and C'' hold, then A' and C' must hold. That is, if event A'' occurs, then the number of symbols trimmed from the left of \mathbf{G}_{I} is strictly less than $\tilde{\ell}/4$, since the above non-zero non-deleted symbol is not trimmed, and this assures that the ‘trimming from the left’ stops before it. A similar claim holds with respect to C'' . Thus, $1 - \Pr(A' \cap C') \leq 1 - \Pr(A'' \cap C'')$.

Since (91) holds for all $s \in \mathcal{S}$, we have by the Markov property that

$$1 - \Pr(A'') < (1 - (1-p)(1-\delta))^{\rho/\tau}. \quad (93)$$

Indeed, if A'' does not hold, this means that we have ‘failed’ on each of the ρ/τ segments, in the sense that each such segment was either all-zero, or its first non-zero symbol was deleted. Since the probability of ‘success’ conditioned on any given string of past failures is always greater than $(1-p)(1-\delta)$, the above follows.

Define

$$\zeta = -\log_e(1 - (1-p)(1-\delta)),$$

and note that $\zeta > 0$. Next, we bound ρ as

$$\begin{aligned}
\rho &> \tilde{\ell}/4 - \tau \\
&= \ell_{n_0+1} \cdot (1-\delta)/8 - \tau \\
&> \left(2^{(1-\xi) \cdot n_0 - 1}\right) \cdot (1-\delta)/8 - \tau,
\end{aligned}$$

where the second inequality follows from (70). Thus,

$$1 - \Pr(A'') < e^{-\frac{\zeta}{\tau} \left(2^{(1-\xi) \cdot n_0 - 1} \cdot (1-\delta)/8 - \tau\right)}.$$

Of course, exactly the same bound holds for $1 - \Pr(C'')$. Hence, by the union bound, and recalling that $A'' \cap C''$ implies $A' \cap C'$, we have that

$$1 - \Pr(A' \cap C') < 2e^{-\frac{\zeta}{\tau} \left(2^{(1-\xi) \cdot n_0 - 1} \cdot (1-\delta)/8 - \tau\right)}. \quad (94)$$

Putting (89), (90), and (94) together, and applying the union bound proves the lemma. \blacksquare

We conclude this section with the proof of our main theorem. Note that both the encoding and decoding schemes are specified in the proof.

Proof of Theorem 1: Our proof is divided into two parts. In the first part, we consider the ‘idealized’ random vectors \mathbf{X} and \mathbf{Y} . That is, \mathbf{X} is drawn from the probability distribution defined in Lemma 23 (there is no encoding of date) and \mathbf{Y} is the result of transmitting $g(\mathbf{X})$ through our deletion channel. We will show that by previously proven lemmas, the rate penalty of expanding \mathbf{X} to $g(\mathbf{X})$ is negligible and the probability of deducing \mathbf{Y}^* from \mathbf{Y} is very high. We conclude the first part by discussing the polarization of $\mathbf{U} = \mathcal{A}(\mathbf{X})$.

In the second part of the proof, we consider the actual case at hand. That is, we show how encoding and decoding are carried out, discuss the encoding and decoding complexity, prove that the rate of our coding scheme approaches the information rate \mathcal{I} , and prove that the probability of misdecoding tends to 0.

Recall that $0 < \nu' < \nu \leq 1/3$ are fixed parameters. We let

$$n_0 = \lfloor \nu n \rfloor \quad (95)$$

and

$$\nu'' = \frac{\nu + \nu'}{2}, \quad (96)$$

implying that

$$0 < \nu' < \nu'' < \nu \leq \frac{1}{3}. \quad (97)$$

Then, set ξ for the guard-band length ℓ_n defined in (69) to

$$\xi = \frac{1 - \frac{1+\nu''/\nu}{2}}{2} = \frac{1 - \nu''/\nu}{4}. \quad (98)$$

Note that by (60),

$$n_1 = n - \lfloor \nu n \rfloor = \lceil (1 - \nu)n \rceil. \quad (99)$$

We start with the first part of the proof: let \mathbf{X} and \mathbf{Y} be defined as in Lemma 23 (as yet, no coding of information).

Sub-claim 1. *The rate penalty incurred by adding guard bands becomes negligible as $n \rightarrow \infty$. Namely, $|g(\mathbf{X})|/|\mathbf{X}|$ tends to 1 as $n \rightarrow \infty$.*

Sub-proof: This follows by Lemma 22, which shows that the rate penalty incurred by adding guard bands becomes negligible as $n_0 \rightarrow \infty$, and the connection between n_0 and n given in (95). \square

Sub-claim 2. *The probability of making a mistake during the partitioning of \mathbf{Y} into the $\Phi = 2^{n-n_0}$ trimmed blocks $\mathbf{Y}(1)^*$, $\mathbf{Y}(2)^*$, ..., $\mathbf{Y}(\Phi)^*$ is less than $\frac{1}{3} \cdot 2^{-2^{\nu''n}}$, for $N = 2^n$ large enough.*

Sub-proof: This follows from Lemma 23 and the union bound. Specifically, recalling the recursive nature of our algorithm to produce \mathbf{Y}^* , we note that an error is made only if the relevant portion of the received vector \mathbf{Y} , after that portion has been trimmed, is such that the middle symbol (rounding down) does not belong to the outermost guard band. Each such probability can be bounded by using Lemma 23. Since we produce Φ blocks, our recursion is applied $\Phi - 1$ times. Hence, for n_0 large enough, the probability of failing to produce \mathbf{Y}^* is at most

$$\begin{aligned} & (\Phi - 1) \cdot 2^{-\theta \cdot 2^{(1-2\xi)n_0}} \\ &= (2^{n-\lfloor \nu n \rfloor} - 1) \cdot 2^{-\theta \cdot 2^{\lfloor \nu n \rfloor \cdot ((1+\nu''/\nu)/2)}}, \quad (100) \end{aligned}$$

where the equality follows from (58) and (95)–(99). Recalling (95), we may take n large enough such that n_0 is indeed large enough for the above to hold. Moreover, since $0 < \nu'' < \nu$, it is straightforward to show that the RHS of (100) is less than $\frac{1}{3} \cdot 2^{-2^{\nu''n}}$ for large enough n , as required. \square

Sub-claim 3. *The fraction of indices i for which U_i contains an information bit tends to \mathcal{I} , as $n \rightarrow \infty$. For $\mathbf{U} = \mathcal{A}(\mathbf{X})$, these indices i are defined as satisfying the following two inequalities: the Bhattacharyya parameter must satisfy*

$$Z(U_i|U_1^{i-1}, \mathbf{Y}(1)^*, \mathbf{Y}(2)^*, \dots, \mathbf{Y}(\Phi)^*) < \frac{1}{3N} \cdot 2^{-2^{\nu''n}} \quad (101)$$

and the total variation parameter (see Definition 8 in the appendix) must satisfy

$$K(U_i|U_1^{i-1}) < \frac{1}{3N} \cdot 2^{-2^{\nu''n}}. \quad (102)$$

Sub-proof: Informally, $H \approx 0$ iff $Z \approx 0$ and $H \approx 1$ iff $K \approx 0$. Specifically, by [21, Lemma 1] we have

$$\begin{aligned} & (Z(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)))^2 \\ & \leq H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) \\ & \leq Z(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) \end{aligned}$$

and

$$\begin{aligned} & 1 - H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) \\ & \leq K(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) \\ & \leq \sqrt{1 - (H(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)))^2}. \end{aligned}$$

Thus, Lemma 21 continues to hold if we replace (65) by the condition

$$\begin{aligned} & Z(V_{i_0}(\phi)|V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)) < \epsilon \quad \text{and} \\ & K(V_{i_0}(\phi)|V_1^{i_0-1}(\phi)) < \epsilon. \quad (103) \end{aligned}$$

That is, at the end of n_0 polarization stages, the fraction of indices $1 \leq i_0 \leq N_0$ satisfying the ‘weak polarization’ in (103) tends to \mathcal{I} for any $\epsilon > 0$. To get from the ‘weak polarization’ implied by (103) to the ‘strong polarization’ implied by (101) and (102), we will employ [31, Lemma 40], which we now state.

Lemma 24. [31, Lemma 40] *Let B_1, B_2, \dots be independent and identically distributed random variables with $\Pr\{B_i = 0\} = \Pr\{B_i = 1\} = 1/2$. Let Z_0, Z_1, \dots be a $[0, 1]$ -valued random process that satisfies*

$$Z_{n+1} \leq \kappa \cdot \begin{cases} Z_n^2, & B_{n+1} = 0, \\ Z_n, & B_{n+1} = 1, \end{cases} \quad n \geq 0,$$

for some $\kappa > 1$. Fix $0 < \beta < 1/2$. Then, for every $\delta > 0$ there exist $\eta > 0$ and n_0 such that if $Z_0 \leq \eta$ then

$$\Pr\{Z_n \leq 2^{-2^{n\beta}} \text{ for all } n \geq n_0\} \geq 1 - \delta.$$

Crucially, η and n_0 depend on the process Z_n only through κ .

We will use the above lemma in a setting in which the initial Z_0 is “close enough” to 0, and thus the lemma implies that with high probability $Z_n \leq 2^{-2^{n\beta}}$. As we will see, the lemma will be used both with respect to the evolution of the Bhattacharyya parameter, as well as the evolution of the total variation.

For $\mathbf{b} = (b_1, b_2, \dots, b_n)$, recall from (4) the definition of $i(\mathbf{b})$, and denote

$$i_0(\mathbf{b}) \triangleq 1 + \sum_{j=1}^{n_0} b_j 2^{n_0-j}.$$

Thus, we may think of the random process by which $i(B_1, B_2, \dots, B_n)$ is chosen as first selecting i_0 , which is in fact a function of B_1, B_2, \dots, B_{n_0} , and then completing the choice of i according to a new process $\tilde{B}_1, \tilde{B}_2, \dots, \tilde{B}_{n_1}$, where

$$\tilde{B}_1 = B_{n_0+1}, \tilde{B}_2 = B_{n_0+2}, \dots, \tilde{B}_{n_1} = B_{n_0+n_1}, \quad (104)$$

recalling that $n_0 + n_1 = n$, by (95) and (99).

Fix $\epsilon > 0$ to a value that will shortly be specified. Next, for now, let us fix an index i_0 for which (103) holds. We define two processes related to (104), denoted $\tilde{Z}_1, \tilde{Z}_2, \dots, \tilde{Z}_{n_1}$ and $\tilde{K}_1, \tilde{K}_2, \dots, \tilde{K}_{n_1}$. Recall that by definition, the $(\mathbf{X}(\phi), \mathbf{Y}(\phi))$ are i.i.d. over $1 \leq \phi \leq \Phi$. Hence, this must also be the case for $(V_{i_0}(\phi), V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi))$, by (63). The first process is the evolution of the conditional Bhattacharyya parameter as we apply the n_1 polar transforms implied by (104), to $(\tilde{X}_\phi, \tilde{Y}_\phi)_{\phi=1}^\Phi$, where

$$\tilde{X}_\phi = V_{i_0}(\phi) \text{ and } \tilde{Y}_\phi = (V_1^{i_0-1}(\phi), \mathbf{Y}^*(\phi)).$$

The second process is defined similarly, but now we consider the evolution of the conditional total variation parameter as we apply n_1 polar transforms to $(\tilde{X}_\phi, \tilde{Y}_\phi)_{\phi=1}^\Phi$, where

$$\tilde{Y}_\phi = V_1^{i_0-1}(\phi).$$

By our assumption of i_0 satisfying (103),

$$\tilde{Z}_1 = Z(\tilde{X}_1|\tilde{Y}_1) < \epsilon \text{ and } \tilde{K}_1 = K(\tilde{X}_1|\tilde{Y}_1) < \epsilon.$$

Since $(\tilde{X}_\phi, \tilde{Y}_\phi)$ are i.i.d. over ϕ , and the same holds for $(\tilde{X}_\phi, \tilde{Y}_\phi)$, we have by [24, Proposition 5] that

$$\tilde{Z}_{t+1} \leq \begin{cases} 2\tilde{Z}_t & \text{if } \tilde{B}_t = 0 \\ \tilde{Z}_t^2 & \text{if } \tilde{B}_t = 1 \end{cases}$$

and by [21, Proposition 4] that

$$\tilde{K}_{t+1} \leq \begin{cases} \tilde{K}_t^2 & \text{if } \tilde{B}_t = 0 \\ 2\tilde{K}_t & \text{if } \tilde{B}_t = 1. \end{cases}$$

Lastly, it follows from (66) that \tilde{Z}_{n_1} equals the LHS of (101) while \tilde{K}_{n_1} equals the LHS of (102), where i is defined in (4), with B_j instead of b_j .

To prove the sub-claim, we must show that, for every $\omega > 0$, there exists a threshold such that, if n is larger than the threshold, then the fraction of indices i satisfying both (101) and (102) is at least $\mathcal{I} - \omega$. We will do this by choosing an ϵ and n_0 such that the fraction of indices satisfying (103) is at least $\mathcal{I} - \omega/3$. Of these weakly polarized indices, we will choose n_1

such that at least a fraction $1 - 2\omega/3$ satisfy both (101) and (102). This is sufficient because $(\mathcal{I} - \omega/3)(1 - 2\omega/3) \geq \mathcal{I} - \omega$. To make a proper argument, however, we will work in reverse.

First, we will set the parameters for strong polarization assuming sufficient weak polarization. In particular, we define

$$\beta = 3(\nu + \nu'')/4 \quad (105)$$

and observe that (97) implies $0 < \beta < 1/2$. Then, we let $\psi = \omega/3$ be the maximum fraction of weakly polarized indices that can fail to strongly polarize and apply Lemma 24 to determine a valid maximum for ϵ and minimum for n_1 (in [31], ψ , ϵ , and n_1 are denoted δ , η , and n , respectively). This lemma implies the existence of an $\epsilon > 0$ such that if (103) holds for an index i_0 , then the fraction of i values $((i_0 - 1) \cdot \Phi + 1 \leq i \leq i_0 \cdot \Phi)$ for which both $\tilde{Z}_i < 2^{-2^{\beta n_1}}$ and $\tilde{K}_i < 2^{-2^{\beta n_1}}$ is at least $1 - 2\omega/3$, for all n_1 large enough¹². Conceptually, we need to apply the lemma twice – once for (101) and once for (102). Thus, the fraction of weakly polarized indices that fail to satisfy both (101) and (102) is at most $2\psi = 2\omega/3$.

Next, for the ϵ determined above, we find the minimum n_0 to guarantee that (103) holds for at least a fraction $\mathcal{I} - \omega/3$ of the i_0 indices. Lastly, we recall that n_0 and n_1 are monotonically increasing functions of n , by (95) and (99). Hence, for all large enough n , the parameters n_0 and n_1 will exceed the bounds computed earlier and the fraction of indices satisfying (101) and (102), where in both cases we replace the RHS by $2^{-2^{\beta n_1}}$, is at least $\mathcal{I} - \omega$.

In order to prove the sub-claim, all that remains is to show that, for all large enough n , we have

$$2^{-2^{\beta n_1}} < \frac{1}{3N} \cdot 2^{-2^{\nu'' n}}, \quad (106)$$

the latter term being RHS of (101) and (102). Indeed, by (99) we have that $n_1 \geq (1 - \nu)n$, and recalling from (97) that $\nu \leq 1/3$, we deduce that $n_1 \geq 2n/3$. Hence, to prove (106), it suffices to show that

$$2^{-2^{2\beta n/3}} < \frac{1}{3N} \cdot 2^{-2^{\nu'' n}}. \quad (107)$$

Indeed, by (97) and (105) we have that $2\beta/3 > \nu''$. Thus, recalling that $N = 2^n$, we deduce that (107) holds for all n large enough. \square

We now move to the second part of our proof. Let us first discuss how data is encoded. We produce $\mathbf{u} = u_1^N$ successively, starting from u_1 and ending in u_N . If the current index i satisfies (101) and (102), then u_i is set to an information bit, where the information bits are assumed i.i.d. and Bernoulli(1/2). Otherwise, u_i is randomly picked according to the distribution $P(U_i = u_i | U_1^{i-1} = u_1^{i-1})$, where u_1^{i-1} are the realizations occurring in previous stages¹³. The random picks in this case are assumed to be from a random source common to both the encoder and the decoder. Typically, this is implemented using a pseudo-random number generator, common to both sides: if the pseudo-random number $0 \leq r_i \leq 1$ drawn for

¹²Crucially, ϵ and the n_1 threshold do not depend on the choice of i_0 .

¹³In practice, the user may want to pick some of these u_i deterministically. That is, pick u_i such that $P(U_i = u_i | U_1^{i-1} = u_1^{i-1})$ is maximized, as suggested in [34] and analyzed in [35].

this stage is such that $P(U_i = 0 | U_1^{i-1} = u_1^{i-1}) \leq r_i$, we set $u_i = 0$. Otherwise, we set $u_i = 1$. These are essentially the ‘frozen-bits’ from the seminal paper [24]. Transforming \mathbf{u} to $\mathbf{x} = \mathcal{A}_n^{-1}(\mathbf{u})$ and adding guard bands to \mathbf{x} is as described before.

The following sub-claim proves a key part of our theorem.

Sub-claim 4. *The rate of our coding scheme approaches \mathcal{I} , as $n \rightarrow \infty$.*

Sub-proof: This is an immediate consequence of Subclaims 1 and 3. That is, Subclaim 1 states that no asymptotic rate penalty is incurred by adding guard bands, while Subclaim 3 shows that the fraction of U_i in which we can place information bits converges to \mathcal{I} . \square

Note that the probability distribution of our encoded \mathbf{u} does *not* generally equal that of the random variable \mathbf{U} used throughout this paper. Namely, denote by \tilde{p} the probability distribution corresponding to the above encoding process: the probability of the encoder producing the vector \mathbf{u} is $\tilde{p}(\mathbf{u})$. Next, denote by p the probability distribution of \mathbf{U} . That is, the probability we would get if we were to set u_i to 0 with probability $P(U_i = 0 | U_1^{i-1} = u_1^{i-1})$, irrespective of whether i satisfies (101) and (102) or not. Our plan is to show that the difference between p and \tilde{p} is ‘small’. However, we must first address a subtle point stemming from this difference in distributions. Specifically, the probability $P(U_i = 0 | U_1^{i-1} = u_1^{i-1})$ used at stage i might be undefined, since we might be conditioning on an event with probability 0. In this case, we define the above probability to be 1/2.

We decode as previously explained: we first recursively partition the received vector into $\mathbf{y}(1)^*, \mathbf{y}(2)^*, \dots, \mathbf{y}(\Phi)^*$. Then, we employ successive cancellation decoding. That is, we produce our estimate $\hat{\mathbf{u}} = \hat{u}_1^N$ of \mathbf{u} by first producing \hat{u}_1 , then \hat{u}_2 , etc., up to \hat{u}_N . If index i is such that both (101) and (102) hold, then we set \hat{u}_i to the value maximizing

$$\begin{aligned} P(U_i = \hat{u}_i | U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y}(1)^* = \mathbf{y}(1)^*, \\ \mathbf{Y}(2)^* = \mathbf{y}(2)^*, \dots, \mathbf{Y}(\Phi)^* = \mathbf{y}(\Phi)^*). \end{aligned} \quad (108)$$

Otherwise, if i does not satisfy both (101) and (102), we set \hat{u}_i is accordance with the common randomness. That is, in the pseudo-random number implementation, we set $\hat{u}_i = 0$ if

$$P(U_i = 0 | U_1^{i-1} = \hat{u}_1^{i-1}) \leq r_i. \quad (109)$$

Otherwise, we set $\hat{u}_i = 1$.

We stress that the probabilities in (108) and (109) are calculated according to the probability distribution of the random vector \mathbf{U} used throughout this paper. That is, although \mathbf{u} has been encoded according to the probability \tilde{p} , we decode it ‘as if’ it had been encoded using p . This discrepancy will shortly be addressed. However, as a first step, the following sub-claim considers the case in which there is no discrepancy.

Sub-claim 5. *If \mathbf{u} were chosen according to the probability distribution p , then the probability of misdecoding would be less than $\frac{2}{3} \cdot 2^{-2^{\nu''n}}$, for large enough n .*

Sub-proof: Note that if the above were the case, then \mathbf{u} and \mathbf{U} would have the same probability distribution. Thus,

Subclaim 2 would apply, and would imply that the probability of our partitioning algorithm failing to produce the correct $\mathbf{y}(1)^*, \mathbf{y}(2)^*, \dots, \mathbf{y}(\Phi)^*$ from the received vector would be less than $\frac{1}{3} \cdot 2^{-2^{\nu''n}}$, for large enough n . Also, if a ‘genie’ were to give us the correct $\mathbf{y}(1)^*, \mathbf{y}(2)^*, \dots, \mathbf{y}(\Phi)^*$, we have from (101) that the probability of misdecoding \mathbf{u} would be less than $\frac{1}{3} \cdot 2^{-2^{\nu''n}}$ for large enough n , using exactly¹⁴ the same arguments as given in [24, Proof of Theorem 2] to bound the probability of the successive cancellation decoder failing. The result follows by applying the union bound. \square

For \mathbf{u} such that $p(\mathbf{u}) > 0$, denote by $P_e(\mathbf{u})$ the probability that our decoder fails, given that \mathbf{u} was encoded. Otherwise, if $p(\mathbf{u}) = 0$, define¹⁵ $P_e(\mathbf{u}) = 1$. We have just shown that for large enough n ,

$$\sum_{\mathbf{u} \in \mathcal{X}^N} p(\mathbf{u}) P_e(\mathbf{u}) < \frac{2}{3} \cdot 2^{-2^{\nu''n}}. \quad (110)$$

However, recall that our ultimate goal is to upper bound the LHS, after $p(\mathbf{u})$ is replaced by $\tilde{p}(\mathbf{u})$. Informally, a similar bound holds for this case as well, since p and \tilde{p} are ‘close’. The two following sub-claims makes this statement precise.

Sub-claim 6.

$$\sum_{\mathbf{u} \in \mathcal{X}^N} |\tilde{p}(\mathbf{u}) - p(\mathbf{u})| < \frac{1}{3} \cdot 2^{-2^{\nu''n}}.$$

Sub-proof: We use the following result from [36, Lemma 3.5]:

$$A_1^N - B_1^N = \sum_{i=1}^N B_1^{i-1} (A_i - B_i) A_{i+1}^N$$

where, here, A_i^j denotes the product $A_i^j = A_i \cdot A_{i+1} \cdots A_j$, and $A_1^0 = A_{N+1}^N \triangleq 1$. We now take

$$A_i = A_i(\mathbf{u}) = \tilde{p}(u_i | u_1^{i-1}) \quad \text{and} \quad B_i = B_i(\mathbf{u}) = p(u_i | u_1^{i-1}).$$

Recall that we have defined B_i to be 1/2 if $p(u_1^{i-1}) = 0$. Similarly, we define A_i to be 1/2 if $\tilde{p}(u_1^{i-1}) = 0$. We deduce that

$$\begin{aligned} \sum_{\mathbf{u} \in \mathcal{X}^N} |\tilde{p}(\mathbf{u}) - p(\mathbf{u})| &= \sum_{\mathbf{u} \in \mathcal{X}^N} |A_1^N - B_1^N| \\ &= \sum_{\mathbf{u} \in \mathcal{X}^N} \left| \sum_{i=1}^N B_1^{i-1} (A_i - B_i) A_{i+1}^N \right| \\ &\leq \sum_{\mathbf{u} \in \mathcal{X}^N} \sum_{i=1}^N |B_1^{i-1} (A_i - B_i) A_{i+1}^N| \end{aligned}$$

¹⁴Since [24] considers the Bhattacharyya parameter for the case of a channel with uniform input, we also need to claim that our Z upper bounds the probability of maximum-a-posteriori misdecoding in the more general setting where the channel input is non-uniform. This is well known, see e.g. [21, Remark 1] for a proof of a slightly stronger claim.

¹⁵Note that we are being conservative. We could have simply defined $P_e(\mathbf{u})$ as the probability that our decoder fails, given that \mathbf{u} was encoded. However, if our input distribution is such that some vectors \mathbf{u} are given a probability of 0, say in order to satisfy a constraint on the input, we should treat the event of the encoder producing a \mathbf{u} not satisfying this constraint as an error.

$$= \sum_{i=1}^N \sum_{\mathbf{u} \in \mathcal{X}^N} |B_1^{i-1}(A_i - B_i)A_{i+1}^N|, \quad (111)$$

where the first equality follows by the chain rule and the first inequality follows from the triangle inequality. Next, fix i , and consider the internal sum in (111),

$$\sum_{\mathbf{u} \in \mathcal{X}^N} |B_1^{i-1}(A_i - B_i)A_{i+1}^N|. \quad (112)$$

If i is an index for which both (101) and (102) hold, then $A_i = A_i(\mathbf{u}) = 1/2$ for all \mathbf{u} . For this case, we get from (102) and Lemma 25 in Appendix A that

$$\begin{aligned} & \sum_{\mathbf{u} \in \mathcal{X}^N} |B_1^{i-1}(A_i - B_i)A_{i+1}^N| \\ &= \sum_{\mathbf{u} \in \mathcal{X}^N} B_1^{i-1} \cdot |A_i - B_i| \cdot A_{i+1}^N \\ &= \sum_{\mathbf{u} \in \mathcal{X}^N} p(u_1^{i-1}) \cdot \left| \frac{1}{2} - p(u_i|u_1^{i-1}) \right| \cdot \tilde{p}(u_{i+1}^n|u_1^i) \\ &= \sum_{u_1^i \in \mathcal{X}^i} p(u_1^{i-1}) \cdot \left| \frac{1}{2} - p(u_i|u_1^{i-1}) \right| \sum_{u_{i+1}^n \in \mathcal{X}^{N-i}} \tilde{p}(u_{i+1}^n|u_1^i) \\ &= \sum_{u_1^i \in \mathcal{X}^i} p(u_1^{i-1}) \cdot \left| \frac{1}{2} - p(u_i|u_1^{i-1}) \right| \\ &= K(U_i|U_1^{i-1}) < \frac{1}{3N} \cdot 2^{-2^{\nu''n}}. \end{aligned}$$

Otherwise, if i is an index for which either (101) or (102) do not hold, then $A_i = B_i$ for all \mathbf{u} , and thus (112) equals 0. \square

We are now ready to state our bound on the probability of misdecoding.

Sub-claim 7. *For large enough n ,*

$$\sum_{\mathbf{u} \in \mathcal{X}^N} \tilde{p}(\mathbf{u})P_e(\mathbf{u}) < 2^{-2^{\nu''n}}.$$

Sub-proof: We use the two previous sub-claims as follows,

$$\begin{aligned} \sum_{\mathbf{u} \in \mathcal{X}^N} \tilde{p}(\mathbf{u})P_e(\mathbf{u}) &= \sum_{\mathbf{u} \in \mathcal{X}^N} (p(\mathbf{u}) + \tilde{p}(\mathbf{u}) - p(\mathbf{u}))P_e(\mathbf{u}) \\ &\leq \sum_{\mathbf{u} \in \mathcal{X}^N} (p(\mathbf{u}) + |\tilde{p}(\mathbf{u}) - p(\mathbf{u})|)P_e(\mathbf{u}) \\ &\leq \sum_{\mathbf{u} \in \mathcal{X}^N} p(\mathbf{u})P_e(\mathbf{u}) + \sum_{\mathbf{u} \in \mathcal{X}^N} |\tilde{p}(\mathbf{u}) - p(\mathbf{u})| \\ &< \frac{2}{3} \cdot 2^{-2^{\nu''n}} + \frac{1}{3} \cdot 2^{-2^{\nu''n}}, \end{aligned}$$

which holds for a large enough n . \square

Recall that in the statement of our theorem, we have denoted the length of our codeword (after adding the guard bands) as Λ . The following subclaim proves another key part of our theorem.

Sub-claim 8. *For large enough n , the probability of misdecoding is less than $2^{-\Lambda^{\nu'}}$.*

Sub-proof: This follows by (96), Subclaim 1, and Subclaim 7. \square

All that remains now is to discuss the encoding and decoding complexity of our algorithms. \blacksquare

Sub-claim 9. *The encoding complexity is $O(\Lambda \log \Lambda)$.*

Sub-proof: If we choose to use an input distribution that is not memoryless, then we must use the generalization of Honda-Yamamoto successive cancellation encoding for input distributions with memory [17], [18]. Fortunately, this successive cancellation encoder does not observe the channel output and does not require the deletion channel trellis. Instead, it must recursively track the distribution of the hidden-Markov input. From [17], [18], we know that the encoding complexity is $O(|\mathcal{S}|^3 N \log N)$ where $|\mathcal{S}|$ is the number of states in the hidden-Markov input process. Thus, for a fixed input distribution, the complexity of producing \mathbf{u} and \mathbf{x} is $O(N \log N)$. Note that adding the guard bands is a simple recursive process whose total time is $O(\Lambda)$. Since $\Lambda \geq N$, the result follows. \square

Sub-claim 10. *The decoding complexity is $O(\Lambda^{1+3\nu})$.*

Sub-proof: The complexity of partitioning the received vector \mathbf{y} into the Φ trimmed blocks $\mathbf{y}(1)^*$, $\mathbf{y}(2)^*$, \dots , $\mathbf{y}(\Phi)^*$ is $O(\Lambda)$. Next, consider step i of the decoding algorithm, in which we decide on the value of \hat{u}_i . The key step is to calculate the probability

$$\begin{aligned} P(U_i = 0 | U_1^{i-1} = \hat{u}_1^{i-1}, \mathbf{Y}(1)^* = \mathbf{y}(1)^*, \\ \mathbf{Y}(2)^* = \mathbf{y}(2)^*, \dots, \mathbf{Y}(\Phi)^* = \mathbf{y}(\Phi)^*). \end{aligned}$$

This is done in two stages. Recall (66) and the discussion below it. First, for each $1 \leq \phi \leq \Phi$, we calculate the probabilities

$$P(V_{i_0}(\phi) = 0 | V_1^{i-1}(\phi) = \hat{v}_1^{i-1}(\phi), \mathbf{Y}(\phi)^* = \mathbf{y}(\phi)^*),$$

where i_0 is the unique integer for which

$$(i_0 - 1)\Phi + 1 \leq i \leq i_0\Phi$$

and $\hat{v}_1^{i_0-1}(\phi)$ is related to \hat{u}_1^{i-1} through (66). That is, we have just calculated the probabilities corresponding to the first n_0 polarization stages. Recall that by Subsection III-C, this can be done using Φ trellises. Next, we apply the remaining $n - n_0$ polarization steps to these probabilities. That is, the standard SC decoder is run for the last $n - n_0$ stages, and can be thought of as effectively operating on a code of length $\Phi = 2^{n_1} = 2^{n-n_0}$.

The total running time of the second stage is well known to be $O(\Phi \log \Phi)$, which is indeed $O(\Lambda^{1+3\nu})$. Recalling the discussion in Subsection IV-D, the total running time of the first stage is

$$O(\Phi \cdot |\mathcal{S}|^3 N_0^4),$$

where $|\mathcal{S}|$ is the number of states in the Markov chain through which the input distribution is defined (and which we treat as a constant), $N_0 = 2^{n_0} = 2^{\lfloor n\nu \rfloor}$ and $\Phi = 2^{n-n_0} = 2^{n-\lfloor n\nu \rfloor}$. Since $N = 2^n \leq \Lambda$, the result follows. \square

To summarize, Subclaim 4 shows that the rate of our coding scheme approaches \mathcal{I} . Subclaim 8 shows that the probability of misdecoding is upper bounded by $2^{-\Lambda^{\nu'}}$, for large enough n . Lastly, Subclaims 9 and 10 state that the encoding and decoding complexities are $O(\Lambda \log \Lambda)$ and $O(\Lambda^{1+3\nu})$, respectively. This is what we have set out to prove. \blacksquare

VIII. SIMULATION RESULTS

In this section, we present simulation results for the proposed coding scheme. Our simulation design is based on the freely available code from [37]. While the complexity of our decoder is decreased by the use of guard bands, it still grows like $\Omega(\Lambda^{1+3\nu})$ for some $\nu > 0$. Therefore, we only present simulations for $n = 10$ and $n = 12$ which correspond to $N = 1024$ and $N = 4096$. The results of these simulations are presented in Fig. 6.

From [8], we know that the mutual information rate for uniform inputs is roughly 0.55 and the achieved rates of 0.226 and 0.274 are somewhat below this. Without guard bands, however, the design rates are 0.384 and 0.422. These numbers are closer to the goal of 0.55. Based on this, we believe that modifying the guard band structure is a promising future direction of research. In particular, our guard band structure was chosen to maximize the strong polarization exponent ν . In practice, one might prefer much shorter guard bands (e.g., blocks of $O((\log N)^2)$ zeros) along with a detector that detects these long runs as guard bands. This is possible because the expected value of the longest run outside a guard band is $O(\log N)$. We also conjecture that the application of successive cancellation list decoding [16] will allow one to approach capacity more quickly because of the improved decoding performance, but we leave this to future work.

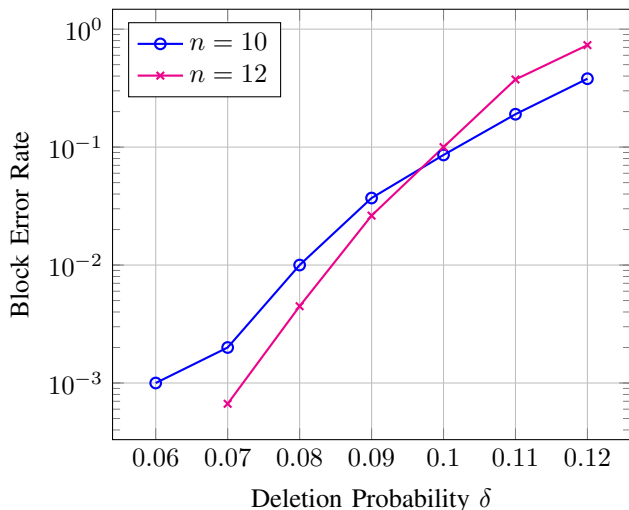


Fig. 6. Simulation results for two codes based on the coding scheme in this paper. The first code has $n = 10$, $n_0 = 6$, and $\xi = 0.2$. It was designed to achieve a block error rate of 0.1 via Monte Carlo using 2000 test blocks with deletion probability $\delta = 0.1$. The result is a code with $K = 394$ information bits, $N = 2^{10} = 1024$, and an overall transmission length of $\Lambda = 1747$ bits (including guard bands), and an overall rate of $394/1747 \approx 0.226$. The second code has $n = 12$, $n_0 = 8$, and $\xi = 0.2$. It was designed to achieve a block error rate of 0.1 via Monte Carlo using 4000 test blocks with deletion probability $\delta = 0.1$. The result is a code with $K = 1729$ information bits, $N = 2^{12} = 4096$, and an overall transmission length of $\Lambda = 6313$ bits (including guard bands), and an overall rate of $1729/6313 \approx 0.274$.

IX. ACKNOWLEDGMENTS

The authors thank Dar Arava and the anonymous reviewers for helpful comments and suggestions.

APPENDIX

A. Conditional Bhattacharyya and Total Variation

In this section we define the conditional Bhattacharyya parameter $Z(X|Y)$ and the conditional total variation $K(X|Y)$. See [21, Section III] for various connections between these and other measures, as well as for their relation to polarization transforms.

Definition 7 (The conditional Bhattacharyya parameter). *Let $X \in \mathcal{X}$ be a binary random variable and $Y \in \mathcal{Y}$ be a discrete random variable. Let their joint distribution be $P_{X,Y}$. We denote*

$$\begin{aligned} Z(X|Y) &= 2 \sum_{y \in \mathcal{Y}} \sqrt{P_{X,Y}(0,y) \cdot P_{X,Y}(1,y)} \\ &= 2 \sum_{y \in \mathcal{Y}} P_Y(y) \sqrt{P_{X|Y}(0|y) \cdot P_{X|Y}(1|y)}. \end{aligned}$$

Definition 8 (The conditional total variation). *Let $X \in \mathcal{X}$ be a binary random variable and $Y \in \mathcal{Y}$ be a discrete random variable. Let their joint distribution be $P_{X,Y}$. We denote*

$$\begin{aligned} K(X|Y) &= \sum_{y \in \mathcal{Y}} |P_{X,Y}(0,y) - P_{X,Y}(1,y)| \\ &= \sum_{y \in \mathcal{Y}} P_Y(y) \cdot |P_{X|Y}(0|y) - P_{X|Y}(1|y)|. \end{aligned}$$

The following lemma shows that if $K(X|Y)$ is ‘small’, then $P(X|Y)$ is ‘close’ to the Bernoulli(1/2) distribution.

Lemma 25. *Let $X \in \mathcal{X}$ be a binary random variable and $Y \in \mathcal{Y}$ be a discrete random variable. Let their joint distribution be $P_{X,Y}$. Then*

$$\sum_{\substack{x \in \mathcal{X} \\ y \in \mathcal{Y}}} P_Y(y) \cdot |P_{X|Y}(x|y) - 1/2| = K(X|Y).$$

Proof.

$$\begin{aligned} &\sum_{\substack{x \in \mathcal{X} \\ y \in \mathcal{Y}}} P_Y(y) \cdot |P_{X|Y}(x|y) - 1/2| \\ &= \sum_{y \in \mathcal{Y}} P_Y(y) \cdot (|P_{X|Y}(0|y) - 1/2| + |P_{X|Y}(1|y) - 1/2|) \\ &= \sum_{y \in \mathcal{Y}} P_Y(y) \cdot (|P_{X|Y}(0|y) - P_{X|Y}(1|y)|) \\ &= K(X|Y), \end{aligned}$$

where the penultimate equality is easily seen to hold if we denote $P_{X|Y}(0|y) = 1/2 + \delta(y)$, from which it follows that $P_{X|Y}(1|y) = 1/2 - \delta(y)$. \square

B. Capacity-Achieving Inputs for the Deletion Channel

In [2], Dobrushin proves a capacity result for a class of synchronization error channels that includes the binary deletion channel. That paper also shows that the capacity can be approached by a sequence of finite-order Markov input distributions. Unfortunately, the Markov input distribution in Dobrushin’s construction is not irreducible [2, Lemma 4]. Thus, Dobrushin’s result falls slightly short of what is required by

the polar coding construction in this paper. In [22], Li and Tan study the capacity of the concatenation of a deletion channel and a finite-state channel. For this setup, they prove a capacity result and show that the capacity can be approached by a sequence of finite-order Markov input distributions that are irreducible and aperiodic. As they note in their paper, their result is sufficient to prove that the polar coding scheme in this paper can achieve capacity.

In this section, we describe a regular hidden-Markov input distribution that also achieves capacity on the deletion channel. Though this is not required, given [22], we include it for completeness and because the argument is somewhat different.

Denote by P_{X^N} an input distribution over binary vectors of length N , which we will shortly optimize over. Let $\underline{X} \triangleq (X_1, \dots, X_N)$ be a random binary vector of length N drawn according to P_{X^N} . Take \underline{X} as the input sequence to a binary deletion channel with deletion probability $\delta \in (0, 1)$ and let $\underline{Y} \triangleq (Y_1, \dots, Y_M)$ be the corresponding output sequence where the random variable M is the output length. The maximum mutual information for a length- N input is denoted by

$$C_N \triangleq \max_{P_{X^N}} \frac{1}{N} I(\underline{X}; \underline{Y}). \quad (113)$$

It is well-known [38, proof of Theorem II.1] that NC_N is a subadditive sequence and this implies [25, Lemma 1.2.1, page 3] that

$$C = \lim_{N \rightarrow \infty} C_N = \inf_{N \geq 1} C_N$$

exists and satisfies $C \leq C_N$ for $N \geq 1$. Thus, for the optimal P_{X^N} we have

$$\frac{1}{N} I(\underline{X}; \underline{Y}) \geq C. \quad (114)$$

We begin with the standard approach [39] of using an optimal P_{X^N} from (113) to generate a length- kN random input $\mathbf{X} = \mathbf{X}(1) \odot \dots \odot \mathbf{X}(k)$ where each $\mathbf{X}(i)$ is a length- N block drawn independently from P_{X^N} and using \odot to represent vector concatenation. For this input, we denote the output by $\mathbf{Y} = \mathbf{Y}(1) \odot \dots \odot \mathbf{Y}(k)$ where $\mathbf{Y}(i)$ contains the output symbols associated with the input $\mathbf{X}(i)$. Thus, for each i , the pair $\mathbf{X}(i), \mathbf{Y}(i)$ has the same distribution as the pair $\underline{X}, \underline{Y}$. The random variables $M_i = |\mathbf{Y}(i)|$, for $i \in [k]$, are chosen to equal the number of output symbols generated by the input block $\mathbf{X}(i)$.

Using the chain rule for mutual information, we note that

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}, M_1^k) &= I(\mathbf{X}; \mathbf{Y}) + I(\mathbf{X}; M_1^k | \mathbf{Y}) \\ &\leq I(\mathbf{X}; \mathbf{Y}) + k \log_2(N+1), \end{aligned}$$

where inequality follows from $I(\mathbf{X}; M_1^k | \mathbf{Y}) \leq \sum_{i=1}^k H(M_i)$ and $0 \leq M_i \leq N$. Thus, it follows that

$$\begin{aligned} I(\mathbf{X}; \mathbf{Y}) &\geq -k \log_2(N+1) + I(\mathbf{X}; \mathbf{Y}, M_1^k) \\ &\stackrel{(a)}{=} -k \log_2(N+1) + I(\mathbf{X}; \mathbf{Y}(1), \dots, \mathbf{Y}(k)) \\ &= -k \log_2(N+1) + \sum_{i=1}^k I(\mathbf{X}; \mathbf{Y}(i) | \mathbf{Y}(1), \dots, \mathbf{Y}(i-1)) \\ &\stackrel{(b)}{=} -k \log_2(N+1) + \sum_{i=1}^k I(\mathbf{X}(i); \mathbf{Y}(i)) \end{aligned}$$

$$\begin{aligned} &= -k \log_2(N+1) + k I(\underline{X}; \underline{Y}) \\ &= kN \left(\frac{1}{N} I(\underline{X}; \underline{Y}) - \frac{\log_2(N+1)}{N} \right) \\ &\stackrel{(c)}{\geq} kN \left(C - \frac{\log_2(N+1)}{N} \right), \end{aligned}$$

where (a) holds because there is an invertible mapping from \mathbf{Y}, M_1^k to $\mathbf{Y}(1), \dots, \mathbf{Y}(k)$, (b) follows from the pairs $(\mathbf{X}(i), \mathbf{Y}(i))_{i=1}^k$ being i.i.d., and (c) follows from (114). After normalizing by the input length, this gives

$$\frac{1}{kN} I(\mathbf{X}; \mathbf{Y}) \geq C - \frac{\log_2(N+1)}{N}.$$

Thus, the information rate can be made arbitrarily close to C by choosing N large enough.

However, the infinite input distribution formed by concatenating length- N blocks cannot be generated by a regular hidden-Markov process. In order to explain how to overcome this, we will first describe this input distribution as a hidden-Markov process with state set

$$\mathcal{S} \triangleq \bigcup_{j=0}^{N-1} \{x \in \{0, 1\}^j \mid P_{X^j}(x) \neq 0\},$$

where the set $\{0, 1\}^i$ represents all possible states after i input symbols from the length- N input distribution P_{X^N} . We denote the initial state by the empty string $\varepsilon \triangleq \{0, 1\}^0$ and let $P_{X^0}(\varepsilon) = 1$ by convention. To generate multiple blocks, we define the underlying Markov chain to start in the ε state and return to the ε state with probability 1 after generating N outputs. Thus, the underlying Markov chain is irreducible because we have only included states with positive probability and there is a path with positive probability from ε to any $x \in \mathcal{S}$.

Notice that the state implicitly encodes the current input position in the length- N block distribution. For example, if $s \in \{0, 1\}^j$, then next symbol is drawn according to $P_{X^{j+1}|X^j}(x|s)$. Thus, the underlying Markov chain is periodic with period N . To make it aperiodic, we will introduce one additional state, which we denote by τ , that is used to dither the input block between length- N and length- $(N+1)$. State τ always outputs a dither bit whose value is 0 and then transitions to state ε . The idea is that, after a length- N input block, a fair coin is used to determine if the next block will start immediately (e.g., the underlying Markov chain transitions to state ε) or be delayed by one symbol (e.g., the underlying Markov chain transitions to state τ). After this, the modified Markov chain will be aperiodic because the transition graph has loops of length N and $N+1$. The period of a Markov chain is the greatest common divisor of the lengths of all loops in the transition graph. Since N and $N+1$ are relatively prime, the period is 1 and the chain is aperiodic. We also note that the new Markov chain is still irreducible because there is still a path with positive probability between any two states.

Let S_0 be initial state of the underlying Markov chain. In the current formulation, we have $S_0 = \varepsilon$ with probability 1 and the Markov chain is not stationary. One can make this Markov chain stationary by drawing the initial state S_0 from the stationary distribution of the underlying Markov chain. After

this change, we have constructed a regular hidden-Markov input derived from our original P_{X^N} block distribution.

Now, let \mathbf{X} be a length- $k(N+1)$ input drawn from the constructed hidden-Markov process. This input can be broken into segments by adding commas before the inputs generated by the state ε . A complete segment is delimited by commas on both sides, and thus has length either N or $N+1$. Note that \mathbf{X} contains at least k segments, and by discarding the first segment we get at least $k-1$ complete segments. We call the length- N prefix of a complete segment a block. Thus, we have at least $k-1$ blocks, $\mathbf{X}(2), \dots, \mathbf{X}(k)$, where each block can be associated with an independent draw from P_{X^N} . Let $T_i \in \{0, 1\}$ be the side-information random variable that indicates, for the i -th (possibly incomplete) segment, whether or not state τ was visited during that segment. Given S_0 and T_1^k , it is always possible to compute the locations of the commas described above and separate \mathbf{X} into the $k-1$ blocks $\mathbf{X}(2), \dots, \mathbf{X}(k)$. This is because S_0 gives the initial offset to the first segment and T_i indicates whether or not each segment has the additional dither bit.

Similarly, the output \mathbf{Y} can be separated into subvectors associated with the above blocks by adding commas to separate outputs generated by different segments and removing any outputs caused by dither bits. Namely, we let $M_i \in \{0, \dots, N+1\}$ be the side-information random variable that indicates the number of outputs generated by the i -th segment and $R_i \in \{0, 1\}$ be the side-information random variable that indicates whether the last output in a subvector is due to a dither bit. Given M_1^k and R_1^k , it is always possible to separate \mathbf{Y} into $\mathbf{Y}(2), \dots, \mathbf{Y}(k)$ where each $\mathbf{Y}(i)$ is the output associated with the block $\mathbf{X}(i)$. Thus, each pair $(\mathbf{X}(i), \mathbf{Y}(i))$ has the same distribution as $(\underline{X}, \underline{Y})$. Using this setup, the chain rule of mutual information and cardinality upper bounds imply that

$$\begin{aligned}
I(\mathbf{X}, T_1^k; \mathbf{Y}, M_1^k, R_1^k | S_0) &= I(\mathbf{X}, T_1^k; \mathbf{Y}, M_1^k, R_1^k) \\
&\quad + I(\mathbf{X}, T_1^k; S_0 | \mathbf{Y}, M_1^k, R_1^k) - I(\mathbf{X}, T_1^k; S_0) \\
&\leq I(\mathbf{X}, T_1^k; \mathbf{Y}, M_1^k, R_1^k) + I(\mathbf{X}, T_1^k; S_0 | \mathbf{Y}, M_1^k, R_1^k) \\
&\stackrel{(a)}{\leq} I(\mathbf{X}, T_1^k; \mathbf{Y}, M_1^k, R_1^k) + N \\
&= I(\mathbf{X}; \mathbf{Y}, M_1^k, R_1^k) + I(T_1^k; \mathbf{Y}, M_1^k, R_1^k | \mathbf{X}) + N \\
&\stackrel{(b)}{\leq} I(\mathbf{X}; \mathbf{Y}, M_1^k, R_1^k) + k + N \\
&= I(\mathbf{X}; \mathbf{Y}) + I(\mathbf{X}; M_1^k, R_1^k | \mathbf{Y}) + k + N \\
&\stackrel{(c)}{\leq} I(\mathbf{X}; \mathbf{Y}) + k \log_2(N+2) + k + k + N, \quad (115)
\end{aligned}$$

where (a) follows from $\log_2 |\mathcal{S}| = \log_2 \left(1 + \sum_{j=0}^{N-1} 2^j\right) = N$, (b) holds because $T_i \in \{0, 1\}$, and (c) follows from $0 \leq M_i \leq N+1$ and $R_i \in \{0, 1\}$.

Based on the decompositions described above, the data processing inequality implies that

$$\begin{aligned}
I(\mathbf{X}, T_1^k; \mathbf{Y}, M_1^k, R_1^k | S_0) &\geq I(\mathbf{X}(2), \dots, \mathbf{X}(k); \mathbf{Y}(2), \dots, \mathbf{Y}(k) | S_0) \\
&= I(\mathbf{X}(2), \dots, \mathbf{X}(k); \mathbf{Y}(2), \dots, \mathbf{Y}(k)) \\
&= \sum_{i=2}^k I(\mathbf{X}(i); \mathbf{Y}(i)) = (k-1)I(\underline{X}; \underline{Y}). \quad (116)
\end{aligned}$$

Combining (114)–(116), we have

$$I(\mathbf{X}; \mathbf{Y}) \geq -k \log_2(N+2) - 2k - N + (k-1)CN.$$

To lower bound the information rate, we can normalize by the input length to see that

$$\begin{aligned}
\frac{1}{k(N+1)} I(\mathbf{X}; \mathbf{Y}) &\geq \frac{(k-1)N}{k(N+1)} \left(C - \frac{1}{k-1} \right) - \frac{2 + \log_2(N+2)}{N+1}.
\end{aligned}$$

By choosing k and N large enough, the information rate can be made arbitrarily close to C . Thus, we have constructed a sequence of regular hidden-Markov input distributions that achieve capacity on the binary deletion channel.

In closing, we note that this argument works without change for channels with independent insertions, deletions, and substitutions.

REFERENCES

- [1] R. G. Gallager, "Sequential decoding for binary channels with noise and synchronization errors," 1961, Lincoln Lab Group Rep. 25G-2, Mass. Inst. of Tech., Lexington MA.
- [2] R. L. Dobrushin, "Shannon's theorems for channels with synchronization errors," *Problemy Peredachi Informatsii*, vol. 3, no. 4, pp. 18–36, 1967.
- [3] M. C. Davey and D. J. MacKay, "Reliable communication over channels with insertions, deletions, and substitutions," *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 687–698, 2001.
- [4] M. Mitzenmacher, "A survey of results for deletion channels and related synchronization channels," *Probability Surveys*, vol. 6, pp. 1–33, 2009.
- [5] D. Fertonani and T. M. Duman, "Novel bounds on the capacity of the binary deletion channel," *IEEE Transactions on Information Theory*, vol. 56, no. 6, pp. 2753–2765, 2010.
- [6] H. Mercier, V. Tarokh, and F. Labeau, "Bounds on the capacity of discrete memoryless channels corrupted by synchronization and substitution errors," *IEEE Transactions on Information Theory*, vol. 58, no. 7, pp. 4306–4330, 2012.
- [7] A. R. Iyengar, P. H. Siegel, and J. K. Wolf, "Modeling and information rates for synchronization error channels," in *Proc. IEEE Int. Sym. on Information Theory*. IEEE, 2011, pp. 380–384.
- [8] —, "On the capacity of channels with timing synchronization errors," *IEEE Trans. Inform. Theory*, vol. 62, no. 2, pp. 793–810, 2015.
- [9] M. Rahmati and T. M. Duman, "Upper bounds on the capacity of deletion channels using channel fragmentation," *IEEE Transactions on Information Theory*, vol. 61, no. 1, pp. 146–156, 2015.
- [10] J. Castiglione and A. Kavcic, "Trellis based lower bounds on capacities of channels with synchronization errors," in *Information Theory Workshop*. Jeju, South Korea: IEEE, 2015, pp. 24–28.
- [11] M. Cheraghchi, "Capacity upper bounds for deletion-type channels," *Journal of the ACM (JACM)*, vol. 66, no. 2, p. 9, 2019.
- [12] E. K. Thomas, V. Y. F. Tan, A. Vardy, and M. Motani, "Polar coding for the binary erasure channel with deletions," *IEEE Communications Letters*, vol. 21, no. 4, pp. 710–713, April 2017.
- [13] K. Tian, A. Fazeli, A. Vardy, and R. Liu, "Polar codes for channels with deletions," in *55th Annual Allerton Conference on Communication, Control, and Computing*, 2017, pp. 572–579.
- [14] K. Tian, A. Fazeli, and A. Vardy, "Polar coding for deletion channels: Theory and implementation," in *IEEE International Symposium on Information Theory*, 2018, pp. 1869–1873.
- [15] —, "Polar coding for channels with deletions," *IEEE Transactions on Information Theory*, vol. 67, no. 11, pp. 7081–7095, November 2021.
- [16] I. Tal and A. Vardy, "List decoding of polar codes," *IEEE Transactions on Information Theory*, vol. 61, no. 5, pp. 2213–2226, May 2015.
- [17] R. Wang, R. Liu, and Y. Hou, "Joint successive cancellation decoding of polar codes over intersymbol interference channels," 2014, arXiv preprint arXiv:1404.3001.
- [18] R. Wang, J. Honda, H. Yamamoto, R. Liu, and Y. Hou, "Construction of polar codes for channels with memory," in *2015 IEEE Information Theory Workshop*, October 2015, pp. 187–191.
- [19] R. A. Wagner and M. J. Fischer, "The string-to-string correction problem," *Journal of the ACM (JACM)*, vol. 21, no. 1, pp. 168–173, 1974.

- [20] E. Şaşoğlu and I. Tal, “Polar coding for processes with memory,” *IEEE Trans. Inform. Theory*, vol. 65, no. 4, pp. 1994–2003, April 2019.
- [21] B. Shuval and I. Tal, “Fast polarization for processes with memory,” *IEEE Trans. Inform. Theory*, vol. 65, no. 4, pp. 2004–2020, April 2019.
- [22] Y. Li and V. Y. F. Tan, “On the capacity of channels with deletions and states,” *arXiv preprint arXiv:1911.04473*, 2019.
- [23] E. Şaşoğlu, “Polar Coding Theorems for Discrete Systems,” Ph.D. dissertation, IC, Lausanne, 2011.
- [24] E. Arıkan, “Channel polarization: a method for constructing capacity-achieving codes for symmetric binary-input memoryless channels,” *IEEE Trans. on Information Theory*, vol. 55, no. 7, pp. 3051–3073, July 2009.
- [25] J. M. Steele, *Probability Theory and Combinatorial Optimization*. Philadelphia, PA: SIAM, 1997, vol. 69, CBMF-NSF Regional Conference Series in Applied Mathematics.
- [26] R. Durrett, *Probability: Theory and Examples*. Cambridge University Press, 2019, vol. 49.
- [27] E. Şaşoğlu, “Polarization and polar codes,” in *Found. and Trends in Commun. and Inform. Theory*, vol. 8, no. 4, 2012, pp. 259–381.
- [28] H. D. Pfister and I. Tal, “Polar codes for channels with insertions, deletions, and substitutions,” in *Proc. IEEE Int’l Symp. Inform. Theory (ISIT’2021)*, Melbourne, Victoria, Australia, 2021, pp. 2554–2559.
- [29] J. Honda and H. Yamamoto, “Polar coding without alphabet extension for asymmetric models,” *IEEE Transactions on Information Theory*, vol. 59, no. 12, pp. 7829–7838, December 2013.
- [30] E. Arıkan and E. Telatar, “On the rate of channel polarization,” in *Proc. IEEE Int. Sym. on Information Theory*, June 2009, pp. 1493–1495.
- [31] B. Shuval and I. Tal, “Universal polarization for processes with memory,” 2018, arXiv:1811.05727v1.
- [32] W. Hoeffding, “Probability inequalities for sums of random variables,” *Journal of the American Statistical Association*, vol. 53, no. 301, pp. 13–30, March 1963.
- [33] M. Mitzenmacher and E. Upfal, *Probability and Computing: Randomization and Probabilistic Techniques in Algorithms and Data Analysis*, 2nd ed. Cambridge, UK: Cambridge University Press, 2005.
- [34] M. Mondelli, S. H. Hassani, and R. Urbanke, “How to achieve the capacity of asymmetric channels,” *IEEE Trans. Inform. Theory*, vol. 64, no. 5, pp. 3371–3393, May 2018.
- [35] A. Chou and M. R. Bloch, “Using deterministic decisions for low-entropy bits in the encoding and decoding of polar codes,” in *Proc. Annu. Allerton Conf. Commun. Control Computing (Allerton’2015)*, New York, NY, USA, 2015.
- [36] S. B. Korada, “Polar codes for channel and source coding,” Ph.D. dissertation, Ecole Polytechnique Fédérale de Lausanne, 2009.
- [37] I. Tal, “polarcub,” <https://github.com/idotal/polarcub>, 2021.
- [38] Y. Kanoria and A. Montanari, “Optimal coding for the binary deletion channel with small deletion probability,” *IEEE Trans. Inform. Theory*, vol. 59, no. 10, pp. 6192–6219, 2013.
- [39] J. Chen and P. H. Siegel, “Markov processes asymptotically achieve the capacity of finite-state intersymbol interference channels,” *IEEE Transactions on Information Theory*, vol. 54, no. 3, pp. 1295–1303, 2008.

Ido Tal (S’05–M’08–SM’18) was born in Haifa, Israel, in 1975. He received the B.Sc., M.Sc., and Ph.D. degrees in computer science from Technion — Israel Institute of Technology, Haifa, Israel, in 1998, 2003 and 2009, respectively. During 2010–2012 he was a postdoctoral scholar at the University of California at San Diego. In 2012 he joined the Electrical and Computer Engineering Department at Technion. His research interests include constrained coding and error-control coding. He received the IEEE Joint Communications Society/Information Theory Society Paper Award (jointly with Alexander Vardy) for the year 2017.

Henry D. Pfister received his Ph.D. in Electrical Engineering in 2003 from the University of California, San Diego and is currently a professor in the Electrical and Computer Engineering Department of Duke University with a secondary appointment in Mathematics. Prior to that, he was an associate professor at Texas A&M University (2006–2014), a post-doctoral fellow at the École Polytechnique Fédérale de Lausanne (2005–2006), and a senior engineer at Qualcomm Corporate R&D in San Diego (2003–2004). His current research interests include information theory, error-correcting codes, quantum computing, and machine learning.

He received the NSF Career Award in 2008 and a Texas A&M ECE Department Outstanding Professor Award in 2010. He is a coauthor of the 2007 IEEE COMSOC best paper in Signal Processing and Coding for Data Storage, a coauthor of a 2016 Symposium on the Theory of Computing (STOC) best paper, and a recipient of the 2021 Information Theory Society Paper Award. He has served the IEEE Information Theory Society as a member of the Board of Governors (2019–2021), an Associate Editor for the IEEE Transactions on Information Theory (2013–2016), and a Distinguished Lecturer (2015–2016). He was the General Chair of the 2016 North American School of Information Theory and a Technical Program Committee Co-Chair of the 2021 International Symposium on Information Theory.

Arman Fazeli (M’14) received the B.S. degree in Electrical Engineering from Sharif University of Technology, Tehran, Iran, in 2012. He has received his Ph.D. degree in Electrical and Computer Engineering at the University of California, San Diego. He was a postdoctoral scholar and a lecturer at the same university until 2021. He is currently a Wireless System Engineer at Apple. His main research interests include coding and information theory, with emphasis on polar codes and coding for distributed storage systems. Arman received silver and bronze medals at the International Mathematical Olympiad (IMO) in 2006 and 2007, when he was in the Iran national mathematics team.

Alexander Vardy (Fellow, IEEE) was born in Moscow, Russia, in 1963. He received the B.Sc. degree (*summa cum laude*) from the Technion–Israel Institute of Technology, Israel, in 1985, and the Ph.D. degree from Tel-Aviv University, Israel, in 1991. From 1985 to 1990, he was with the Israeli Air Force, where he worked on electronic counter measures systems and algorithms. From 1992 to 1993, he was a Visiting Scientist with the IBM Almaden Research Center, San Jose, CA, USA. From 1993 to 1998, he was with the University of Illinois at Urbana–Champaign, first as an Assistant Professor then as an Associate Professor. Since 1998, he has been with the University of California at San Diego (UCSD), where he is currently the Jack Keil Wolf Chair Professor with the Department of Electrical and Computer Engineering and the Department of Computer Science. While on sabbatical from UCSD, he has held long-term visiting appointments with CNRS, France, the EPFL, Switzerland, the Technion–Israel Institute of Technology, and Nanyang Technological University, Singapore. His research interests include error-correcting codes, algebraic and iterative decoding algorithms, lattices and sphere packings, coding for storage systems, cryptography, computational complexity theory, and fun math problems.

He has been a member of the Board of Governors of the IEEE Information Theory Society from 1998 to 2006 and from 2011 to 2017. He received an IBM Invention Achievement Award in 1993 and NSF Research Initiation and CAREER Awards in 1994 and 1995. In 1996, he was appointed as a Fellow in the Center for Advanced Study, University of Illinois, and received the Xerox Award for Faculty Research. In 1996, he became a Fellow of the David and Lucile Packard Foundation. He received the IEEE Information Theory Society Paper Award (jointly with Ralf Koetter) in 2004. In 2005, he received the Fulbright Senior Scholar Fellowship, and the Best Paper Award at the IEEE Symposium on Foundations of Computer Science (FOCS). In 2017, his work on polar codes was recognized by the IEEE Communications and Information Theory Societies Joint Paper Award. From 1995 to 1998, he was an Associate Editor for *Coding Theory* and from 1998 to 2001, he was the Editor-in-Chief of the IEEE TRANSACTIONS ON INFORMATION THEORY.