

General Strong Polarization

Jarosław Błasiok* Venkatesan Guruswami† Preetum Nakkiran‡
 Atri Rudra§ Madhu Sudan¶

Abstract

Arıkan’s exciting discovery of polar codes has provided an altogether new way to efficiently achieve Shannon capacity. Given a (constant-sized) invertible matrix M , a family of polar codes can be associated with this matrix and its ability to approach capacity follows from the *polarization* of an associated $[0, 1]$ -bounded martingale, namely its convergence in the limit to either 0 or 1 with probability 1. Arıkan showed appropriate polarization of the martingale associated with the matrix $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ to get capacity achieving codes. His analysis was later extended to all matrices M which satisfy an obvious necessary condition for polarization.

While Arıkan’s theorem does not guarantee that the codes achieve capacity at small block-lengths (specifically in length which is a polynomial in $1/\varepsilon$ where ε is the difference between the capacity of a channel and the rate of the code), it turns out that a “strong” analysis of the polarization of the underlying martingale would lead to such constructions. Indeed for the martingale associated with G_2 such a strong polarization was shown in two independent works ([Guruswami and Xia, IEEE IT ’15] and [Hassani et al., IEEE IT ’14]), thereby resolving a major theoretical challenge associated with the efficient attainment of Shannon capacity.

In this work we extend the result above to cover martingales associated with all matrices that satisfy the necessary condition for (weak) polarization. In addition to being vastly more general, our proofs of strong polarization are (in our view) also much simpler and modular. Key to our proof is a notion of *local polarization* that only depends on the evolution of the martingale in a single time step. We show that local polarization always implies strong polarization. We then apply relatively simple reasoning about conditional entropies to prove local polarization in very general settings. Specifically, our result shows strong polarization over all prime fields and leads to efficient capacity-achieving source codes for compressing arbitrary i.i.d. sources, and capacity-achieving channel codes for arbitrary symmetric memoryless channels.

*John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: jblasio@g.harvard.edu. Supported by ONR grant N00014-15-1-2388.

†Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Some of this work was done when the author was visiting the School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore. venkatg@cs.cmu.edu. Research supported in part by NSF grants CCF-1422045 and CCF-1563742.

‡John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: preetum@cs.harvard.edu. Work supported in part by a Simons Investigator Award, NSF Awards CCF 1565641 and CCF 1715187, and the NSF Graduate Research Fellowship Grant No. DGE1144152.

§Computer Science and Engineering Department, University at Buffalo, SUNY. atri@buffalo.edu. Research supported in part by NSF grant CCF-1717134.

¶Harvard John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford Street, Cambridge, MA 02138, USA. Email: madhu@cs.harvard.edu. Work supported in part by a Simons Investigator Award and NSF Awards CCF 1565641 and CCF 1715187.

Contents

1	Introduction	2
1.1	Polarization of $[0, 1]$ -martingales	2
1.2	Results I: Local Polarization and Implication	3
1.3	The Arikan martingale and Polar codes	4
1.4	Results II: Local polarization of Arikan martingales	6
1.5	Comparison with previous analyses of (strong) polarization	6
2	Preliminaries and Notation	8
2.1	Notation	8
2.1.1	Probability Notation	8
2.1.2	Tensor Notation	8
2.1.3	Tensor Product Recursion	9
2.2	Information Theory Preliminaries	9
2.2.1	Channels	10
2.3	Basic Probabilistic Inequalities	10
3	Local to global polarization	14
4	Arikan Martingale	17
5	Proof of Local Polarization	19
5.1	Proof Overview	19
5.2	Entropic Lemmas in the 2×2 Case	20
5.3	Local polarization of $k \times k$ mixing matrices	21
5.3.1	Reduction to the 2×2 case	22
5.3.2	Proof of Theorem 1.10	24
6	Proofs of Entropic Lemmas	26
6.1	Suction at the upper end	26
6.2	Suction at the lower end	28
A	Codes from Polarization	35
A.1	Polar Encoder	35
A.2	The Successive-Cancellation Decoder	36
A.2.1	Decoding Analysis	36
A.2.2	Fast Decoder	38
A.2.3	Arikan Martingale and Polar Coding	42

1 Introduction

Polar codes, proposed in Arikan’s remarkable work [2], gave a fresh information-theoretic approach to construct linear codes that achieve the Shannon capacity of symmetric channels, together with efficient encoding and decoding algorithms. About a decade after their discovery, there is now a vast and extensive body of work on polar coding spanning hundreds of papers, and polar codes are also being considered as one of the candidates for use in 5G wireless (e.g., see [7] and references therein). The underlying concept of polarizing transforms has emerged as a versatile tool to successfully attack a diverse collection of information-theoretic problems beyond the original channel and source coding applications, including wiretap channels [16], the Slepian-Wolf, Wyner-Ziv, and Gelfand-Pinsker problems [14], broadcast channels [9], multiple access channels [22, 1], and interference networks [24]. We recommend the survey by Şaşıoğlu [21] for a nice treatment of the early work on polar codes.

The algorithmic interest in polar codes emerges from a consequence shown in the works [11, 12, 10] who show that this approach leads to a family of codes of rate $C - \varepsilon$ for transmission over a channel of (Shannon) capacity C , where the block length of the code and the decoding time grow only polynomially in $1/\varepsilon$. In contrast, for all previous constructions of codes, the decoding algorithms required time exponential in $1/\varepsilon$. Getting a polynomial running time in $1/\varepsilon$ was arguably one of the most important theoretical challenges in the field of algorithmic coding theory, and polar codes were the first to overcome this challenge. The analyses of polar codes turn into questions about *polarizations* of certain *martingales*. The vast class of polar codes alluded to in the previous paragraph all build on polarizing martingales, and the results of [11, 12, 10] show that for one of the families of polar codes, the underlying martingale polarizes “extremely fast” — a notion we refer to as *strong polarization* (which we will define shortly).

The primary goal of this work is to understand the process of polarization of martingales, and in particular to understand when a martingale polarizes strongly. In attempting to study this question, we come up with a local notion of polarization and show that this local notion is sufficient to imply strong polarization. Applying this improved understanding to the martingales arising in the study of polar codes we show that a simple necessary condition for weak polarization of such martingales is actually sufficient for strong polarization. This allows us to extend the results of [11, 12, 10] to a broad class of codes and show essentially that all polarizing codes lead to polynomial convergence to capacity. Below we formally describe the notion of polarization of martingales and our results.

1.1 Polarization of $[0, 1]$ -martingales

Our interest is mainly in the (rate of) polarization of a specific family of martingales that we call the Arikan martingales. We will define these objects later, but first describe the notion of polarization for general $[0, 1]$ -bounded martingales. Recall that a sequence of random variables X_0, \dots, X_t, \dots is said to be a *martingale* if for every t and a_0, \dots, a_t it is the case that $\mathbb{E}[X_{t+1} | X_0 = a_0, \dots, X_t = a_t] = a_t$. We say that a martingale is *$[0, 1]$ -bounded* (or simply a $[0, 1]$ -martingale) if $X_t \in [0, 1]$ for all $t \geq 0$.

Definition 1.1 (Weak Polarization). *A $[0, 1]$ -martingale sequence $X_0, X_1, \dots, X_t, \dots$ is defined to be weakly polarizing if $\lim_{t \rightarrow \infty} \{X_t\}$ exists with probability 1, and this limit is either 0 or 1 (and so the limit is a Bernoulli random variable with expectation X_0).*

Thus a polarizing martingale does not converge to a single value with probability 1, but rather converges to one of its extreme values. For the applications to constructions of polar codes, we need

more explicit bounds on the rates of convergence leading to the notions of (regular) polarization and strong polarization defined below in Definition 1.3 and 1.4 respectively.

Definition 1.2 ((τ, ε) -Polarization). *For functions $\tau, \varepsilon : \mathbb{Z}^+ \rightarrow \mathbb{R}^{\geq 0}$, a $[0, 1]$ -martingale sequence $X_0, X_1, \dots, X_t, \dots$ is defined to be (τ, ε) -polarizing if for all t we have*

$$\Pr(X_t \in (\tau(t), 1 - \tau(t))) < \varepsilon(t).$$

Definition 1.3 (Regular Polarization). *A $[0, 1]$ -martingale sequence $X_0, X_1, \dots, X_t, \dots$ is defined to be regular polarizing if for all constant $\gamma > 0$, there exist $\varepsilon(t) = o(1)$, such that X_t is $(\gamma^t, \varepsilon(t))$ -polarizing.*

We refer to the above as being “sub-exponentially” close to the limit (since it holds for every $\gamma > 0$). While weak polarization by itself is an interesting phenomenon, regular polarization (of Arikan martingales) leads to capacity-achieving codes (though without explicit bounds on the length of the code as a function of the gap to capacity) and thus regular polarization is well-explored in the literature and tight necessary and sufficient conditions are known for regular polarization of Arikan martingales [3, 15].

To get codes of block length polynomially small in the gap to capacity, an even stronger notion of polarization is needed, where we require that the sub-exponential closeness to the limit happens with *all but exponentially small probability*. We define this formally next.

Definition 1.4 (Strong Polarization). *A $[0, 1]$ -martingale sequence $X_0, X_1, \dots, X_t, \dots$ is defined to be strongly polarizing if for all $\gamma > 0$ there exist $\eta < 1$ and $\beta < \infty$ such that martingale X_t is $(\gamma^t, \beta \cdot \eta^t)$ -polarizing.*

In contrast to the rich literature on regular polarization, results on strong polarization are quite rare, reflecting a general lack of understanding of this phenomenon. Indeed (roughly) an Arikan martingale can be associated with every invertible matrix over any finite field \mathbb{F}_q , and the only matrix for which strong polarization is known is $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ [11, 12, 10].¹

Part of the reason behind the lack of understanding of strong polarization is that polarization is a “limiting phenomenon” in that one tries to understand $\lim_{t \rightarrow \infty} X_t$, whereas most stochastic processes, and the Arikan martingales in particular, are defined by local evolution, i.e., one that relates X_{t+1} to X_t . The main contribution of this work is to give a local definition of polarization (Definition 1.5) and then showing that this definition implies strong polarization (Theorem 1.6). Later we show that Arikan martingales polarize locally whenever they satisfy a simple condition that is necessary even for weak polarization. As a consequence we get strong polarization for all Arikan martingales for which previously only regular polarization was known.

1.2 Results I: Local Polarization and Implication

Before giving the definition of local polarization, we give some intuition using the following martingale: Let $Z_0 = 1/2$, and $Z_{t+1} = Z_t + Y_{t+1}2^{-(t+2)}$ where Y_1, \dots, Y_t, \dots are chosen uniformly and

¹An exception is the work by Pfister and Urbanke [19] who showed that for the q -ary erasure channel for large enough q , the martingale associated with a $q \times q$ Reed-Solomon based matrix proposed in [18] polarizes strongly. A recent (unpublished) work [8] shows that for the binary erasure channel, martingales associated with large random matrices polarize strongly. Both these results obtain an optimal value of η for (specific/random) large matrices. However, they only apply to the erasure channel, which is simple to error correct via Gaussian elimination and therefore not really reflective of the general capacity-achieving power of polar codes.

independently from $\{-1, +1\}$. Clearly this sequence is not polarizing (the limit of Z_t is uniform in $[0, 1]$). One reason why this happens is that as time progresses, the martingale slows down and stops varying much. We would like to prevent this, but this is also inevitable if a martingale is polarizing. In particular, a polarizing martingale would be slowed at the boundary and cannot vary much. The first condition in our definition of local polarization insists that this be the only reason a martingale slows down (we refer to this as *variance in the middle*).

Next we consider what happens when a martingale is close to the boundary. For this part consider a martingale $Z_0 = 1/2$ and $Z_{t+1} = Z_t + \frac{1}{2}Y_{t+1} \min\{Z_t, 1 - Z_t\}$. This martingale does polarize and even shows regular polarization, but it can also be easily seen that the probability that $Z_t < \frac{1}{2} \cdot 2^{-t}$ is zero (whereas we would like probability of being less than say 10^{-t} to go to 1). So this martingale definitely does not show strong polarization. This is so since even in the best case the martingale is approaching the boundary at a fixed exponential rate, and not a sub-exponential one. To overcome this obstacle we require that when the martingale is close to the boundary, with a fixed constant probability it should get much closer in a single step (a notion we refer to as *suction at the ends*).

The definition below makes the above requirements precise.

Definition 1.5 (Local Polarization). *A $[0, 1]$ -martingale sequence X_0, \dots, X_j, \dots , is locally polarizing if the following conditions hold:*

1. **(Variance in the middle):** *For every $\tau > 0$, there is a $\theta = \theta(\tau) > 0$ such that for all j , we have: If $X_j \in (\tau, 1 - \tau)$ then $\mathbb{E}[(X_{j+1} - X_j)^2 | X_j] \geq \theta$.*
2. **(Suction at the ends):** *There exists an $\alpha > 0$, such that for all $c < \infty$, there exists a $\tau = \tau(c) > 0$, such that:*
 - (a) *If $X_j \leq \tau$ then $\Pr[X_{j+1} \leq X_j/c | X_j] \geq \alpha$.*
 - (b) *Similarly, if $1 - X_j \leq \tau$ then $\Pr[(1 - X_{j+1}) \leq (1 - X_j)/c | X_j] \geq \alpha$.*

We refer to condition (a) above as Suction at the low end and condition (b) as Suction at the high end.

When we wish to be more explicit, we refer to the sequence as $(\alpha, \tau(\cdot), \theta(\cdot))$ -locally polarizing.

As such this definition is neither obviously sufficient for strong polarization, nor is it obviously satisfiable by any interesting martingale. In the rest of the paper, we address these concerns. Our first technical contribution is a general theorem connecting local polarization to strong polarization.

Theorem 1.6 (Local vs. Strong Polarization). *If a $[0, 1]$ -martingale sequence X_0, \dots, X_t, \dots , is locally polarizing, then it is also strongly polarizing.*

It remains to show that the notion of local polarization is not vacuous. Next, we show that in fact Arıkan martingales polarize locally (under simple necessary conditions). First we give some background on Polar codes.

1.3 The Arıkan martingale and Polar codes

The setting of polar codes considers an arbitrary *symmetric memoryless channel* and yields codes that aim to achieve the *capacity* of this channel. These notions are reviewed in Section 2.2.1. Given any q -ary memoryless channel $\mathcal{C}_{Y|Z}$ and invertible matrix $M \in \mathbb{F}_q^{k \times k}$, the theory of polar codes implicitly defines a martingale, which we call the Arıkan martingale associated with $(M, \mathcal{C}_{Y|Z})$ and studies its polarization. (An additional contribution of this work is that we give an explicit compact definition of this martingale, see Definition 4.1. Since we do not need this definition for

the purposes of this section, we defer it for Section 4). The consequences of regular polarization are described by the following remarkable theorem. (Below we use $M \otimes N$ to denote the tensor product of the matrix M and N . Further, we use $M^{\otimes t}$ to denote the tensor of a matrix M with itself t times.)

Theorem 1.7 (Implied by Arıkan [2]). *Let \mathcal{C} be a q -ary symmetric memoryless channel and let $M \in \mathbb{F}_q^{k \times k}$ be an invertible matrix. If the Arıkan martingale associated with (M, \mathcal{C}) polarizes regularly, then given $\varepsilon > 0$ and $c < \infty$ there is a t_0 such that for every $t \geq t_0$ there is a code $C \subseteq \mathbb{F}_q^n$ for $n = k^t$ of dimension at least $(\text{Capacity}(\mathcal{C}) - \varepsilon) \cdot n$ such that C is an affine code generated by the restriction of $(M^{-1})^{\otimes t}$ to a subset of its rows and an affine shift. Moreover there is a polynomial time decoding algorithm for these codes that has failure probability bounded by n^{-c} .²*

For $n = 2^t$, Arıkan and Telatar [3] proved that the martingale associated with the matrix $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, polarizes regularly over any binary input symmetric channel (Arıkan’s original paper [2] proved a weaker form of regular polarization with $\tau(t) < 2^{-5t/4}$ which also sufficed for decoding error going to 0). Subsequent work generalized this to other matrices with the work of Korada, Şaşıođlu, and Urbanke [15] giving a precise characterization of matrices M for which the Arıkan martingale polarizes (again over binary input channels). We will refer to such matrices as *mixing*.

Definition 1.8 (Mixing Matrix). *A matrix $M \in \mathbb{F}_q^{k \times k}$ is said to be mixing, if it is invertible and none of the permutations of the rows of M yields an upper triangular matrix, i.e., for every permutation $\pi : [k] \rightarrow [k]$ there exists $i, j \in [k]$ with $j < \pi(i)$ such that $M_{i,j} \neq 0$.*

It is not too hard to show that the Arıkan martingale associated with non-mixing matrices do not polarize (even weakly). In contrast [15] shows that every mixing matrix over \mathbb{F}_2 polarizes regularly. Mori and Tanaka [18] show that the same result holds for all prime fields, and give a slightly more complicated criterion that characterizes (regular) polarization for general fields. (These works show that the decoding failure probability of the resulting polar codes is at most 2^{-n^β} for some positive β determined by the structure of the mixing matrix — this follows from an even stronger decay in the first of the two parameters in the definition of polarization. However, they do *not* show strong polarization, which is what we achieve.)

As alluded to earlier, strong polarization leads to even more effective code constructions and this is captured by the following theorem.

Theorem 1.9 ([2, 11, 12]). *Let \mathcal{C} be a q -ary symmetric memoryless channel and let $M \in \mathbb{F}_q^{k \times k}$ be an invertible matrix. If the Arıkan martingale associated with (M, \mathcal{C}) polarizes strongly, then for every c there exists $t_0(x) = O(\log x)$ such that for every $\varepsilon > 0$ and every $t \geq t_0(1/\varepsilon)$ there is an affine code C , that is generated by the rows of $(M^{-1})^{\otimes t}$ and an affine shift, with the property that the rate of C is at least $\text{Capacity}(\mathcal{C}) - \varepsilon$, and C can be encoded and decoded in time $O(n \log n)$ where $n = k^t$ and failure probability of the decoder is at most n^{-c} .*

This theorem is implicit in the works above, but for completeness we include a proof of this theorem in Appendix A. As alluded to earlier, the only Arıkan martingales that were known to polarize strongly were those where the underlying matrix was $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. Specifically Guruswami

²We remark that the encoding and decoding are not completely uniform as described above, since the subset of rows and the affine shift that are needed to specify the code are only guaranteed to exist. In the case of additive channels, where the shift can be assumed to be zero, the work of Tal and Vardy [23] (or [11, Sec. V]) removes this non-uniformity by giving a polynomial time algorithm to find the subset.

and Xia [11] and Hassani et al. [12] show strong polarization of the Arıkan martingale associated with this matrix over any binary input symmetric channel, and Guruswami and Velingker [10] extended to the case of q -ary input channels for prime q . By using the concept of local polarization we are able to extend these results to all mixing matrices.

1.4 Results II: Local polarization of Arıkan martingales

In our second main result, we show that every mixing matrix gives rise to an Arıkan martingale that is locally polarizing:

Theorem 1.10. *For every prime q , for every mixing matrix $M \in \mathbb{F}_q^{k \times k}$, and for every symmetric memoryless channel $\mathcal{C}_{Y|Z}$ over \mathbb{F}_q , the associated Arıkan martingale sequence is locally polarizing.*

As a consequence of Theorems 1.9, 1.6, and 1.10, we have the following theorem.

Theorem 1.11. *For every prime q , every mixing matrix $M \in \mathbb{F}_q^{k \times k}$, every symmetric memoryless channel \mathcal{C} over \mathbb{F}_q , and every $c < \infty$, there exists $t_0(x) = O(\log x)$ such that for every $\varepsilon > 0$, for every $t \geq t_0(1/\varepsilon)$, there is an affine code C , that is generated by the rows of $(M^{-1})^{(\otimes t)}$ and an affine shift, with the property that the rate of C is at least $\text{Capacity}(\mathcal{C}) - \varepsilon$, and C can be encoded and decoded in time $O(n \log n)$ where $n = k^t$ and failure probability of the decoder is at most n^{-c} .*

The above theorem shows that all polar codes associated with every mixing matrix achieves the Shannon capacity of a symmetric memoryless channel efficiently, thus, vastly expanding on the class of polar codes known to satisfy this condition.

Our primary motivation in this work is to develop a general approach to proving polarization that applies to all matrices (matching the simple necessary condition for polarization) and is strong enough for the desired coding theory conclusion (convergence to capacity at polynomial block lengths, the distinguishing feature of polar codes). At the same time, our proof is arguably simpler and brings to light exactly what drives strong polarization — namely some simple local polarization conditions that hold for the single step evolution. One concrete motivation to consider polar codes with different choice of mixing matrices M is that an appropriate choice can lead to decoding error probability of $\exp(-n^\beta)$ for any $\beta < 1$ (as opposed to $\beta < 1/2$ for G_2) [15, 18], where $n = k^t$ is the block length of the code.

1.5 Comparison with previous analyses of (strong) polarization

While most of the ingredients going into our eventual analysis of strong polarization are familiar in the literature on polar codes, our proofs end up being much simpler and modular. We describe some of the key steps in our proofs and contrast them with those in previous works.

Definition of Local Polarization. While we are not aware of a definition similar to local polarization being explicit in the literature before, such notions have been considered implicitly before. For instance, for the variation in the middle (where we require that $\mathbb{E}[(X_{t+1} - X_t)^2] \geq \theta$ if $X_t \in (\tau, 1 - \tau)$) the previous analyses in [11, 10] required θ be quadratic in τ . Indeed this was the most significant technical hurdle in the analysis for prime case in [10]. In contrast, our requirement on the variation is very weak and qualitative, allowing any function $\theta(\tau) > 0$. Similarly, our requirement in the *suction at the ends* case is relative mild and qualitative. In previous analyses the requirements were of the form “if $X_t \leq \tau$ then $X_{t+1} \leq X_t^2$ with positive probability.” This high demand on the suction case prevented the analyses from relying only on the local behavior of the

martingale X_0, \dots, X_t, \dots and instead had to look at other parameters associated with it which essentially depend on the entire sequence. (For the reader familiar with previous analyses, this is where the Bhattacharyya parameters enter the picture.) Our approach, in contrast, only requires arbitrarily large constant factor drop, and thereby works entirely with the local properties of X_t .

Local Polarization implies Strong Polarization. Our proof that local polarization implies strong polarization is short (about 3 pages) and comes in two parts. The first part uses a simple variance argument to show that X_t is exponentially close (in t) to the limit except with probability exponentially small in t . The second part then amplifies X_t 's proximity to $\{0, 1\}$ to sub-exponentially small values using the suction at the end guarantee of each local step, coupled with Doob's martingale inequality and standard concentration inequalities. Such a two-part breakdown of the analysis is not new; however, our technical implementation is more abstract, more general and more compact all at the same time.

Local Polarization of Arkan martingales. We will elaborate further on the approach for this after defining the Arkan martingales, but we can say a little bit already now: First we essentially reduce the analysis of the polarization of Arkan martingale associated with an arbitrary mixing matrix M to the analysis when $M = G_2$. This reduction loses in the parameters $(\alpha, \tau(\cdot), \theta(\cdot))$ specifying the level of local polarization, but since our strong polarization theorem works for any function, such loss in performance does not hurt the eventual result. Finally, local polarization for the case where the matrix is G_2 is of course standard, but even here our proofs (which we include for completeness) are simpler since they follow from known entropic inequalities on sums of two independent random variables. We stress that even quantitatively weak forms of these inequalities meet our requirements of local polarization, and we do not need strong forms of such inequalities (like Mrs. Gerber's lemma for the binary case [5, 11] and an ad hoc one for the prime case [10]).

Some weakness in our analyses. We first point out two weaknesses in our analyses. First, in contrast to the result of Mori and Tanaka [18] who characterize the set of matrices that lead to regular polarization over all fields, we only get a characterization over prime fields. Second, our definition of strong polarization only allows us to bound the failure probability of decoding by an arbitrarily small polynomial in the block length whereas results such as those in [3] actually get exponentially small (2^{-n^β} for some $\beta > 0$) failure probability.

In both cases we do not believe that these limitations are inherent to our approach. In particular the extension to general fields will probably involve more care, but should not run into major technical hurdles. Reducing the failure probability will lead to new technical challenges, but we do believe they can be overcome. Specifically, this requires stronger suction which is not true for the Arkan martingale if one considers a single step evolution, but it seems plausible that multiple steps (even two) might show strong enough suction. We hope to investigate this in future work.

Organization of the rest of this paper. We first introduce some of the notation and probabilistic preliminaries used to define and analyze the Arkan martingale in Section 2. We then prove Theorem 1.6 showing that local polarization implies strong polarization in Section 3. This is followed by the formal definition of the Arkan martingale in Section 4. Section 5.1 gives an overview of the proof of Theorem 1.10 which asserts that the Arkan martingale is locally polarizing (under appropriate conditions). Section 5.2 then states the local polarization conditions for sums of two independent variables, with proofs deferred to Section 6. Section 5.3 reduces the analysis of local polarization of general mixing matrices to the conditions studied in Section 5.2 and uses this reduction to prove Theorem 1.10. Finally in Appendix A we show (for completeness) how the Arkan martingale (and its convergence) can be used to construct capacity achieving codes.

2 Preliminaries and Notation

In this section we introduce the notation needed to define the Arikan martingale (which will be introduced in the following section). We also include information-theoretic and probabilistic inequalities that will be necessary for the subsequent analysis.

2.1 Notation

The Arikan martingale is based on a recursive construction of a vector valued random variable. To cleanly describe this construction it is useful to specify our notational conventions for vectors, tensors and how to view the tensor products of matrices. These notations will be used extensively in the following sections.

2.1.1 Probability Notation

Throughout this work, all random variables involved will be discrete. For a probability distribution D and random variable X , we write $X \sim D$ to mean that X is distributed according to D , and independent of all other variables. Similarly, for a set S , we write $X \sim S$ to mean that X is independent and uniform over S . For a set S , let $\Delta(S)$ denote the set of probability distributions over S .

We occasionally abuse notation by treating distributions as random variables. That is, for $D \in \Delta(\mathbb{F}_q^k)$ and a matrix $M \in \mathbb{F}_q^{k \times k}$, we write DM to denote the distribution of the random variable $\{XM\}_{X \sim D}$. For a distribution D and an event E , we write $D|E$ to denote the conditional distribution of D conditioned on E .

2.1.2 Tensor Notation

Here we introduce useful notation for dealing with scalars, vectors, tensors, and tensor-products.

All scalars will be non-boldfaced, for example: $X \in \mathbb{F}_q$.

Any tensors of order ≥ 1 (including vectors) will be boldfaced, for example: $\mathbf{Y} \in \mathbb{F}_q^k$. One exception to this is the matrix M used in the polarization transforms, which we do not boldface.

Subscripts are used to index tensors, with indices starting from 1. For example, for \mathbf{Y} as above, $\mathbf{Y}_i \in \mathbb{F}_q$. Matrices and higher-order tensors are indexed with multiple subscripts: For $\mathbf{Z} \in (\mathbb{F}_q^k)^{\otimes 3}$, we may write $\mathbf{Z}_{1,2,1} \in \mathbb{F}_q$. We often index tensors by tuples (*multiindices*), which will be boldfaced: For $\mathbf{i} = (1, 2, 1) \in [k]^3$, we write $\mathbf{Z}_{\mathbf{i}} = \mathbf{Z}_{1,2,1}$. Let \prec be the lexicographic order on these indexing tuples.

When an index into a tensor is the concatenation of multiple tuples, we emphasize this by using brackets in the subscript. For example: for tensor \mathbf{Z} as above, and $\mathbf{i} = (1, 2)$ and $j = 1$, we may write $\mathbf{Z}_{[\mathbf{i}, j]} = \mathbf{Z}_{1,2,1}$.

For a given tensor \mathbf{Z} , we can consider fixing some subset of its indices, yielding a *slice* of \mathbf{Z} (a tensor of lower order). We denote this with brackets, using \cdot to denote unspecified indices. For example for tensor $\mathbf{Z} \in (\mathbb{F}_q^k)^{\otimes 3}$ as above, we have $\mathbf{Z}_{[1,2,\cdot]} \in \mathbb{F}_q^k$ and $\mathbf{Z}_{[\cdot,1]} \in (\mathbb{F}_q^k)^{\otimes 2}$.

We somewhat abuse the indexing notation, using $\mathbf{Z}_{\prec \mathbf{i}}$ to mean the set of variables $\{\mathbf{Z}_{\mathbf{j}} : \mathbf{j} \prec \mathbf{i}\}$. Similarly, $\mathbf{Z}_{[\mathbf{i}, \prec j]} := \{\mathbf{Z}_{[\mathbf{i}, k]} : k \prec j\}$.

We occasionally unwrap tensors into vectors, using the correspondence between $(\mathbb{F}_q^k)^{\otimes t}$ and $\mathbb{F}_q^{k^t}$. Here, we unwrap according to the lexicographic order \prec on tuples.

Finally, for matrices specifically, $M_{i,j}$ specifies the entry in the i -th row and j -th column of matrix M . Throughout, all vectors will be row-vectors by default.

2.1.3 Tensor Product Recursion

The construction of polar codes and analysis of the Arıkan martingale rely crucially on the recursive structure of the tensor product. Here we review the definition of the tensor product, and state its recursive structure.

For a linear transform $M : \mathbb{F}_q^k \rightarrow \mathbb{F}_q^k$, let $M^{\otimes t} : (\mathbb{F}_q^k)^{\otimes t} \rightarrow (\mathbb{F}_q^k)^{\otimes t}$ denote the t -fold tensor power of M . Explicitly (fixing basis for all the spaces involved), this operator acts on tensors $\mathbf{X} \in (\mathbb{F}_q^k)^{\otimes t}$ as:

$$[M^{\otimes t}(\mathbf{X})]_j = \sum_{\mathbf{i} \in [k]^t} X_{\mathbf{i}} M_{i_1, j_1} M_{i_2, j_2} \cdots M_{i_t, j_t}.$$

The tensor product has the following recursive structure: $M^{\otimes t} = (M^{\otimes t-1}) \otimes M$, which corresponds explicitly to:

$$[M^{\otimes t}(\mathbf{X})]_{[a, j_t]} = \sum_{i_t \in [k]} M_{i_t, j_t} [M^{\otimes t-1}(\mathbf{X}_{[\cdot, i_t]})]_a. \quad (1)$$

In the above, if we define tensor

$$\mathbf{Y}^{(i_t)} := M^{\otimes t-1}(\mathbf{X}_{[\cdot, i_t]})$$

then this becomes

$$[M^{\otimes t}(\mathbf{X})]_{[a, \cdot]} = M((\mathbf{Y}_a^{(1)}, \mathbf{Y}_a^{(2)}, \dots, \mathbf{Y}_a^{(k)})) \quad (2)$$

where the vector $(\mathbf{Y}_a^{(1)}, \mathbf{Y}_a^{(2)}, \dots, \mathbf{Y}_a^{(k)}) \in \mathbb{F}_q^k$.

Finally, we use that $(M^{\otimes t})^{-1} = (M^{-1})^{\otimes t}$.

2.2 Information Theory Preliminaries

For the sake of completeness we include the information-theoretic concepts and tools we use in this paper.

For a discrete random variable X , let $H(X)$ denote its binary entropy:

$$H(X) := \sum_{a \in \text{Support}(X)} p_X(a) \log\left(\frac{1}{p_X(a)}\right)$$

where $p_X(a) := \Pr[X = a]$ is the probability mass function of X . Throughout, $\log(\cdot)$ by default denotes $\log_2(\cdot)$.

For $p \in [0, 1]$, we overload this notation, letting $H(p)$ denote the entropy $H(X)$ for $X \sim \text{Bernoulli}(p)$.

For arbitrary random variables X, Y , let $H(X|Y)$ denote the conditional entropy:

$$H(X|Y) = \mathbb{E}_Y[H(X|Y = y)].$$

For a q -ary random variable $X \in \mathbb{F}_q$, let $\overline{H}(X) \in [0, 1]$ denote its q -ary entropy:

$$\overline{H}(X) := \frac{H(X)}{\log(q)}.$$

Finally, the *mutual information* between jointly distributed random variables X, Y is:

$$I(X; Y) := H(X) - H(X|Y) = H(Y) - H(Y|X)$$

We will use the following standard properties of entropy:

1. **(Adding independent variables increases entropy):** For any random variables X, Y, Z such that X, Y are conditionally independent given Z , we have

$$H(X + Y|Z) \geq H(X|Z) \tag{3}$$

2. **(Transforming Conditioning):** For any random variables X, Y , any function f , and any bijection σ , we have

$$H(X|Y) = H(X + f(Y)|Y) = H(X + f(Y)|\sigma(Y)) \tag{4}$$

3. **(Chain rule):** For arbitrary random variables X, Y : $H(X, Y) = H(X) + H(Y|X)$.
4. **(Conditioning does not increase entropy):** For X, Y, Z arbitrary random variables, $H(X|Y, Z) \leq H(X|Y)$.
5. **(Monotonicity):** For $p \in [0, 1/2)$, the binary entropy $H(p)$ is non-decreasing with p . And for $p \in (1/2, 1]$, the binary entropy $H(p)$ is non-increasing with p .
6. **(Deterministic postprocessing does not increase entropy):** For arbitrary random variables X, Y and function f we have $H(X|Y) \geq H(f(X)|Y)$.
7. **(Conditioning on independent variables):** For random variables X, Y, Z where Z is independent from (X, Y) , we have $H(X|Y) = H(X|Y, Z)$.

2.2.1 Channels

Given a finite field \mathbb{F}_q , and output alphabet \mathcal{Y} , a q -ary channel $\mathcal{C}_{Y|Z}$ is a probabilistic function from \mathbb{F}_q to \mathcal{Y} . Equivalently, it is given by q probability distributions $\{\mathcal{C}_{Y|\alpha}\}_{\alpha \in \mathbb{F}_q}$ supported on \mathcal{Y} . We use notation $\mathcal{C}(Z)$ to denote the channel operating on inputs Z . A *memoryless channel* maps \mathbb{F}_q^n to \mathcal{Y}^n by acting independently (and identically) on each coordinate. A *symmetric channel* is a memoryless channel where for every $\alpha, \beta \in \mathbb{F}_q$ there is a bijection $\sigma : \mathcal{Y} \rightarrow \mathcal{Y}$ such that for every $y \in \mathcal{Y}$ it is the case that $\mathcal{C}_{Y=y|\alpha} = \mathcal{C}_{Y=\sigma(y)|\beta}$, and moreover for any pair $y_1, y_2 \in \mathcal{Y}$, we have $\sum_{x \in \mathbb{F}_q} \mathcal{C}_{Y=y_1|x} = \sum_{x \in \mathbb{F}_q} \mathcal{C}_{Y=y_2|x}$ (see, for example, [4, Section 7.2]). As shown by Shannon every memoryless channel has a finite capacity, denoted $\text{Capacity}(\mathcal{C}_{Y|Z})$. For symmetric channels, this is the mutual information $I(Y; Z)$ between the input Z and output Y where Z is drawn uniformly from \mathbb{F}_q and Y is drawn from $\mathcal{C}_{Y|Z}$ given Z .

2.3 Basic Probabilistic Inequalities

We first show that a random variable with small-enough entropy will usually take its most-likely value:

Lemma 2.1. *Let $X \in \mathbb{F}_q$ be a random variable. Then there exist \hat{x} such that*

$$\Pr[X \neq \hat{x}] \leq H(X)$$

and therefore

$$\Pr[X \neq \hat{x}] \leq \overline{H}(X) \log q.$$

Proof. Let $\alpha := H(X)$ and let $p_i := \Pr_X[X = i]$. Let $\hat{x} = \operatorname{argmax}_i\{p_i\}$ be the value maximizing this probability. Let $p_{\hat{x}} = 1 - \gamma$. We wish to show that $\gamma \leq \alpha$. If $\gamma \leq 1/2$ we have

$$\begin{aligned}
\alpha &= H(X) \\
&= \sum_i p_i \log \frac{1}{p_i} \\
&\geq \sum_{i \neq \hat{x}} p_i \log \frac{1}{p_i} && \text{(Since all summands are non-negative)} \\
&\geq \sum_{i \neq \hat{x}} p_i \log \frac{1}{\sum_{j \neq \hat{x}} p_j} && \text{(Since } p_i \leq \sum_{j \neq \hat{x}} p_j \text{.)} \\
&= \left(\sum_{i \neq \hat{x}} p_i \right) \cdot \log \left(\frac{1}{\sum_{j \neq \hat{x}} p_j} \right) \\
&= \gamma \cdot \log 1/\gamma \\
&\geq \gamma && \text{(Since } \gamma \leq 1/2 \text{ and so } \log 1/\gamma \geq 1 \text{)}
\end{aligned}$$

as desired. Now if $\gamma > 1/2$ we have a much simpler case since now we have

$$\begin{aligned}
\alpha &= H(X) \\
&= \sum_i p_i \log \frac{1}{p_i} \\
&\geq \sum_i p_i \log \frac{1}{p_{\hat{x}}} && \text{(Since } p_i \leq p_{\hat{x}} \text{)} \\
&= \log \frac{1}{p_{\hat{x}}} && \text{(Since } \sum_i p_i = 1 \text{)} \\
&= \log \frac{1}{1 - \gamma} \\
&\geq 1. && \text{(Since } \gamma \geq 1/2 \text{)}
\end{aligned}$$

But γ is always at most 1 so in this case also we have $\alpha \geq 1 \geq \gamma$ as desired. ■

For the decoder, we will need a conditional version of Lemma 2.1, saying that if a variable X has low conditional entropy conditioned on Y , then X can be predicted well given the instantiation of variable Y .

Lemma 2.2. *Let X, Y be arbitrary discrete random variables with range \mathcal{X}, \mathcal{Y} respectively. Then there exists a function $\hat{X} : \mathcal{Y} \rightarrow \mathcal{X}$ such that*

$$\Pr_{X,Y}[X \neq \hat{X}(Y)] \leq H(X|Y)$$

In particular, the following estimator satisfies this:

$$\hat{X}(y) := \operatorname{argmax}_x \{\Pr[X = x|Y = y]\}$$

Proof. For every setting of $Y = y$, we can bound the error probability of this estimator using

Lemma 2.1 applied to the conditional distribution $X|Y = y$:

$$\begin{aligned} \Pr_{X,Y}[X \neq \hat{X}(Y)] &= \mathbb{E}_Y[\Pr_{X|Y}[\hat{X}(Y) \neq X]] \\ &\leq \mathbb{E}_Y[H(X|Y = y)] && \text{(Lemma 2.1)} \\ &= H(X|Y) \end{aligned}$$

■

We will need an inverse to the usual Chebychev inequality. Recall that Chebychev shows that variables with small variance are concentrated close to their expectation. The Paley-Zygmund inequality below can be used to invert it (somewhat) — for a random variable W with comparable fourth and second central moment, by applying the lemma below to $Z = (W - \mathbb{E}[W])^2$ we can deduce that W has positive probability of deviating noticeably from the mean.

Lemma 2.3 (Paley-Zygmund). *If $Z \geq 0$ is a random variable with finite variance, then*

$$\Pr(Z > \lambda \mathbb{E}[Z]) \geq (1 - \lambda)^2 \frac{\mathbb{E}[Z]^2}{\mathbb{E}[Z^2]}.$$

Next, we define the notion of a sequence of random variables being adapted to another sequence of variables, which will be useful in our later proofs.

Definition 2.4. *We say that a sequence $Y_1, Y_2 \dots$ of random variables is adapted to the sequence $X_1, X_2 \dots$ if and only if for every t , Y_t is completely determined given X_1, \dots, X_t . We will use $\mathbb{E}[Z|X_{[1:t]}]$ as a shorthand $\mathbb{E}[Z|X_1, \dots, X_t]$, and $\Pr[E|X_{[1:t]}]$ as a shorthand for $\mathbb{E}[\mathbb{1}_E|X_1, \dots, X_t]$. If the underlying sequence X is clear from context, we will skip it and write just $\mathbb{E}[Z|\mathcal{F}_t]$.*

Lemma 2.5. *Consider a sequence of non-negative random variables $Y_1, Y_2, \dots, Y_t, \dots$ adapted to the sequence X_t . If for every t we have $\Pr(Y_{t+1} > \lambda | X_{[1:t]}) < \exp(-\lambda)$, then for every $T > 0$:*

$$\Pr\left(\sum_{i \leq T} Y_i > CT\right) \leq \exp(-\Omega(T))$$

for some universal constant C .

Proof. First, observe that

$$\begin{aligned} \mathbb{E}[\exp(Y_{t+1}/2)|\mathcal{F}_t] &= \int_0^\infty \Pr(\exp(Y_{t+1}/2) > \lambda | \mathcal{F}_t) d\lambda \\ &\leq 1 + \int_1^\infty \exp(-2 \log \lambda) d\lambda \\ &= 1 + \int_1^\infty \lambda^{-2} d\lambda \\ &\leq \exp(C_0) \end{aligned} \tag{5}$$

for some constant C_0 . On the other hand, we have decomposition (where we apply (5) in the first

equality):

$$\begin{aligned}
\mathbb{E}[\exp(\sum_{i \leq T} \frac{Y_i}{2})] &= \mathbb{E}[\mathbb{E}[\exp(\sum_{i \leq T} \frac{Y_i}{2}) | \mathcal{F}_{T-1}]] \\
&= \mathbb{E}[\exp(\sum_{i \leq T-1} Y_i/2) \mathbb{E}[\exp(Y_T/2) | \mathcal{F}_{T-1}]] \\
&\leq \mathbb{E}[\exp(\sum_{i \leq T-1} Y_i/2)] \exp(C_0) \\
&\leq \dots \\
&\leq \exp(C_0 T).
\end{aligned}$$

Now we can apply Markov inequality to obtain desired tail bounds:

$$\Pr(\sum_{i \leq T} Y_i > 4C_0 T) = \Pr(\exp(\frac{1}{2} \sum_{i \leq T} Y_i) > \exp(2C_0 T)) \leq \mathbb{E} \left[\exp(\frac{1}{2} \sum_{i \leq T} Y_i) \right] \exp(-2C_0 T) \leq \exp(-C_0 T).$$

■

Lemma 2.6. *Consider a sequence of random variables Y_1, Y_2, \dots with $Y_i \in \{0, 1\}$, adapted to the sequence X_t . If $\Pr(Y_{t+1} = 1 | X_{[1:t]}) > \mu_{t+1}$ for some deterministic value μ_t , then for $\mu := \sum_{t \leq T} \mu_t$ we have*

$$\Pr(\sum_{t \leq T} Y_t < \mu/2) \leq \exp(-\Omega(\mu))$$

Proof. Let $M_{t+1} := \mathbb{E}[Y_{t+1} | X_{[1:t]}]$, we know that $M_t > \mu_t$ with probability 1. Standard calculation involving Markov inequality yields following bound

$$\begin{aligned}
\Pr(\sum_{t \leq T} Y_t < \sum_{t \leq T} M_t/2) &= \Pr(\exp(-\sum_{t \leq T} Y_t + \sum_{t \leq T} M_t/2) > 1) \\
&\leq \mathbb{E}[\exp(\sum_{t \leq T} (-Y_t + M_t/2))] \\
&= \mathbb{E}[\mathbb{E}[\exp(\sum_{t \leq T} (-Y_t + M_t/2)) | X_{[1:T-1]}]] \\
&\leq \mathbb{E}[\exp(\sum_{t \leq T-1} (-Y_t + M_t/2)) \mathbb{E}[\exp(-Y_T + M_T/2) | X_{[1:T-1]}]]. \quad (6)
\end{aligned}$$

We now observe that for any random variable $\tilde{Y} \in \{0, 1\}$ with $\mathbb{E}[\tilde{Y}] = p$, we have

$$\log \mathbb{E}[\exp(-\tilde{Y} + p/2)] = \frac{p}{2} + \log[(1-p) + \frac{p}{e}] \leq \frac{p}{2} - p + \frac{p}{e} \leq -cp$$

with constant $c = (1 - \frac{1}{2} - \frac{1}{e}) > 0$. In particular $\mathbb{E}[\exp(-Y_T + M_T/2) | X_{[1:T-1]}] \leq \exp(-cM_T) \leq \exp(-c\mu_T)$. Plugging this back to (6), we get

$$\begin{aligned}
\Pr(\sum_{t \leq T} Y_t < \sum_{t \leq T} M_t/2) &\leq \mathbb{E}[\exp(\sum_{t \leq T-1} (-Y_t + M_t/2))] \exp(-c'\mu_T) \\
&\leq \dots \\
&\leq \exp(-c \sum_{t \leq T} \mu_t) \\
&= \exp(-c\mu)
\end{aligned}$$

And moreover, since $M_t > \mu_t$ deterministically, we have $\Pr(\sum_{t \leq T} Y_t < \mu/2) \leq \Pr(\sum_{t \leq T} Y_t < \sum M_t/2) \leq \exp(-\Omega(\mu))$ as desired. ■

Finally, we will use the well-known Doob's martingale inequality:

Lemma 2.7 (Doob's martingale inequality [6, Theorem 5.4.2]). *If a sequence X_0, X_1, \dots is a martingale, then for every T we have*

$$\Pr(\sup_{t \leq T} X_t > \lambda) \leq \frac{\mathbb{E}[|X_T|]}{\lambda}$$

Corollary 2.8. *If X_0, X_1, \dots is a nonnegative martingale, then for every T we have*

$$\Pr(\sup_{t \leq T} X_t > \lambda) \leq \frac{\mathbb{E}[X_0]}{\lambda}$$

3 Local to global polarization

In this section we prove Theorem 1.6, which asserts that every locally polarizing $[0, 1]$ -martingale is also strongly polarizing. The proofs in this section depend on some basic probabilistic concepts and inequalities that we have seen in Section 2.3.

The proof of this statement is implemented in two main steps: first, we show that any locally polarizing martingale, is $((1 - \frac{\nu}{2})^t, (1 - \frac{\nu}{4})^t)$ -polarizing for *some* constant ν depending only on the parameters α, τ, θ of local polarization. This means that, except with exponentially small probability, $\min\{X_{t/2}, 1 - X_{t/2}\}$ is exponentially small in t , which we can use to ensure that X_s for all $\frac{t}{2} \leq s \leq t$ stays in the range where the conditions of *suction at the ends* apply (again, except with exponentially small failure probability). Finally, we show that if the martingale stays in the *suction at the ends* regime, it will polarize strongly — i.e. if we have a $[0, 1]$ -martingale, such that in each step it has probability at least α to decrease by a factor of C , we can deduce that at the end we have $\Pr(X_T > C^{-\alpha T/4}) \leq \exp(-\Omega(\alpha T))$.

We start by showing that in the first $t/2$ steps we do get exponentially small polarization, with all but exponentially small failure probability. This is proved using a simple potential function $\min\{\sqrt{X_t}, \sqrt{1 - X_t}\}$ which we show shrinks by a constant factor, $1 - \nu$ for some $\nu > 0$, in expectation at each step. Previous analyses in [11, 10] tracked $\sqrt{X_t(1 - X_t)}$ (or some tailormade algebraic functions [13, 17]) as potential functions, and relied on quantitatively strong forms of variance in the middle to demonstrate that the potential diminishes by a constant factor in each step. While such analyses can lead to sharper bounds on the parameter ν , which in turn translate to better *scaling exponents* in the polynomial convergence to capacity, e.g. see [13, Thm. 18] or [17, Thm. 1], these analyses are more complex, and less general.

Lemma 3.1. *If a $[0, 1]$ -martingale sequence X_0, \dots, X_t, \dots , is $(\alpha, \tau(\cdot), \theta(\cdot))$ -locally polarizing, then there exist $\nu > 0$, depending only on α, τ, θ , such that*

$$\mathbb{E}[\min(\sqrt{X_t}, \sqrt{1 - X_t})] \leq (1 - \nu)^t.$$

Proof. Take $\tau_0 = \tau(4), \theta_0 = \theta(\tau_0)$. We will show that $\mathbb{E}[\min(\sqrt{X_{t+1}}, \sqrt{1 - X_{t+1}}) | X_t] \leq (1 - \nu) \min(\sqrt{X_t}, \sqrt{1 - X_t})$, for some $\nu > 0$ depending on τ_0, θ_0 and α . The statement of the lemma will follow by induction.

Let us condition on X_t , and first consider the case $X_t \in (\tau_0, 1 - \tau_0)$. We know that

$$\mathbb{E}[\min(\sqrt{X_{t+1}}, \sqrt{1 - X_{t+1}})] \leq \min(\mathbb{E}[\sqrt{X_{t+1}}], \mathbb{E}[\sqrt{1 - X_{t+1}}]),$$

we will show that $\mathbb{E}[\sqrt{X_{t+1}}] \leq (1 - \nu)\sqrt{X_t}$. The proof of $\mathbb{E}[\sqrt{1 - X_{t+1}}] \leq (1 - \nu)\sqrt{1 - X_t}$ is symmetric.

Indeed, let us take $T := \sqrt{\frac{X_{t+1}}{X_t}}$. Because $(X_t)_t$ is a martingale, we have $\mathbb{E}[T^2] = 1$, and by Jensen's inequality, we have that $\mathbb{E}[T] \leq \sqrt{\mathbb{E}[T^2]} \leq 1$, where all the expectations above are conditioned on X_t . Take δ such that $\mathbb{E}[T] = 1 - \delta$. We will show a lower bound on δ in terms of θ_0, τ_0 and α_0 .

The high-level idea of the proof is that we can show that local polarization criteria implies that T is relatively far from 1 with noticeable probability, but if $\mathbb{E}[T]$ were close to one, by Chebyshev inequality we would be able to deduce that T is far from its mean with much smaller probability. This implies that mean of T has to be bounded away from 1.

More concretely, observe first that by Chebyshev inequality, we have $\Pr(|T - \mathbb{E}[T]| > \lambda) < \frac{\text{Var}(T)}{\lambda^2} = \frac{2\delta - \delta^2}{\lambda^2} \leq \frac{2\delta}{\lambda^2}$, hence, for $C_0 = 4$, we have:

$$\Pr\left(|T - 1| \geq \delta + C_0\sqrt{\delta}\theta_0^{-1}\tau_0^{-1}\right) \leq \frac{1}{8}\theta_0^2\tau_0^2. \quad (7)$$

On the other hand, because of the *Variation in the middle condition* of local polarization, we have

$$\text{Var}(T^2) = \frac{\mathbb{E}[X_{t+1}^2] - X_t^2}{X_t^2} \geq \frac{\theta_0}{X_t^2} \geq \theta_0,$$

where the last inequality follows since $X_t \leq 1$. Moreover $T < \frac{1}{\sqrt{\tau_0}}$, because $\sqrt{X_{t+1}} < 1$ and $\sqrt{X_t} > \sqrt{\tau_0}$.

Let us now consider $Z = (T^2 - 1)^2$. We have $\mathbb{E}[Z] = \text{Var}(T^2) \geq \theta_0$, and moreover $\mathbb{E}[Z^2] < \tau_0^{-2}$ (because T is bounded and $\tau_0 \leq 1$), hence by Lemma 2.3 (for $C_1 = 1/2$)

$$\Pr\left((1 - T^2)^2 > C_1\theta_0\right) \geq \frac{1}{4}\theta_0^2\tau_0^2.$$

And also $1 - T^2 = -(1 - T)^2 + 2(1 - T) < 2(1 - T)$, hence if $(1 - T^2)^2 > C_1\theta_0$ then $|1 - T| > \frac{\sqrt{C_1}}{2}\sqrt{\theta_0}$, which implies (for the choice of $C_2 = \sqrt{C_1}/2$):

$$\Pr(|T - 1| > C_2\sqrt{\theta_0}) \geq \frac{1}{4}\theta_0^2\tau_0^2. \quad (8)$$

By comparing (7) and (8), we deduce that $\delta \geq C_4\theta_0^3\tau_0^2$, (for $C_4 = C_2^2/(4C_0^2)$ – note that with our choice of parameters, we have $C_0\sqrt{\delta}\theta_0^{-1}\tau_0^{-1} \geq \delta$) and by the definition of δ we have $\mathbb{E}[\sqrt{X_{t+1}}|X_t] \leq (1 - \delta)\sqrt{X_t}$. The same argument applies to show that $\mathbb{E}[\sqrt{1 - X_{t+1}}|X_t] \leq (1 - C_4\theta_0^3\tau_0^2)\sqrt{1 - X_t}$.

Consider now the case when $X_t < \tau_0$. For T, δ as above (and again after conditioning on X_t), we have $\text{Var}(T) < 2\delta$ (note that the argument for this inequality from the previous case also holds here), and hence by Chebyshev inequality (for the choice of $C_5 = 2$):

$$\Pr\left(|T - 1| \geq \delta + C_5\sqrt{\frac{\delta}{\alpha}}\right) \leq \frac{\alpha}{2}. \quad (9)$$

On the other hand, because of the *suction at the end* condition of local polarization, we know that with probability α , we have $T \leq \frac{1}{2}$, which means $|T - 1| \geq \frac{1}{2}$ and by comparing this with (9),

we deduce $\delta \geq C_6\alpha$ (for $C_6 = \frac{1}{16C_5^2}$). Therefore, in the case $X_t < \tau_0$, we have $\mathbb{E}[\sqrt{X_{t+1}}|X_t] \leq (1 - C_6\alpha)\sqrt{X_t} = (1 - C_6\alpha) \min(\sqrt{X_t}, \sqrt{1 - X_t})$. The case $X_t > 1 - \tau_0$ is symmetric and is omitted.

This implies the statement of the lemma with $\nu = \min(C_6\alpha, C_4\theta_0^3\tau_0^2)$. \blacksquare

Corollary 3.2. *If a $[0, 1]$ -martingale sequence X_0, \dots, X_t, \dots , is $(\alpha, \tau(\cdot), \theta(\cdot))$ -locally polarizing, then there exist $\nu > 0$, depending only on α, τ, θ , such that*

$$\Pr \left[\min(X_{t/2}, 1 - X_{t/2}) > \lambda(1 - \frac{\nu}{2})^t \right] \leq (1 - \frac{\nu}{4})^t \frac{1}{\sqrt{\lambda}}.$$

Proof. By applying Markov Inequality to the bound from Lemma 3.1 (with $t/2$ instead of t), we get

$$\begin{aligned} \Pr \left[\min(X_{t/2}, 1 - X_{t/2}) > \lambda(1 - \frac{\nu}{2})^t \right] &= \Pr \left[\min(\sqrt{X_{t/2}}, \sqrt{1 - X_{t/2}}) > \sqrt{\lambda}(1 - \frac{\nu}{2})^{t/2} \right] \\ &\leq (1 - \nu)^{t/2} (1 - \frac{\nu}{2})^{-t/2} \frac{1}{\sqrt{\lambda}} \\ &\leq (1 - \frac{\nu}{4})^t \frac{1}{\sqrt{\lambda}}. \end{aligned}$$

\blacksquare

The next lemma will be used to show that if a $[0, 1]$ -martingale indeed stays at all steps $s \geq \frac{t}{2}$ in the *suction at the ends* range, i.e. in each step it has constant probability α of dropping by some large constant factor C , then expect it to be $(C^{-\alpha t/8}, \exp(-\Omega(\alpha t)))$ -polarized.

Lemma 3.3. *There exists $c < \infty$, such that for all K, α with $K\alpha \geq c$ the following holds. Let X_t be a martingale satisfying $\Pr(X_{t+1} < e^{-K}X_t | X_t) \geq \alpha$, where $X_0 \in (0, 1)$. Then $\Pr(X_T > \exp(-\alpha KT/4)) \leq \exp(-\Omega(\alpha T))$.*

Proof. Consider $Y_{t+1} := \log \frac{X_{t+1}}{X_t}$, and note that sequence Y_t is adapted to sequence X_t in the sense of Definition 2.4. We have the following bounds on the upper tails of Y_{t+1} , conditioned on $X_{[1:t]}$, given by Markov inequality

$$\Pr(Y_{t+1} > \lambda | \mathcal{F}_t) = \Pr \left(\frac{X_{t+1}}{X_t} > \exp(\lambda) \mid X_{[1:t]} \right) = \Pr(X_{t+1} > \exp(\lambda)X_t | X_{[1:t]}) \leq \exp(-\lambda).$$

Let us decompose $Y_{t+1} =: (Y_{t+1})_+ + (Y_{t+1})_-$, where $(Y_{t+1})_+ := \max(Y_{t+1}, 0)$. By Lemma 2.5,

$$\Pr \left(\sum_{t \leq T} (Y_{t+1})_+ > CT \right) \leq \exp(-\Omega(T)).$$

On the other hand, let E_{t+1} be the indicator of $Y_{t+1} \leq -K$. It is again adapted to the sequence X_t , and we know that $\Pr(E_{t+1} | X_{[1:t]}) \geq \alpha$, hence by Lemma 2.6 with probability at most $\exp(-\Omega(\alpha T))$ at most $\alpha T/2$ of those events holds. Note that $(Y_t)_- \leq 0$, which implies that if at least $\alpha T/2$ of the events E_t hold then we have $\sum_{t \leq T} (Y_t)_- \leq -\alpha KT/2$. Thus, we have $\Pr(\sum_{t \leq T} (Y_t)_- > -\alpha KT/2) \leq \exp(-\Omega(\alpha T))$. Therefore, as long as $\alpha K/4 > C$, we can conclude

$$\Pr \left(\sum_{t \leq T} Y_t > -\alpha KT/4 \right) \leq \exp(-\Omega(T)) + \exp(-\Omega(\alpha T)) \leq \exp(-\Omega(\alpha T)).$$

The proof is complete by noting that $\sum_{t \leq T} Y_t = \log(X_T/X_0)$ and recalling that $X_0 \leq 1$. \blacksquare

Proof of Theorem 1.6. For given γ , we take K to be large enough so that $\exp(-\alpha K/8) \leq \gamma$, and moreover αK to be large enough to satisfy assumptions of Lemma 3.3. Let us also take $\tau_0 = \tau(e^K)$. We consider ν as in Corollary 3.2. We have

$$\Pr \left(\min(X_{t/2}, 1 - X_{t/2}) > \left(1 - \frac{\nu}{2}\right)^t \tau_0 \right) \leq \left(1 - \frac{\nu}{4}\right)^{-t} \frac{1}{\sqrt{\tau_0}}.$$

Now Doob's martingale inequality (Corollary 2.8) implies that, conditioned on $X_{t/2} < (1 - \frac{\nu}{4})^t \tau_0$, we have $\Pr \left(\sup_{i \in (t/2, t)} X_i > \tau_0 \right) \leq (1 - \frac{\nu}{4})^t$.

Finally, after conditioning on $X_i \leq \tau_0, \forall t/2 \leq i \leq t$, process X_i for $i \in (t/2, t)$ satisfies conditions of Lemma 3.3, because X_i always stays below τ_0 and as such *suction at the end* condition of local polarization corresponds exactly to the assumption in this Lemma. Therefore we can conclude that except with probability $\exp(-\Omega(\alpha t))$, we have $X_t < \exp(-\alpha K t/8) = \gamma^t$. The other case ($1 - X_{t/2} < (1 - \frac{\nu}{2})^t \tau_0$) is symmetric, and in this case we get $1 - X_t < \exp(-\alpha K t/8)$ except with probability $\exp(-\Omega(\alpha t))$. ■

4 Arikan Martingale

We now formally describe the Arikan martingale associated with an invertible matrix $M \in \mathbb{F}_q^{k \times k}$ and a channel $\mathcal{C}_{Y|Z}$. Briefly, this martingale measures at time t , the conditional entropy of a random variable \mathbf{A}'_i , conditioned on the values of a vector of variables \mathbf{B}' and on the values of \mathbf{A}'_j for j smaller than i for a random choice of the index i . Here \mathbf{A}' is a vector of k^t random variables taking values in \mathbb{F}_q while $\mathbf{B}' \in \mathcal{Y}^{k^t}$. The exact construction of the joint distribution of these $2k^t$ variables is the essence of the Arikan construction of codes, and we describe it shortly. The hope with this construction is that eventually (for large values of t) the conditional entropies are either very close to 0, or very close to $\log q$ for most choices of i .

When $t = 1$, the process starts with k independent and identical pairs of variables $\{(A_i, B_i)\}_{i \in [k]}$ where $A_i \sim \mathbb{F}_q$ and $B_i \sim \mathcal{C}_{Y|Z=A_i}$. (So each pair corresponds to an independent input/output pair from transmission of a uniformly random input over the channel $\mathcal{C}_{Y|Z}$.) Let $\mathbf{A} = (A_1, \dots, A_k)$ and $\mathbf{B}' = (B_1, \dots, B_k)$, and note that the conditional entropies $H(A_i | \mathbf{A}_{\prec i}, \mathbf{B}')$ are all equal, and this entropy, divided by $\log_2 q$, will be our value of X_0 . On the other hand, if we now let $\mathbf{A}' = \mathbf{A} \cdot M$ then the conditional entropies $H(\mathbf{A}'_i | \mathbf{A}'_{\prec i}, \mathbf{B}')$ are no longer equal (for most, and in particular for all mixing, matrices M). On the other hand, conservation of conditional entropy on application of an invertible transformation tells us that $\mathbb{E}_{i \sim [k]} [H(\mathbf{A}'_i | \mathbf{A}'_{\prec i}, \mathbf{B}') / \log_2 q] = X_0$. Thus letting $X_1 = H(\mathbf{A}'_i | \mathbf{A}'_{\prec i}, \mathbf{B}') / \log_2 q$ (for random i) gives us the martingale at time $t = 1$.

While this one step of multiplication by M *differentiates* among the k (previously identical) random variables, it doesn't yet polarize. The hope is by iterating this process one can get polarization³. But to get there we need to describe how to iterate this process. This iteration is conceptually simple (though notationally still complex) and illustrated in Figure 1. Roughly the idea is that at the beginning of stage t , we have defined a joint distribution of k^t dimensional vectors (\mathbf{A}, \mathbf{B}) along with a multi-index $\mathbf{i} \in [k]^t$. We now sample k independent and identically distributed pairs of these random variables $\{(\mathbf{A}^{(\ell)}, \mathbf{B}^{(\ell)})\}_{\ell \in [k]}$ and view $(\mathbf{A}^{(\ell)})_{\ell \in [k]}$ as a $k^t \times k$ matrix which we multiply by M to get a new $k^t \times k$ matrix. Flattening this matrix into a k^{t+1} -dimensional vector gives us a sample from the distribution of $\mathbf{A}' \in \mathbb{F}_q^{k^{t+1}}$. \mathbf{B}' is simply the concatenation of all the vectors $(\mathbf{B}^{(\ell)})_{\ell \in [k]}$. And finally the new index $\mathbf{j} \in [k]^{t+1}$ is simply obtained by extending $\mathbf{i} \in [k]^t$

³In the context of Polar coding, *differentiation* and *polarization* are good events, and hence our "hope."

with a $(t + 1)$ th coordinate distributed uniformly at random in $[k]$. X_{t+1} is now defined to be $H(\mathbf{A}'_j | \mathbf{A}'_{\prec j}, \mathbf{B}') / \log_2 q$. The formal description is below.

Definition 4.1 (Arikan martingale). *Given an invertible matrix $M \in \mathbb{F}_q^{k \times k}$ and a channel description $C_{Y|Z}$ for $Z \in \mathbb{F}_q, Y \in \mathcal{Y}$, the Arikan-martingale X_0, \dots, X_t, \dots associated with it is defined as follows. For every $t \in \mathbb{N}$, let D_t be the distribution on pairs $\mathbb{F}_q^{k^t} \times \mathcal{Y}^{k^t}$ described inductively below:*

A sample (A, B) from D_0 supported on $\mathbb{F}_q \times \mathcal{Y}$ is obtained by sampling $A \sim \mathbb{F}_q$, and $B \sim C_{Y|Z=A}$. For $t \geq 1$, a sample $(\mathbf{A}', \mathbf{B}') \sim D_t$ supported on $\mathbb{F}_q^{k^t} \times \mathcal{Y}^{k^t}$ is obtained as follows:

- Draw k independent samples $(\mathbf{A}^{(1)}, \mathbf{B}^{(1)}), \dots, (\mathbf{A}^{(k)}, \mathbf{B}^{(k)}) \sim D_{t-1}$.
- Let \mathbf{A}' be given by $\mathbf{A}'_{[i, \cdot]} = (\mathbf{A}_i^{(1)}, \dots, \mathbf{A}_i^{(k)}) \cdot M$ for all $i \in [k]^{t-1}$ and $\mathbf{B}' = (\mathbf{B}^{(1)}, \mathbf{B}^{(2)}, \dots, \mathbf{B}^{(k)})$.

Then, the sequence X_t is defined as follows: For each $t \in \mathbb{N}$, sample $i_t \in [k]$ iid uniformly. Let $\mathbf{j} = (i_1, \dots, i_t)$ and let $X_t := H(\mathbf{A}_{\mathbf{j}} | \mathbf{A}_{\prec \mathbf{j}}, \mathbf{B}) / \log_2 q$, where the entropies are with respect to the distribution $(\mathbf{A}, \mathbf{B}) \sim D_t$.⁴

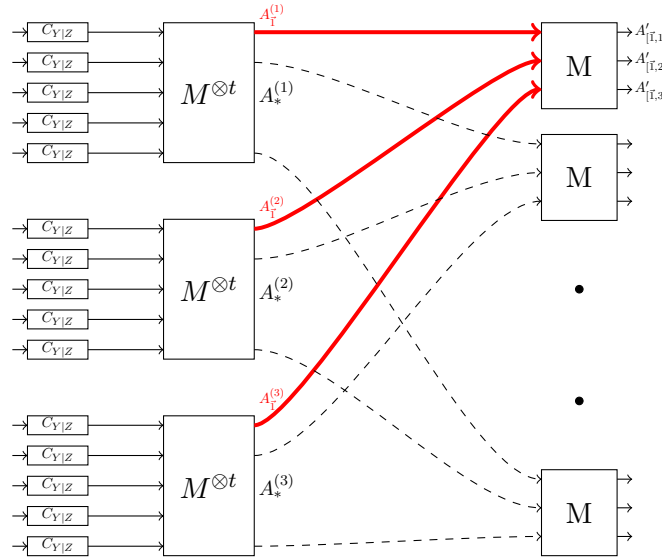


Figure 1: Evolution of Arikan martingale for 3×3 matrix M .

Figure 1 illustrates the definition by highlighting the construction of the vector \mathbf{A}' , and in particular highlights the recursive nature of the construction.

It is easy (and indeed no different than in the case $t = 1$) to show that $\mathbb{E}[X_{t+1} | X_t] = X_t$ and so the Arikan martingale is indeed a martingale. This is shown below.

Proposition 4.2. *For every matrix M and channel $C_{Y|Z}$, the Arikan martingale is a martingale and in particular a $[0, 1]$ -martingale.*

Proof. The fact that $X_t \in [0, 1]$ follows from the fact for $0 \leq H(\mathbf{A}_i | \mathbf{A}_{\prec i}, \mathbf{B}) \leq H(\mathbf{A}_i) \leq \log_2 q$ and so $0 \leq X_t = H(\mathbf{A}_i | \mathbf{A}_{\prec i}, \mathbf{B}) / \log_2 q \leq 1$.

⁴We stress that the only randomness in the evolution of X_t is in the choice of i_1, \dots, i_t, \dots . The process of sampling \mathbf{A} and \mathbf{B} is only used to define the distributions for which we consider the conditional entropies $H(\mathbf{A}_{\mathbf{j}} | \mathbf{A}_{\prec \mathbf{j}}, \mathbf{B})$.

We turn to showing that $\mathbb{E}[X_{t+1}|X_t = a] = a$. To this end, consider a sequence of indices $\mathbf{i} = (i_1, \dots, i_t)$, such that $\overline{H}(\mathbf{A}_{\mathbf{i}} | \mathbf{A}_{\prec \mathbf{i}}, \mathbf{B}) = a$. We wish to show that $\mathbb{E}_{i_{t+1} \sim [k]}[\overline{H}(\mathbf{A}'_{[i, i_{t+1}]} | \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}')] = a$.

Since the pairs $(\mathbf{A}^{(s)}, \mathbf{B}^{(s)})$ are independent, note that for any s , we have $\overline{H}(\mathbf{A}_{\mathbf{i}}^{(s)} | \mathbf{A}_{\prec \mathbf{i}}^{(s)}, \mathbf{B}^{(s)}) = a$. Furthermore, because of the same independence, we have

$$\begin{aligned} \overline{H}(\mathbf{A}_{\mathbf{i}}^{(s)} | \mathbf{A}_{\prec \mathbf{i}}^{(s)}, \mathbf{B}^{(s)}) &= \overline{H}(\mathbf{A}_{\mathbf{i}}^{(s)} | \cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}, \cup_{j \in [k]} \mathbf{B}^{(j)}) \\ \text{and } \overline{H}(\mathbf{A}_{\mathbf{i}}^{(1)}, \dots, \mathbf{A}_{\mathbf{i}}^{(k)} | \cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}, \cup_{j \in [k]} \mathbf{B}^{(j)}) &= k \cdot a. \end{aligned}$$

By the invertibility of M we have

$$\overline{H}(\mathbf{A}'_{[i, 1]}, \dots, \mathbf{A}'_{[i, k]} | \cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}, \cup_{j \in [k]} \mathbf{B}^{(j)}) = \overline{H}(\mathbf{A}_{\mathbf{i}}^{(1)}, \dots, \mathbf{A}_{\mathbf{i}}^{(k)} | \cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}, \cup_{j \in [k]} \mathbf{B}^{(j)}) = k \cdot a.$$

We can apply again invertibility of the matrix M to deduce that conditioning on $\cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}$ is the same as conditioning on $\mathbf{A}'_{\prec [i, 1]}$ — i.e. for any multiindex $\mathbf{i}' \prec \mathbf{i}$ variables $\mathbf{A}_{\mathbf{i}'}^{(1)}, \dots, \mathbf{A}_{\mathbf{i}'}^{(k)}$ and $\mathbf{A}'_{[\mathbf{i}', 1]}, \dots, \mathbf{A}'_{[\mathbf{i}', k]}$ are related via invertible transform M . This yields

$$\overline{H}(\mathbf{A}'_{[i, 1]}, \dots, \mathbf{A}'_{[i, k]} | \mathbf{A}'_{\prec [i, 1]}, \mathbf{B}') = \overline{H}(\mathbf{A}'_{[i, 1]}, \dots, \mathbf{A}'_{[i, k]} | \cup_{j \in [k]} \mathbf{A}_{\prec \mathbf{i}}^{(j)}, \cup_{j \in [k]} \mathbf{B}^{(j)}) = ka.$$

Finally by the Chain rule of entropy we have

$$\begin{aligned} \overline{H}(\mathbf{A}'_{[i, 1]}, \dots, \mathbf{A}'_{[i, k]} | \mathbf{A}'_{\prec [i, 1]}, \mathbf{B}') &= \sum_{i_{t+1}=1}^k \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} | \mathbf{A}'_{[i, \prec i_{t+1}]}, \mathbf{A}'_{\prec [i, 1]}, \mathbf{B}') \\ &= \sum_{i_{t+1}=1}^k \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} | \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') \end{aligned}$$

Putting these together, we have $\mathbb{E}[X_{t+1}|X_t = a] = \mathbb{E}_{i_{t+1}}[H(\mathbf{A}'_{[i, i_{t+1}]} | \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}')] = a$. ■

Finally, we remark that based on the construction it is not too hard to see that if M were an identity matrix, or more generally a non-mixing matrix, then X_t would deterministically equal X_0 . (There is no differentiation and thus no polarization.) The thrust of this paper is to show that in all other cases we have strong polarization.

5 Proof of Local Polarization

In this section we prove Theorem 1.10, which states that the Arikan martingale is locally polarizing, modulo some entropic inequalities. In Section 6, we prove these inequalities.

We first start with an overview of the proof.

5.1 Proof Overview

Here we describe the overall structure of the proof that the Arikan martingale is locally polarizing. First we recall the theorem we would like to prove:

Theorem 1.10. *For every prime q , for every mixing matrix $M \in \mathbb{F}_q^{k \times k}$, and for every symmetric memoryless channel $\mathcal{C}_{Y|Z}$ over \mathbb{F}_q , the associated Arikan martingale sequence is locally polarizing.*

The main ideas are roughly as follows. Let $\mathbf{A} \in \mathbb{F}_q^k$ be a random vector, and let W be an arbitrary random variable. Suppose the entries of \mathbf{A} are independent and identically-distributed, conditioned on W . Let $X_0 := \overline{H}(\mathbf{A}_1|W)$ be the conditional entropy of each entry of \mathbf{A} .

Let $\mathbf{A}' := \mathbf{A} \cdot M$, corresponding to variables in the next step of polarization. Local polarization of the Arikan martingale boils down to showing that for a random index $i \in [k]$, the conditional entropies of the transformed variables $X_1 \sim \overline{H}(\mathbf{A}'_i|\mathbf{A}'_{\prec i}, W)$ satisfy the local polarization conditions. Note that by conservation of entropy (since M is invertible), $\mathbb{E}_i[X_1] = X_0$.

In particular, it is sufficient to show that:

1. **(Variance in the middle):** There is some index $j \in [k]$ for which the following holds: For every $\tau > 0$, there exists $\varepsilon > 0$ such that if $X_0 := \overline{H}(\mathbf{A}_1|W) \in (\tau, 1 - \tau)$, then

$$\overline{H}(\mathbf{A}'_j|\mathbf{A}'_{\prec j}, W) \geq \overline{H}(\mathbf{A}_1|W) + \varepsilon.$$

This implies variance-in-the-middle, since with constant probability ($1/k$) the index j will be chosen, and in this case $X_1 = \overline{H}(\mathbf{A}'_j|\mathbf{A}'_{\prec j}, W) \geq X_0 + \varepsilon$. Thus $\mathbb{E}[(X_1 - X_0)^2|X_0] \geq \varepsilon'$ for some constant ε' .

2. **(Suction at the lower end):** There is some index $j \in [k]$ for which the following holds: For every $c < \infty$, there exists $\tau > 0$ such that if $X_0 := \overline{H}(\mathbf{A}_1|W) < \tau$ then

$$\overline{H}(\mathbf{A}'_j|\mathbf{A}'_{\prec j}, W) \leq \frac{1}{c} \overline{H}(\mathbf{A}_1|W).$$

This implies suction at the low end, since with constant probability ($1/k$) the index j will be chosen, in which case the entropy drops by at least c . Thus $\Pr_i[X_1 \leq X_0/c] \geq 1/k$.

3. **(Suction at the high end):** Analogously to suction at the low end, it is sufficient to show that there is some index $j \in [k]$ for which the following holds: For every $c < \infty$, there exists $\tau > 0$ such that if $X_0 := \overline{H}(\mathbf{A}_1|W) > 1 - \tau$ then

$$1 - \overline{H}(\mathbf{A}'_j|\mathbf{A}'_{\prec j}, W) \leq \frac{1}{c}(1 - \overline{H}(\mathbf{A}_1|W)).$$

In Section 5.2 we state three inequalities relating conditional entropy of a sum of random two random variables, with entropy of each of those random variables — these inequalities can be used to show that Arikan martingale for matrix $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$ is locally polarizing — each condition of local polarization can be deduced directly from the corresponding entropic inequality.

In Section 5.3 we show that using Gaussian Elimination we can reduce showing local polarization of any mixing $k \times k$ matrix to the very same entropic inequalities that are proven in Section 5.2.

5.2 Entropic Lemmas in the 2×2 Case

In this section, we state entropic inequalities which hold the key to proving the local polarization property. For 2×2 matrix $G_2 = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$, each condition of the local polarization can be almost directly deduced from the corresponding entropic inequalities. Later in Section 5.3 we will see that with some extra effort we can reduce local polarization conditions of any $k \times k$ mixing matrix to the same entropic inequalities, naturally arising in the proof of local polarization of Arikan martingale for G_2 .

The proofs of these lemmas are deferred to Section 6.

The following lemma corresponds to *Suction at the upper end* (for $X_t > 1 - \tau$).

Lemma 5.1. *For every finite field \mathbb{F}_q and every $\gamma > 0$, there exist τ , such that if (X_1, A_1) and (X_2, A_2) are independent random variables with $X_i \in \mathbb{F}_q$, and such that $1 - \overline{H}(X_1 | A_1) \leq \tau$ and $1 - \overline{H}(X_2 | A_2) \leq \tau$, then*

$$1 - \overline{H}(X_1 + X_2 | A_1, A_2) \leq \gamma(1 - \overline{H}(X_1 | A_1)).$$

Analogous inequality for low entropic variables corresponds to *Suction at the lower end* (for $X_t < \tau$).

Lemma 5.2. *For every finite field \mathbb{F}_q and every $\gamma > 0$, there exist τ such that the following holds. Let (X_1, A_1) and (X_2, A_2) be any pair of independent random variables with $X_i \in \mathbb{F}_q$, and such that A_1, A_2 are identically distributed, and moreover for every a we have $\overline{H}(X_1 | A_1 = a) = \overline{H}(X_2 | A_2 = a)$. Then if $\overline{H}(X_1 | A_1) = \overline{H}(X_2 | A_2) \leq \tau$, we have*

$$H(X_1 | X_1 + X_2, A_1, A_2) \leq \gamma \overline{H}(X_1 | A_1).$$

Finally, the following lemma, corresponding to the *Variance in the middle* was already present in the literature, and we will not reproduce the proof — we state it here for future reference.

Lemma 5.3 ([5, Lemma 4.2]). *For every $\tau > 0$ and prime finite field \mathbb{F}_q , there exist $\varepsilon > 0$ such that if (X_1, Y_1) and (X_2, Y_2) are independent pairs of random variables (but not necessarily identically distributed), with $X_i \in \mathbb{F}_q$ for some prime q . Then*

$$\overline{H}(X_1 | Y_1), \overline{H}(X_2 | Y_2) \in (\tau, 1 - \tau)$$

implies

$$\overline{H}(X_1 + X_2 | Y_1, Y_2) \geq \max\{\overline{H}(X_1 | Y_1), \overline{H}(X_2 | Y_2)\} + \varepsilon.$$

5.3 Local polarization of $k \times k$ mixing matrices

In this section we prove Theorem 1.10, that $k \times k$ mixing matrices locally polarize, essentially by reducing to the entropic inequalities of the 2×2 case from Section 5.2.

The high-level strategy for showing local polarization $k \times k$ mixing matrix M is as follows. For simplicity, let us ignore the conditioning on \mathbf{B} . At the t -th step of polarization, for some fixed index $i \in [q]^t$, consider the random vector U with iid coordinates, $\mathbf{U} := (\mathbf{A}_i^{(1)}, \dots, \mathbf{A}_i^{(k)})$, where $\mathbf{A}^{(j)} \sim D_{t-1}$ as in Definition 4.1. And let the linearly-transformed variables in the next step be $\mathbf{V} := \mathbf{U} \cdot M$. In Section 5.3.1 we will show that:

1. There is some index $j \in [k]$ and some $\alpha \in \mathbb{F}_q^*$ for which

$$\overline{H}(\mathbf{V}_j | \mathbf{V}_{<j}) \geq \overline{H}(\mathbf{U}_1 + \alpha \mathbf{U}_2)$$

2. There is some index $j \in [k]$ and some $\alpha \in \mathbb{F}_q^*$ for which

$$\overline{H}(\mathbf{V}_j | \mathbf{V}_{<j}) \leq \overline{H}(\mathbf{U}_1 | \mathbf{U}_1 + \alpha \mathbf{U}_2)$$

Those two, together with entropic inequalities stated in Section 5.2, are enough to show local polarization of a given matrix: we can use condition 1, together with a lower bound $\overline{H}(\mathbf{U}_1 + \alpha \mathbf{U}_2) \geq \overline{H}(\mathbf{U}_1) + \varepsilon$ given by Lemma 5.3 to deduce *Variance in the middle*; condition 1 together with a lower bound from Lemma 5.1 to deduce *Suction at the upper end* (with $\gamma = 1/c$); and condition 2 together with an upper bound from Lemma 5.2 to deduce *Suction at the lower end* (with $\gamma = 1/c$). This is made formal in the proof of Theorem 1.10, which is proved in Section 5.3.2.

5.3.1 Reduction to the 2×2 case

This section will be devoted to proving following two lemmas, which are proven essentially by applying Gaussian Elimination.

Lemma 5.4 (Reduction for Suction at the upper end). *Let (\mathbf{U}, W) be a joint distribution where $\mathbf{U} \in \mathbb{F}_q^k$ (with U_i for $i \in [k]$ being independent) and let M be any mixing matrix. Then, there exist three indices $j, \ell, s \in [k]$, and $\alpha \in \mathbb{F}_q^*$, such that*

$$\overline{H}((\mathbf{U}M)_j \mid (\mathbf{U}M)_{<j}, W) \geq \overline{H}(\mathbf{U}_\ell + \alpha \mathbf{U}_s \mid W).$$

Lemma 5.5 (Reduction for Suction at the lower end). *Let (\mathbf{U}, W) be a joint distribution, where $\mathbf{U} \in \mathbb{F}_q^k$, and let M be any mixing matrix. Then, there exist three indices $j, \ell, s \in [k]$, and $\alpha \in \mathbb{F}_q^*$, such that*

$$\overline{H}((\mathbf{U}M)_j \mid (\mathbf{U}M)_{<j}, W) \leq \overline{H}(\mathbf{U}_\ell \mid \mathbf{U}_\ell + \alpha \mathbf{U}_s, W).$$

Note that for both statements Lemma 5.4 and Lemma 5.5, are invariant under row permutations of M : i.e. if they are true for M' , which is a row permutation of M , the same statement is true for M itself, with different choice of indices ℓ, s .

First, consider performing column-wise Gaussian Elimination on M , and mark the k pivot elements. Let σ be an appropriate row-permutation of M , which brings all the pivot elements to the diagonal. Let $M' := \sigma(M)$ be this row-permuted matrix. From now on we will focus on this matrix instead.

The main idea is that $\overline{H}((\mathbf{U}M)_j \mid (\mathbf{U}M)_{<j}) = \overline{H}((\mathbf{U}M)_j + f((\mathbf{U}M)_{<j}) \mid P((\mathbf{U}M)_{<j}))$ for any f and any full-rank linear transform P (Equation (4)). Thus, we can equivalently consider entropies after “forward-eliminating” variables. The following definition will be useful.

Definition 5.6. *For $j \in [k]$, we define the matrix $M^{(j)}$ as follows. $M^{(j)}$ is the result of applying the following operations to M' :*

1. *Perform column-wise Gaussian Elimination on the first $j - 1$ columns. That is, perform both the “forward” and “backward” pass of Gaussian Elimination.*
2. *Forward-eliminate the j -th column, using the previous $j - 1$ columns.*

Notice the matrices $M^{(j)}$ have the following properties:

1. We can equivalently consider entropies after forward-elimination. That is, for any arbitrary random variable Z , we have the following “forward-elimination identity”:

$$H((\mathbf{U} \cdot M')_j \mid (\mathbf{U} \cdot M')_{<j}, Z) = H((\mathbf{U} \cdot M^{(j)})_j \mid (\mathbf{U} \cdot M^{(j)})_{<j}, Z). \quad (10)$$

2. By definition of M' and Gaussian Elimination, $M^{(s)}$ has all ones on the diagonal, and the top left $(s - 1) \times (s - 1)$ sub-matrix will be the identity matrix. That is, $[M^{(s)}]_{i,j} = \delta_{i,j}$ for $i, j \leq s - 1$. Further, by the forward elimination, $[M^{(s)}]_{i,s} = 0$ for all $i < s$.
3. If M (and hence M') is a mixing matrix, then for all $j \in [k]$, $M^{(j)}$ is not upper-triangular.

For example, for $j = 3$, matrices $M^{(j)}$ will have the following form:

$$M^{(3)} = \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ \star & \star & 1 & \dots \\ \star & \star & \star & \dots \\ \star & \star & \star & \dots \end{bmatrix},$$

where \star denotes any arbitrary element in \mathbb{F}_q .

We are ready now to show the Proof of Lemma 5.4.

Proof of Lemma 5.4. Take $\mathbf{V} := \mathbf{U}M'$ (recall that proving the result for M' is enough to prove our result for M) and consider the first index $j \in [k]$ such that the j -th column of $M^{(j)}$ has support larger than 1. There must exist such an index j , because $M^{(j)}$ is not upper-triangular for at least one $j \in [k]$. Let $s > j$ be the index of the non-diagonal element, such that $\alpha := [M^{(j)}]_{s,j} \neq 0$. Then, we have from (10):

$$\begin{aligned} \overline{H}(\mathbf{V}_j | \mathbf{V}_{<j}, W) &= \overline{H}((\mathbf{U} \cdot M^{(j)})_j | (\mathbf{U} \cdot M^{(j)})_{<j}, W) \\ &= \overline{H}\left(\sum_{i \in [k]} (M^{(j)})_{i,j} \mathbf{U}_i | \mathbf{U}_{<j}, W\right) \\ &\geq \overline{H}(\mathbf{U}_j + (M^{(j)})_{s,j} \mathbf{U}_s | \mathbf{U}_{<j}, W) && \text{(by (3))} \\ &= \overline{H}(\mathbf{U}_j + \alpha \mathbf{U}_s | W) \end{aligned}$$

The final equality uses the fact that $\{\mathbf{U}_i\}_{i \in [k]}$ are independent, and $s > j$. ■

Proof of Lemma 5.5. Note that here, and throughout, all vectors are row-vectors by default. Let $\mathbf{V} := \mathbf{U} \cdot M'$ (recall that proving the result for M' is enough to prove our result for M). Let $j \in [k]$ be the last index where the span of the first $(j-1)$ columns of M' does not equal $\text{span}\{\mathbf{e}_1^T, \dots, \mathbf{e}_{j-1}^T\}$, where $\{\mathbf{e}_i\}$ are the standard basis vectors.

Such an index must exist, because otherwise the matrix M' is upper-triangular (Recall that M' has been row-permuted to place the pivots on the diagonal; thus if for every j , the span of the first $(j-1)$ columns of M' is exactly $\text{span}\{\mathbf{e}_1^T, \dots, \mathbf{e}_{j-1}^T\}$, then M' is upper-triangular). Further, since M (and hence M') is invertible, we also have $j \leq k$.

Now by definition of j , the span of the first j columns of M' must exactly equal $\text{span}\{\mathbf{e}_1^T, \dots, \mathbf{e}_j^T\}$. Thus, all of the first $(j-1)$ columns of $M^{(j)}$ can only be supported on coordinates $\{1, \dots, j\}$. Further, by Definition 5.6, the j -th column of $M^{(j)}$ must be exactly \mathbf{e}_j^T . Finally, because the span of the first $(j-1)$ columns is not $\text{span}\{\mathbf{e}_1^T, \dots, \mathbf{e}_{j-1}^T\}$, there must exist some column $\ell < j$ of $M^{(j)}$ that is supported on coordinate j . In fact, the ℓ -th column of $M^{(j)}$ must be exactly $(\mathbf{e}_\ell^T + \alpha \mathbf{e}_j^T)$ for some $\alpha \in \mathbb{F}_q^*$, due to the Gaussian Elimination. (Recall that we have made sure that the top left $(j-1) \times (j-1)$ submatrix is the identity matrix.)

For example, if $j = 3$, then $M^{(3)}$ must have the form:

$$M^{(3)} = \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ \alpha & \star & 1 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \end{bmatrix} \quad \text{OR} \quad \begin{bmatrix} 1 & 0 & 0 & \dots \\ 0 & 1 & 0 & \dots \\ \star & \alpha & 1 & \dots \\ 0 & 0 & 0 & \dots \\ 0 & 0 & 0 & \dots \end{bmatrix}$$

Now we can see that this j, ℓ, α satisfies the statement of the Lemma:

$$\begin{aligned} H(\mathbf{V}_j | \mathbf{V}_{<j}, W) &= H((\mathbf{U} \cdot M^{(j)})_j | (\mathbf{U} \cdot M^{(j)})_{<j}, W) && \text{(From (10))} \\ &= H(\langle \mathbf{U}, \mathbf{e}_j^T \rangle | (\mathbf{U} \cdot M^{(j)})_{<j}, W) \\ &\leq H(\langle \mathbf{U}, \mathbf{e}_j^T \rangle | (\mathbf{U} \cdot M^{(j)})_\ell, W) \\ &\hspace{10em} \text{(Conditioning does not increase entropy, and } \ell < j) \\ &= H(\mathbf{U}_j | \mathbf{U}_\ell + \alpha \mathbf{U}_j, W). \quad \blacksquare \end{aligned}$$

This concludes our analysis of the reductions. Combined with the entropic inequalities of Section 5.2, this is sufficient to show local polarization of $k \times k$ mixing matrices.

5.3.2 Proof of Theorem 1.10

We begin with a lemma that will be useful in the proof of Theorem 1.10:

Lemma 5.7. *Let $\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(k)}$, and \mathbf{A}' be defined as in Definition 4.1, and let V, W be arbitrary random variables. Then for any multiindex $\mathbf{i} \in [k]^t$ and any $i_{t+1} \in [k]$ we have*

$$\overline{H}(V \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, W) = \overline{H}(V \mid \mathbf{A}'_{[\mathbf{i}, < i_{t+1}]}, \mathbf{A}'_{\prec \mathbf{i}}^{(1)}, \mathbf{A}'_{\prec \mathbf{i}}^{(2)}, \dots, \mathbf{A}'_{\prec \mathbf{i}}^{(k)}, W) .$$

Proof. Observe first that by definition of the order \prec we have that $\mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]} = (\mathbf{A}'_{\prec[\mathbf{i}, 1]}, \mathbf{A}'_{[\mathbf{i}, < i_{t+1}]})$, hence

$$\overline{H}(V \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, W) = \overline{H}(V \mid \mathbf{A}'_{[\mathbf{i}, < i_{t+1}]}, \mathbf{A}'_{\prec[\mathbf{i}, 1]}, W) .$$

The definition of the sequence \mathbf{A}' in terms of \mathbf{A} (in Definition 4.1) reads

$$\mathbf{A}'_{[j, \cdot]} = (\mathbf{A}_j^{(1)}, \dots, \mathbf{A}_j^{(k)})M .$$

Note that if random variables B, B' are related by invertible function $B = f(B')$, then $\overline{H}(A|B) = \overline{H}(A|B')$. By definition of mixing matrix, M is invertible, and hence variables $\mathbf{A}'_{\prec[\mathbf{i}, 1]}$ and variables $\mathbf{A}'_{\prec \mathbf{i}}^{(1)}, \dots, \mathbf{A}'_{\prec \mathbf{i}}^{(k)}$ are indeed related by invertible (linear) transformation, which yields

$$\overline{H}(V \mid \mathbf{A}'_{[\mathbf{i}, < i_{t+1}]}, \mathbf{A}'_{\prec[\mathbf{i}, 1]}, W) = \overline{H}(V \mid \mathbf{A}'_{[\mathbf{i}, < i_{t+1}]}, \mathbf{A}'_{\prec \mathbf{i}}^{(1)}, \mathbf{A}'_{\prec \mathbf{i}}^{(2)}, \dots, \mathbf{A}'_{\prec \mathbf{i}}^{(k)}, W) . \quad \blacksquare$$

Proof of Theorem 1.10. Consider a mixing matrix $M \in \mathbb{F}_q^{k \times k}$, let us condition on a choice of sequence of indices i_1, \dots, i_t , and let us pick a sequence $(\mathbf{A}^{(1)}, \mathbf{B}^{(1)}), \dots, (\mathbf{A}^{(k)}, \mathbf{B}^{(k)}) \sim D_t$ as in Definition 4.1. For future convenience we will use $\mathbf{i} := (i_1, \dots, i_t)$, and $h := \overline{H}(\mathbf{A}_i^{(1)} \mid \mathbf{A}'_{\prec \mathbf{i}}, \mathbf{B}^{(1)})$. For any other $s \in [k]$ we also have $\overline{H}(\mathbf{A}_i^{(s)} \mid \mathbf{A}'_{\prec \mathbf{i}}, \mathbf{B}^{(s)}) = h$, because all the pairs $(\mathbf{A}^{(s)}, \mathbf{B}^{(s)})$ are iid. Let us also take independently $(\mathbf{A}, \mathbf{B}) \sim D_t$, and $(\mathbf{A}', \mathbf{B}')$ constructed from $(\mathbf{A}^{(1)}, \mathbf{B}^{(1)}), \dots, (\mathbf{A}^{(s)}, \mathbf{B}^{(s)})$ as in Definition 4.1. In particular $(\mathbf{A}', \mathbf{B}') \sim D_{t+1}$. Note that with this notation, we have

$$X_t = \overline{H}(\mathbf{A}_i^{(s)} \mid \mathbf{A}'_{\prec \mathbf{i}}, \mathbf{B}^{(s)}) = h,$$

and for a random choice of i_{t+1} , we have

$$X_{t+1} = \overline{H}(\mathbf{A}'_{[\mathbf{i}, i_{t+1}]} \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, \mathbf{B}').$$

We will first show the *Variance in the middle* condition: if for given i_1, \dots, i_t we have $h \in (\tau, 1 - \tau)$, then $\text{Var}_{i_{t+1} \sim [k]}(\overline{H}(\mathbf{A}'_{[\mathbf{i}, i_{t+1}]} \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, \mathbf{B}') - \overline{H}(\mathbf{A}_i \mid \mathbf{A}'_{\prec \mathbf{i}}, \mathbf{B})) > \theta(\tau)$. Note that by the martingale property, we have $\mathbb{E}_{i_{t+1} \sim [k]}[\overline{H}(\mathbf{A}'_{[\mathbf{i}, i_{t+1}]} \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, \mathbf{B}') - \overline{H}(\mathbf{A}_i \mid \mathbf{A}'_{\prec \mathbf{i}}, \mathbf{B})] = 0$, and as such to obtain the lower bound on the variance it is enough to show that

$$\Pr_{i_{t+1} \sim [k]}(\overline{H}(\mathbf{A}'_{[\mathbf{i}, i_{t+1}]} \mid \mathbf{A}'_{\prec[\mathbf{i}, i_{t+1}]}, \mathbf{B}') \geq h + \varepsilon(\tau)) \geq \frac{1}{k} . \quad (11)$$

This would allow us to deduce that the variance above is lower bounded by $\varepsilon(\tau)^2/k$. (Note that this lower bound is true for every h and hence the actual variance needed in the statement of the *Variance in the middle* condition is also true.)

We apply Lemma 5.4 with $\mathbf{U} = (\mathbf{A}_i^{(1)}, \dots, \mathbf{A}_i^{(k)})$ (note that each entry is independent as required). Now consider the triple of indices j, ℓ, s and $\alpha \in \mathbb{F}_q^*$ guaranteed by Lemma 5.4. We have $\Pr(i_{t+1} = j) = \frac{1}{k}$, and if this happens, by Lemma 5.7 we have

$$\overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') = \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{[i, \prec i_{t+1}]}, \mathbf{A}'_{\prec i}^{(1)}, \dots, \mathbf{A}'_{\prec i}^{(k)}, \mathbf{B}'),$$

and by applying Lemma 5.4 we get (recall that $(UM)_j = \mathbf{A}'_{[i, j]}$)

$$\begin{aligned} \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{[i, \prec i_{t+1}]}, \mathbf{A}'_{\prec i}^{(1)}, \dots, \mathbf{A}'_{\prec i}^{(k)}, \mathbf{B}') &\geq \overline{H}(\mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(s)} \mid \mathbf{A}_{\prec i}^{(1)}, \dots, \mathbf{A}_{\prec i}^{(k)}, \mathbf{B}') \\ &= \overline{H}(\mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(s)} \mid \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{A}_{\prec i}^{(s)}, \mathbf{B}^{(\ell)}, \mathbf{B}^{(s)}) \end{aligned}$$

where the equality is deduced by dropping conditioning on random variables independent with $\mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(s)}$ (i.e. if the pair (X, Y) is independent from Y' , then $\overline{H}(X|Y) = \overline{H}(X|Y, Y')$).

Now we can apply Lemma 5.3, to deduce that

$$\overline{H}(\mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(s)} \mid \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{A}_{\prec i}^{(s)}, \mathbf{B}^{(\ell)}, \mathbf{B}^{(s)}) > h + \varepsilon(\tau).$$

Indeed assumptions of the Lemma 5.3 are satisfied: $(\mathbf{A}^{(\ell)}, \mathbf{B}^{(\ell)})$ and $(\mathbf{A}^{(s)}, \mathbf{B}^{(s)})$ are independent by construction, and $\overline{H}(\alpha \mathbf{A}_i^{(s)} \mid \mathbf{B}^{(s)}, \mathbf{A}_{\prec i}^{(s)}) = \overline{H}(\mathbf{A}_i^{(s)} \mid \mathbf{B}^{(s)}, \mathbf{A}_{\prec i}^{(s)}) = \overline{H}(\mathbf{A}_i^{(\ell)} \mid \mathbf{B}^{(\ell)}, \mathbf{A}_{\prec i}^{(\ell)}) = h \in (\tau, 1 - \tau)$. (The first equality used the fact that α is non-zero.) This proves inequality (11), and therefore shows *variation in the middle* for Arkan martingale.

Now we move to the *suction at the upper end* condition. That is, we wish to show that for every c if $1 - h < \tau(c)$, then with probability at least $\frac{1}{k}$, over the choice of i_{t+1} we will have $1 - \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') < \frac{1}{c}(1 - h)$.

The first phase of the proof is analogous to the previous one showing *variation in the middle*. Let us take again triple of indices j, ℓ, s and $\alpha \in \mathbb{F}_q^*$ implied by Lemma 5.4, and as above — with probability $\frac{1}{k}$ we have $i_{t+1} = j$, in which case by Lemma 5.7 and Lemma 5.4, and dropping out superfluous conditioning we get

$$\overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') \geq \overline{H}(\mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(s)} \mid \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{A}_{\prec i}^{(s)}, \mathbf{B}^{(\ell)}, \mathbf{B}^{(s)}).$$

Now applying Lemma 5.1 to pairs $(\mathbf{A}_i^{(\ell)}, (\mathbf{B}^{(\ell)}, \mathbf{A}_{\prec j}^{(\ell)}))$ and $(\alpha \mathbf{A}_i^{(s)}, (\mathbf{A}_{\prec j}^{(s)}, \mathbf{B}^{(s)}))$ (with $\gamma = 1/c$ and τ is picked accordingly), we get desired inequality

$$1 - \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') \leq \frac{1}{c} \cdot (1 - h).$$

Hence, we have shown the desired suction at the upper end (with probability at least $\alpha = \frac{1}{k}$).

Finally, we will show that the Arkan martingale for a mixing matrix M satisfies *suction at the lower end*. That is, we want to show that if $h < \tau(c)$ for τ the same $\tau(c)$ as in the Lemma 5.2 (with $\gamma = 1/c$), then with probability at least $\frac{1}{k}$ we have

$$\overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') \leq \frac{h}{c}. \tag{12}$$

Consider triple of indices j, ℓ, s and $\alpha \in \mathbb{F}_q^*$ as in Lemma 5.5. With probability $\frac{1}{k}$ we have $i_{t+1} = j$, in which case we can use Lemma 5.7 and Lemma 5.5 to deduce

$$\begin{aligned} \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{\prec [i, i_{t+1}]}, \mathbf{B}') &= \overline{H}(\mathbf{A}'_{[i, i_{t+1}]} \mid \mathbf{A}'_{[i, \prec i_{t+1}]}, \mathbf{A}_{\prec i}^{(1)}, \dots, \mathbf{A}_{\prec i}^{(k)}, \mathbf{B}') \\ &\leq \overline{H}(\mathbf{A}_i^{(\ell)} \mid \mathbf{A}_i^{(s)} + \alpha \mathbf{A}_i^{(\ell)}, \mathbf{A}_{\prec i}^{(1)}, \dots, \mathbf{A}_{\prec i}^{(k)}, \mathbf{B}') \\ &= \overline{H}(\mathbf{A}_i^{(\ell)} \mid \mathbf{A}_i^{(s)} + \alpha \mathbf{A}_i^{(\ell)}, \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{A}_{\prec i}^{(s)}, \mathbf{B}^{(\ell)}, \mathbf{B}^{(s)}), \end{aligned}$$

where, again, the last identity is justified by dropping conditioning on variables that are independent from the rest of the expression.

Now, pairs $(\alpha \mathbf{A}_i^{(s)}, (\mathbf{A}_{\prec i}^{(s)}, \mathbf{B}^{(s)}))$ and $(\mathbf{A}_i^{(\ell)}, (\mathbf{A}_{\prec i}^{(\ell)}, \mathbf{B}^{(\ell)}))$ satisfy assumptions of Lemma 5.2 (because pairs $(\mathbf{A}^{(s)}, \mathbf{B}^{(s)})$ and $(\mathbf{A}^{(\ell)}, \mathbf{B}^{(\ell)})$ are iid, and moreover we are assuming $\overline{H}(\mathbf{A}_i^{(\ell)} \mid \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{B}^{(\ell)}) < \tau(c)$ and $\gamma = 1/c$), and therefore

$$\overline{H}(\mathbf{A}_i^{(j)} \mid \mathbf{A}_i^{(\ell)} + \alpha \mathbf{A}_i^{(j)}, \mathbf{A}_{\prec i}^{(j)}, \mathbf{A}_{\prec i}^{(\ell)}, \mathbf{B}^{(j)}, \mathbf{B}^{(\ell)}) \leq \frac{1}{c} \overline{H}(\mathbf{A}_i^{(j)} \mid \mathbf{A}_{\prec i}^{(j)}, \mathbf{B}^{(j)}) = \frac{h}{c}.$$

This shows the last property of local polarization, and concludes the proof of the lemma (with $\tau(c)$ chosen small enough to satisfy all the corresponding conditions in Lemmas 5.1, 5.2 and 5.3). \blacksquare

6 Proofs of Entropic Lemmas

Here we prove the entropic lemmas stated in Section 5.2.

6.1 Suction at the upper end

To establish Lemma 5.1, we will first show similar kind of statement for unconditional entropies. To this end, we first show that for random variables taking values in *small* set, having entropy close to maximal is essentially the same as being close to uniform with respect to L_2 distance. The L_2 distance of a probability distribution to uniform is controlled by the sum of squares of non-trivial Fourier coefficients of the distribution, and all the non-trivial Fourier coefficients are significantly reduced after adding two independent variables close to the uniform distribution.

Finally a simple averaging argument is sufficient to lift this result to conditional entropies, establishing Lemma 5.1.

Lemma 6.1. *If $X \in \mathbb{F}_q$ is a random variable with a distribution \mathcal{D}_X , then*

$$d_2(\mathcal{D}_X, U)^2 \frac{1}{2 \log q} \leq 1 - \overline{H}(X) \leq d_2(\mathcal{D}_X, U)^2 \mathcal{O}(q),$$

where U is a uniform distribution over \mathbb{F}_q , and $d_p(\mathcal{D}_1, \mathcal{D}_2) := \left(\sum_{x \in \mathbb{F}_q} (\mathcal{D}_1(x) - \mathcal{D}_2(x))^p \right)^{1/p}$.

Proof. Pinskers inequality [20] yields $d_1(\mathcal{D}_X, U) \leq \sqrt{2 \log q} \cdot \sqrt{1 - \overline{H}(X)}$, and by standard relations between ℓ_p norms, we have $d_2(\mathcal{D}_X, U) \leq d_1(\mathcal{D}_X, U)$, which after rearranging yields the bound $d_2(\mathcal{D}_X, U)^2 \leq (2 \log q)(1 - \overline{H}(X))$, which in turn proves the claimed lower bound.

For the upper bound, given $i \in \mathbb{F}_q$ let us take δ_i such that $\mathcal{D}_X(i) = \frac{1 + \delta_i}{q}$. We have $\sum_{i \in \mathbb{F}_q} \delta_i = 0$, and $d_2(\mathcal{D}_X, U)^2 = \frac{1}{q^2} \sum_i \delta_i^2$. Now

$$1 - \overline{H}(X) = \frac{1}{\log q} \sum_{i \in \mathbb{F}_q} \frac{(1 + \delta_i)}{q} \log(1 + \delta_i).$$

By Taylor expansion we have $\log(1 + \delta_i) = \delta_i + \mathcal{E}(\delta_i)$ with some error term $\mathcal{E}(\delta_i)$ such that $|\mathcal{E}(\delta_i)| \leq 2\delta_i^2$ for $|\delta_i| < 1$. Therefore in the case when all $\delta_i < 1$, we have (for some constant C):

$$\begin{aligned} 1 - \overline{H}(X) &= \frac{1}{q \log q} \sum_{i \in \mathbb{F}_q} (1 + \delta_i)(\delta_i + \mathcal{E}(\delta_i)) \\ &\leq \frac{1}{q \log q} \sum_{i \in \mathbb{F}_q} [\delta_i + \delta_i^2 + \mathcal{O}(\delta_i^2)] \\ &\leq \frac{1}{q \log q} \left[\sum_{i \in \mathbb{F}_q} \delta_i + C \sum_{i \in \mathbb{F}_q} \delta_i^2 \right] \\ &\leq C q d_2(\mathcal{D}_X, U)^2. \end{aligned}$$

If some $\delta_i \geq 1$, then the inequality is satisfied trivially: $d_2(\mathcal{D}_X, U) \geq \frac{1}{q}$, hence $1 - \overline{H}(X) \leq q d_2(\mathcal{D}_X, U)^2$. \blacksquare

Lemma 6.2. *If $X, Y \in \mathbb{F}_q$ are independent random variables, then $1 - \overline{H}(X + Y) \leq \text{poly}(q)(1 - \overline{H}(X))(1 - \overline{H}(Y))$.*

Proof. By Lemma 6.1 it is enough to show that $d_2(\mathcal{D}_{X+Y}, U)^2 \leq \text{poly}(q) d_2(\mathcal{D}_X, U)^2 d_2(\mathcal{D}_Y, U)^2$. For a distribution \mathcal{D}_X , consider a Fourier transform of this distribution given by $\hat{\mathcal{D}}_X(k) = \mathbb{E}_{j \sim \mathcal{D}_X} \omega^{jk}$, where $\omega = \exp(-2\pi i/q)$. As usual, we have $\hat{\mathcal{D}}_{X+Y}(k) = \hat{\mathcal{D}}_X(k) \hat{\mathcal{D}}_Y(k)$.

Moreover, by Parseval's identity we will show that $d_2(\mathcal{D}_X, U)^2 = \frac{1}{q} \sum_{k \neq 0} \hat{\mathcal{D}}_X(k)^2$. Indeed — as in the proof of Lemma 6.1, define $\mathcal{D}_X(i) =: \frac{1 + \delta_i}{q}$. Then by Parseval's identity we have

$$\frac{1}{q} \cdot \sum_{k \in \mathbb{F}_q} \hat{\mathcal{D}}_X(k)^2 = \sum_{i \in \mathbb{F}_q} \frac{(1 + \delta_i)^2}{q^2} = \frac{1}{q} + d_2(\mathcal{D}_X, U)^2,$$

which implies the claimed bound by noting that $\hat{\mathcal{D}}_X(0) = 1$.

This yields

$$\begin{aligned} q \cdot d_2(\mathcal{D}_{X+Y}, U)^2 &= \sum_{k \neq 0} \hat{\mathcal{D}}_X(k)^2 \hat{\mathcal{D}}_Y(k)^2 \\ &\leq \left(\sum_{k \neq 0} \hat{\mathcal{D}}_X(k)^2 \right) \left(\sum_{k \neq 0} \hat{\mathcal{D}}_Y(k)^2 \right) = q^2 d_2(\mathcal{D}_X, U)^2 d_2(\mathcal{D}_Y, U)^2. \end{aligned}$$

Lemma 6.3. *Let $X_1, X_2 \in \mathbb{F}_q$ be a pair of random variables, and let A_1, A_2 be pair of discrete random variables, such that (X_1, A_1) and (X_2, A_2) are independent. Then*

$$1 - \overline{H}(X_1 + X_2 | A_1, A_2) \leq (1 - \overline{H}(X_1 | A_1))(1 - \overline{H}(X_2 | A_2)) \text{poly}(q).$$

Proof. We have

$$\begin{aligned}
& 1 - \overline{H}(X_1 + X_2 | A_1, A_2) \\
&= \sum_{a_1, a_2} \Pr(A_1 = a_1) \Pr(A_2 = a_2) (1 - \overline{H}(X_1 + X_2 | A_1 = a_1, A_2 = a_2)) \\
&\leq \text{poly}(q) \sum_{a_1, a_2} \Pr(A_1 = a_1) \Pr(A_2 = a_2) (1 - \overline{H}(X_1 | A_1 = a_1, A_2 = a_2)) (1 - \overline{H}(X_2 | A_1 = a_1, A_2 = a_2)) \\
&= \text{poly}(q) \sum_{a_1, a_2} \Pr(A_1 = a_1) (1 - \overline{H}(X_1 | A_1 = a_1)) \Pr(A_2 = a_2) (1 - \overline{H}(X_2 | A_2 = a_2)) \\
&= \text{poly}(q) \left(\sum_{a_1} \Pr(A_1 = a_1) (1 - \overline{H}(X_1 | A_1 = a_1)) \right) \left(\sum_{a_2} \Pr(A_2 = a_2) (1 - \overline{H}(X_2 | A_2 = a_2)) \right) \\
&= \text{poly}(q) (1 - \overline{H}(X_1 | A_1)) (1 - \overline{H}(X_2 | A_2)),
\end{aligned}$$

where the inequality follows from Lemma 6.2 and the second equality follows from independence of (X_1, A_1) and (X_2, A_2) . \blacksquare

Proof of Lemma 5.1. Given γ, q , take $\tau < \gamma/P(q)$ where $P(q)$ is the polynomial appearing in the statement of Lemma 6.3. By applying the conclusion of Lemma 6.3, we have

$$\begin{aligned}
1 - \overline{H}(X_1 + X_2 | A_1, A_2) &\leq (1 - \overline{H}(X_1 | A_1)) (1 - \overline{H}(X_2 | A_2)) P(q) \\
&\leq (1 - \overline{H}(X_1 | A_1)) \tau P(q) \\
&\leq \gamma (1 - \overline{H}(X_1 | A_1)).
\end{aligned}$$

\blacksquare

6.2 Suction at the lower end

In this subsection will show Lemma 5.2. To this end, we want to show that for pairs (X_1, A_1) and (X_2, A_2) with low conditional entropy $\overline{H}(X_1 | A_1) < \tau, \overline{H}(X_2 | A_2) < \tau$, the entropy of the sum is almost as big as sum of corresponding entropies, i.e. $\overline{H}(X_1 + X_2 | A_1, A_2) \geq (1 - \gamma)(\overline{H}(X_1 | A_1) + \overline{H}(X_2 | A_2))$ — and the statement of Lemma 5.2 will follow by application of chain rule. To this end, we first show the same type of statement for non-conditional entropies, i.e. if $\overline{H}(X_1) < \tau, \overline{H}(X_2) < \tau$, then $\overline{H}(X_1 + X_2) > (1 - \gamma)(\overline{H}(X_1) + \overline{H}(X_2))$ — this fact can be deduced by reduction to the analogous fact for binary random variables, where it becomes just a simple computation. Then we proceed by lifting this statement to the corresponding statement about conditional entropies — this requires somewhat more effort than in Lemma 5.1.

Lemma 6.4. *Let X, Y be independent random variables in \mathbb{F}_q . For any $\gamma < 1$, there exists $\alpha = \alpha(\gamma)$ such that: if $\overline{H}(X) \leq \alpha$ and $\overline{H}(Y) \leq \alpha$, then*

$$\overline{H}(X + Y) \geq (1 - \gamma)(\overline{H}(X) + \overline{H}(Y)).$$

First, we will show some preliminary useful lemmas.

Assumption 6.5. *In the following, without loss of generality, let 0 be the most likely value for both random variables X, Y . (This shifting does not affect entropies).*

Lemma 6.6. *Let X be a random variable over \mathbb{F}_q , such that 0 is the most-likely value of X . Then for any q , there exists a function $\alpha_2(\gamma) := \exp(-1/\gamma)$ such that for any $\gamma < 1$ we have*

$$\overline{H}(X) \leq \alpha_2(\gamma) \implies \Pr[X \neq 0] \leq \gamma \overline{H}(X).$$

Proof. Let $\beta := \Pr[X \neq 0]$, and $\alpha := \overline{H}(X)$. We have

$$\alpha \log q = H(X) \geq H(\overline{\delta}(X)) = H(\beta) \geq \beta \log(1/\beta).$$

In the above the inequality follows from the fact that applying a deterministic function to a random variable can only decrease its entropy. Thus,

$$\begin{aligned} \Pr[X \neq 0] = \beta &\leq \frac{\alpha \log q}{\log(1/\beta)} \\ &\leq \frac{\alpha \log q}{\log(1/\alpha) - \log \log q} \end{aligned}$$

where we used the fact that $\beta \leq \alpha \log q$ from Lemma 2.1.

Hence, as soon as $\log \frac{1}{\alpha} > \frac{\log q}{\gamma} + \log \log q$, the statement of the lemma holds. \blacksquare

Lemma 6.7 (Suction-at-lower-end in the Binary Case). *Let U, V be independent binary random variables. There exists a function $\alpha_0(\gamma)$ such that, for any $0 < \gamma < 1$,*

$$H(U), H(V) \leq \alpha_0(\gamma) \implies H(U \oplus V) \geq (1 - \gamma)(H(U) + H(V)).^5$$

Proof. Let p_1 and p_2 be the biases of U, V respectively, such that $U \sim \text{Bernoulli}(p_1)$ and $V \sim \text{Bernoulli}(p_2)$. Let $p_1 \circ p_2 = p_1(1 - p_2) + (1 - p_1)p_2$ be the bias of $U \oplus V$, that is $U \oplus V \sim \text{Bernoulli}(p_1 \circ p_2)$.

We first describe some useful bounds on $H(p)$. On the one hand we have $H(p) \geq p \log 1/p$. For $p \leq 1/2$ we also have $-(1 - p) \log(1 - p) \leq (1/\ln 2)(1 - p)(p + p^2) \leq (1/\ln 2)p \leq 2p$. And so we have $H(p) \leq p(2 + \log 1/p)$.

Summarizing, we have $p \log(1/p) \leq H(p) \leq p \log(1/p) + 2p$. Suppose $H(p_1), H(p_2) \leq \tau$. We now consider $H(p_1) + H(p_2) - H(p_1 \circ p_2)$. WLOG assume $p_1 \leq p_2$. We have

$$\begin{aligned} &H(p_1) + H(p_2) - H(p_1 \circ p_2) \\ &\leq p_1(\log(1/p_1) + 2) + p_2(\log(1/p_2) + 2) - (p_1 \circ p_2) \log(1/(p_1 \circ p_2)) \\ &\leq p_1(\log(1/p_1) + 2) + p_2(\log(1/p_2) + 2) - (p_1 + p_2 - 2p_1p_2) \log(1/(2p_2)) \\ &= p_1 \log(2p_2/p_1) + p_2 \log(2p_2/p_2) + 2p_1p_2 \log(1/(2p_2)) + 2(p_1 + p_2) \\ &\leq p_1 \log(p_2/p_1) + 2p_1p_2 \log(1/(p_2)) + 6p_2 \\ &\leq 2p_1H(p_2) + 7p_2 \quad (\text{Using } p_1 \log(p_2/p_1) \leq p_2) \\ &\leq 2p_1H(p_2) + 7H(p_2)/\log(1/p_2) \\ &\leq 9H(p_2)/\log(1/\tau). \end{aligned}$$

In the above, the last inequality follows from the assumption that $\tau \leq 1/8$ (which will be true in our case). Indeed, note that with this assumption $\tau \log(1/\tau) \leq 1$ (which along with the fact that $p_1 \leq \tau$ implies $p_1 \leq 1/\log(1/\tau)$) and $p_2 \leq \tau$ (since we have $p_2 \log(1/p_2) \leq \tau$). Thus, we have

$$H(U), H(V) \leq \tau \implies H(U) + H(V) - H(U \oplus V) \leq 9H(V)/\log(1/\tau)$$

This implies the desired statement, for $\alpha_0(\gamma) := 2^{-9/\gamma}$. \blacksquare

⁵We note that we could have replaced \oplus by just $+$ as those operations are over \mathbb{F}_2 but we chose to keep $+$ for addition over reals.

Let $\bar{\delta} : \mathbb{F}_q \rightarrow \{0, 1\}$ be the complemented Kronecker-delta function, $\bar{\delta}(x) := \mathbb{1}\{x \neq 0\}$. For small enough entropies, the entropy $H(\bar{\delta}(X))$ is comparable to $H(X)$:

Lemma 6.8. *There exists a function $\alpha_1(\gamma)$ such that for any given $0 < \gamma < 1$, and any arbitrary random variable $X \in \mathbb{F}_q$,*

$$\bar{H}(X) \leq \alpha_1(\gamma) \implies \bar{H}(X) \geq \frac{1}{\log q} H(\bar{\delta}(X)) \geq (1 - \gamma) \bar{H}(X).$$

Proof. The first inequality $\bar{H}(X) \log q = H(X) \geq H(\bar{\delta}(X))$ always holds, by properties of entropy. Thus, we will now show the second bound: that for small enough entropies, $\frac{1}{\log q} H(\bar{\delta}(X)) \geq (1 - \gamma) \bar{H}(X)$. This is equivalent with showing that $H(\bar{\delta}(X)) \geq (1 - \gamma) H(X)$. Given γ , let $\alpha_1 := \alpha_2(\gamma)$ be the entropy guaranteed by Lemma 6.6, such that if $\bar{H}(X) \leq \alpha_2(\gamma)$ then $\Pr[\bar{\delta}(X) = 1] = \Pr[X \neq 0] \leq \gamma \bar{H}(X)$. Now, for $H(X) \leq \alpha_1$, we have

$$\begin{aligned} H(X) &= H(X, \bar{\delta}(X)) \\ &= H(\bar{\delta}(X)) + H(X|\bar{\delta}(X)) && \text{(Chain rule)} \\ &= H(\bar{\delta}(X)) + H(X|\bar{\delta}(X) = 1) \Pr[\bar{\delta}(X) = 1] && \text{(because } H(X|\bar{\delta}(X) = 0) = 0) \\ &\leq H(\bar{\delta}(X)) + \log(q) \Pr[\bar{\delta}(X) = 1] && \text{(because } X \in \mathbb{F}_q, \text{ so } H(X) \leq \log(q)) \\ &\leq H(\bar{\delta}(X)) + \log(q) \gamma \bar{H}(X) && \text{(by Lemma 6.6)} \\ &\leq H(\bar{\delta}(X)) + \gamma H(X). \end{aligned}$$

Thus, if $H(X) \leq \alpha_1$, then $(1 - \gamma)H(X) \leq H(\bar{\delta}(X))$ as desired. ■

Now, by combining these, we can reduce suction-at-the-lower-end from \mathbb{F}_q to the binary case.

Proof of Lemma 6.4. Given γ , we will set $\alpha \leq 1/4$, to be determined later. Notice that we have

$$\bar{H}(X + Y) = \frac{1}{\log q} H(X + Y) \geq \frac{1}{\log q} H(\bar{\delta}(X + Y)). \quad (13)$$

We will proceed to show first that

$$H(\bar{\delta}(X + Y)) \geq H(\bar{\delta}(X) \oplus \bar{\delta}(Y)). \quad (14)$$

This inequality is justified by comparing the distributions of $\bar{\delta}(X + Y)$ and $\bar{\delta}(X) \oplus \bar{\delta}(Y)$, both binary random variables, and noticing that

$$\Pr[\bar{\delta}(X + Y) = 0] = \Pr[X + Y = 0] \leq \Pr[\{X = 0, Y = 0\} \cup \{X \neq 0, Y \neq 0\}] = \Pr[\bar{\delta}(X) \oplus \bar{\delta}(Y) = 0].$$

Moreover, let us observe that $\Pr[\bar{\delta}(X + Y) = 0] = \Pr[X + Y = 0] \geq 1/2$. Indeed,

$$\Pr[X + Y \neq 0] \leq H(X + Y) \leq H(X, Y) \leq H(X) + H(Y) \leq 2\alpha \leq 1/2.$$

In the above, the second inequality follows since $X + Y$ is a deterministic function of X, Y and the third inequality follows from the chain rule and the fact that conditioning can only decrease entropy. Therefore, by monotonicity of the binary entropy function $H(p)$ for $1/2 \leq p \leq 1$, and since $\Pr[\bar{\delta}(X + Y) = 0] \leq \Pr[\bar{\delta}(X) \oplus \bar{\delta}(Y) = 0]$ we have

$$H(\bar{\delta}(X + Y)) \geq H(\bar{\delta}(X) \oplus \bar{\delta}(Y)).$$

This justifies Equation (14).

Now we conclude by using the suction-lemma in the binary case, applied to $\bar{\delta}(X) \oplus \bar{\delta}(Y)$.

Let γ' be a small enough constant, such that $(1 - \gamma')^2 \geq (1 - \gamma)$. Let $\alpha_0 := \alpha_0(\gamma')$ be the entropy bound provided by Lemma 6.7, and let $\alpha_1 := \alpha_1(\gamma')$ be the entropy bound provided by Lemma 6.8. Set $\alpha := \min\{\alpha_0, \alpha_1, 1/4\}$.

Then, for $\bar{H}(X), \bar{H}(Y) \leq \alpha$, we have

$$\begin{aligned} \bar{H}(X + Y) \log q &\geq H(\bar{\delta}(X + Y)) && \text{(Equation (13))} \\ &\geq H(\bar{\delta}(X) \oplus \bar{\delta}(Y)) && \text{(Equation (14))} \\ &\geq (1 - \gamma')(H(\bar{\delta}(X)) + H(\bar{\delta}(Y))) && \text{(Lemma 6.7 and } \bar{H}(\bar{\delta}(Z)) \leq \bar{H}(Z) \text{ for r.v. } Z) \\ &\geq (1 - \gamma')^2(\bar{H}(X) + \bar{H}(Y)) \log q. && \text{(Lemma 6.8)} \end{aligned}$$

With our setting of γ' , this concludes the proof. \blacksquare

We will now see how Lemma 6.4 implies its strengthening for conditional entropies.

Lemma 6.9. *Let (X_1, A_1) and (X_2, A_2) be independent random variables with $X_i \in \mathbb{F}_q$, and such that A_1, A_2 are identically distributed, and moreover for every a we have $\bar{H}(X_1|A_1 = a) = \bar{H}(X_2|A_2 = a)$. Then for every $\gamma > 0$, there exist τ such that if $\bar{H}(X_1|A_1) \leq \tau$, done]technically A and X have not been defined. then*

$$\bar{H}(X_1 + X_2|A_1, A_2) \geq (1 - \gamma)(\bar{H}(X_1|A_1) + \bar{H}(X_2|A_2)). \quad (15)$$

Proof. Let us take $\alpha := \bar{H}(X_1|A_1) = \bar{H}(X_2|A_2)$. For given γ we shall find τ such that if $\alpha < \tau$ then inequality (15) is satisfied. Let us now consider $G_A := \{a : \bar{H}(X_1|A_1 = a) < \alpha_1\}$, for $\alpha_1 = \frac{\alpha}{\gamma}$. (In the remainder of the proof when we want to talk about a random variable from the identical distribution from which A_1 and A_2 are drawn, we will denote it by A .) By Markov inequality

$$\Pr(A \notin G_A) \leq \frac{\alpha}{\alpha_1} = \gamma.$$

Let us fix now τ which appears in the statement of this lemma to be smaller than γ and moreover small enough so that when $\alpha < \tau$ for every $a_1, a_2 \in G_A$ we can apply Lemma 6.4 to distributions $(X_1|A_1 = a_1)$ and $(X_2|A_2 = a_2)$ to ensure that $H(X_1 + X_2|A_1 = a_1, A_2 = a_2) \geq (1 - \gamma)(H(X_1|A_1 = a_1) + H(X_2|A_2 = a_2))$.

Let us use shorthand $S(a_1, a_2) = \bar{H}(X_1 + X_2|A_1 = a_1, A_2 = a_2) \Pr(A_1 = a_1, A_2 = a_2)$. We have

$$\begin{aligned} \bar{H}(X_1 + X_2|A_1, A_2) &= \sum_{a_1, a_2} S(a_1, a_2) \\ &\geq \sum_{\substack{a_1 \in G_A \\ a_2 \in G_A}} S(a_1, a_2) + \sum_{\substack{a_1 \notin G_A \\ a_2 \in G_A}} S(a_1, a_2) + \sum_{\substack{a_1 \in G_A \\ a_2 \notin G_A}} S(a_1, a_2). \end{aligned} \quad (16)$$

If both a_1 and a_2 are in G_A , then by Lemma 6.4 we have

$$S(a_1, a_2) \geq (1 - \gamma)(\bar{H}(X_1|A_1 = a_1) + H(X_2|A_2 = a_2)) \Pr(A_1 = a_1, A_2 = a_2),$$

therefore

$$\sum_{a_1 \in G_A, a_2 \in G_A} S(a_1, a_2) \geq 2(1 - \gamma) \Pr(A \in G_A) \sum_{a_1 \in G_A} H(X_1|A_1 = a_1) \Pr(A_1 = a_1), \quad (17)$$

where in the above we have used the fact that A_1 and A_2 are identically distributed.

On the other hand, for $a_1 \notin G_A, a_2 \in G_A$ let us bound

$$\begin{aligned} S(a_1, a_2) &= \overline{H}(X_1 + X_2 | A_1 = a_1, A_2 = a_2) \Pr(A_1 = a_1, A_2 = a_2) \\ &\geq \overline{H}(X_1 + X_2 | A_1 = a_1, A_2 = a_2, X_2) \Pr(A_1 = a_1, A_2 = a_2) \\ &= \overline{H}(X_1 | A_1 = a_1) \Pr(A_1 = a_1, A_2 = a_2) \end{aligned}$$

where the inequality follows from the fact that additional conditioning decreases entropy (and we also used the fact that since X_1 and X_2 are independent, $\overline{H}(X_1 + X_2 | A_1 = a_1, A_2 = a_2, X_2) = \overline{H}(X_1 | A_1 = a_1, A_2 = a_2, X_2) = \overline{H}(X_1 | A_1 = a_1, A_2 = a_2) = \overline{H}(X_1 | A_1 = a_1)$). Summing this bound over all such pairs yields

$$\sum_{a_1 \notin G_A, a_2 \in G_A} S(a_1, a_2) \geq \Pr(A \in G_A) \sum_{a_1 \notin G_A} \overline{H}(X_1 | A_1 = a_1) \Pr(A_1 = a_1) \quad (18)$$

and symmetrically for the third summand, we get

$$\sum_{a_1 \in G_A, a_2 \notin G_A} S(a_1, a_2) \geq \Pr(A \in G_A) \sum_{a_2 \notin G_A} \overline{H}(X_2 | A_2 = a_2) \Pr(A_2 = a_2). \quad (19)$$

Plugging in (17), (18) and (19) into (16) (and the fact that A_1 and A_2 are identically distributed) we find

$$\begin{aligned} \overline{H}(X_1 + X_2 | A_1, A_2) &\geq 2(1 - \gamma) \Pr(A \in G_A) \sum_{a_1} \overline{H}(X_1 | A_1 = a_1) \Pr(A_1 = a_1) \\ &= 2(1 - \gamma) \Pr(A \in G_A) \overline{H}(X_1 | A_1). \end{aligned}$$

We have $\Pr(A \in G_A) \geq (1 - \gamma)$, which yields

$$\overline{H}(X_1 + X_2 | A_1, A_2) \geq 2(1 - \gamma)^2 \alpha \geq 2(1 - 2\gamma)\alpha$$

and the statement of the lemma follows, after rescaling γ by half. ■

Proof of Lemma 5.2. By chain rule we have

$$\begin{aligned} \overline{H}(X_1 | X_1 + X_2, A_1, A_2) &= \overline{H}(X_1, X_1 + X_2 | A_1, A_2) - \overline{H}(X_1 + X_2 | A_1, A_2) \\ &= \overline{H}(X_1, X_2 | A_1, A_2) - \overline{H}(X_1 + X_2 | A_1, A_2) \\ &= 2\overline{H}(X_1 | A_1) - \overline{H}(X_1 + X_2 | A_1, A_2), \end{aligned}$$

where the last equality follows from the independence of (X_1, A_1) and (X_2, A_2) . Now we can apply Lemma 6.9 to get

$$\overline{H}(X_1 | X_1 + X_2, A_1, A_2) \leq 2\overline{H}(X_1 | A_1) - (1 - \gamma)(2\overline{H}(X_1 | A_1)) = 2\gamma\overline{H}(X_1 | A_1)$$

and the statement follows directly from Lemma 6.9 and rescaling γ by half. ■

References

- [1] Emmanuel Abbe and Emre Telatar. Polar codes for the m -user multiple access channel. *IEEE Transactions on Information Theory*, 58(8):5437–5448, 2012.

- [2] Erdal Arıkan. Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels. *IEEE Transactions on Information Theory*, pages 3051–3073, July 2009.
- [3] Erdal Arıkan and Emre Telatar. On the rate of channel polarization. In *Proceedings of 2009 IEEE International Symposium on Information Theory*, pages 1493–1495, 2009.
- [4] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. John Wiley and Sons, Hoboken, NJ, USA, 2nd edition, 2005.
- [5] Eren Şaşıođlu. Polarization and polar codes. *Foundations and Trends in Communications and Information Theory*, 8(4):259–381, 2012.
- [6] Rick Durrett. *Probability: Theory and examples*, 2011.
- [7] Furkan Ercan, Carlo Condo, Seyyed Ali Hashemi, and Warren J. Gross. On error-correction performance and implementation of polar code list decoders for 5G. In *55th Annual Allerton Conference on Communication, Control, and Computing*, 2017.
- [8] Arman Fazeli, S. Hamed Hassani, Marco Mondelli, and Alexander Vardy. Binary linear codes with optimal scaling and quasi-linear complexity. *Personal communication*, October 2017.
- [9] Naveen Goela, Emmanuel Abbe, and Michael Gastpar. Polar codes for broadcast channels. In *Proceedings of the 2013 IEEE International Symposium on Information Theory, Istanbul, Turkey, July 7-12, 2013*, pages 1127–1131, 2013.
- [10] Venkatesan Guruswami and Ameya Velingker. An entropy sunset inequality and polynomially fast convergence to Shannon capacity over all alphabets. In *Proceedings of 30th Conference on Computational Complexity*, pages 42–57, 2015.
- [11] Venkatesan Guruswami and Patrick Xia. Polar codes: Speed of polarization and polynomial gap to capacity. *IEEE Trans. Information Theory*, 61(1):3–16, 2015. Preliminary version in Proc. of FOCS 2013.
- [12] Seyed Hamed Hassani, Kasra Alishahi, and Rüdiger L. Urbanke. Finite-length scaling for polar codes. *IEEE Trans. Information Theory*, 60(10):5875–5898, 2014.
- [13] S.H. Hassani, K. Alishahi, and R. Urbanke. Finite-length scaling for polar codes. *Information Theory, IEEE Transactions on*, PP(99):1–1, 2014.
- [14] Satish Babu Korada. Polar codes for Slepian-Wolf, Wyner-Ziv, and Gelfand-Pinsker. In *Proceedings of the 2010 IEEE Information Theory Workshop*, pages 1–5, 2010.
- [15] Satish Babu Korada, Eren Sasoglu, and Rüdiger L. Urbanke. Polar codes: Characterization of exponent, bounds, and constructions. *IEEE Transactions on Information Theory*, 56(12):6253–6264, 2010.
- [16] Hessam MahdaviFar and Alexander Vardy. Achieving the secrecy capacity of wiretap channels using polar codes. *IEEE Transactions on Information Theory*, 57(10):6428–6443, 2011.
- [17] Marco Mondelli, S. Hamed Hassani, and Rüdiger L. Urbanke. Unified scaling of polar codes: Error exponent, scaling exponent, moderate deviations, and error floors. *IEEE Trans. Information Theory*, 62(12):6698–6712, 2016.

- [18] Ryuhei Mori and Toshiyuki Tanaka. Source and channel polarization over finite fields and reed-solomon matrices. *IEEE Trans. Information Theory*, 60(5):2720–2736, 2014.
- [19] Henry D. Pfister and Rüdiger L. Urbanke. Near-optimal finite-length scaling for polar codes over large alphabets. In *IEEE International Symposium on Information Theory, ISIT*, pages 215–219, 2016.
- [20] M.S. Pinsker. *Information and Information Stability of Random Variables and Processes*. Holden-Day series in time series analysis. Holden-Day, 1964.
- [21] Eren Sasoglu. Polarization and polar codes. *Foundations and Trends in Communications and Information Theory*, 8(4):259–381, 2012.
- [22] Eren Sasoglu, Emre Telatar, and Edmund M. Yeh. Polar codes for the two-user multiple-access channel. *IEEE Transactions on Information Theory*, 59(10):6583–6592, 2013.
- [23] Ido Tal and Alexander Vardy. How to construct polar codes. *IEEE Transactions on Information Theory*, 59(10):6562–6582, Oct 2013.
- [24] Lele Wang and Eren Sasoglu. Polar coding for interference networks. In *2014 IEEE International Symposium on Information Theory, Honolulu, HI, USA, June 29 - July 4, 2014*, pages 311–315, 2014.

A Codes from Polarization

In this section, we describe the construction of polar codes, and analyze the failure probability of decoders by corresponding them to the Arikan martingale. This proves Theorem 1.9.

Specifically, we first describe the polar encoder along with a fast $\mathcal{O}(n \log n)$ -time implementation, where n is the blocklength. Then, in Section A.2 we define the (inefficient) successive-cancellation decoder, and analyze its failure probability assuming a correspondence between polar coding and the Arikan martingale. In Section A.2.2, we describe a fast $\mathcal{O}(n \log n)$ -time decoder that is functionally equivalent to the successive-cancellation decoder. Finally, in Section A.2.3, we prove the required correspondence between polar coding and the Arikan martingale.

Throughout this section, fix parameters $k \in \mathbb{N}$ as the dimension of the mixing matrix $M \in \mathbb{F}_q^{k \times k}$, \mathbb{F}_q as a finite field, and $n = k^t$ as the codeword length.

A.1 Polar Encoder

Given a set $S \subseteq [n]$ and a fixing $\alpha \in \mathbb{F}_q^{|S^c|}$,⁶ we define the polar code of dimension $|S|$ by giving the encoder mapping $\mathbb{F}_q^S \rightarrow \mathbb{F}_q^n$ as follows:

Algorithm 1 Polar Encoder

Constants: $M \in \mathbb{F}_q^{k \times k}$, $S \subseteq [n]$, $\alpha \in \mathbb{F}_q^{S^c}$

Input: $U \in \mathbb{F}_q^S$

Output: $Z \in \mathbb{F}_q^n$

- 1: **procedure** POLAR-ENCODER($U; \alpha$)
 - 2: Extend U to $\bar{U} \in \mathbb{F}_q^n$ by letting $(\bar{U}_i)_{i \notin S} = \alpha$ for coordinates not in S
 - 3: **Return** $Z = \bar{U} \cdot (M^{-1})^{\otimes t}$
-

The above gives a polynomial time algorithm for encoding. An $\mathcal{O}_q(n \log n)$ algorithm can also be obtained by using the recursive structure imposed by the tensor powers.

Below, we switch to considering vectors in $\mathbb{F}_q^{k^t}$ as tensors in $(\mathbb{F}_q^k)^{\otimes t}$, indexed by multiindices $i \in [k]^t$. The following encoder takes as input the “extended” message \bar{U} , as defined above.

Algorithm 2 Fast Polar Encoder

Constants: $M \in \mathbb{F}_q^{k \times k}$

Input: $\bar{U} \in (\mathbb{F}_q^k)^{\otimes t}$

Output: $Z = \bar{U} \cdot (M^{-1})^{\otimes t}$

- 1: **procedure** FAST-POLAR-ENCODER _{t} (\bar{U})
 - 2: **for all** $j \in [k]$ **do**
 - 3: $Z^{(j)} \leftarrow$ FAST-POLAR-ENCODER _{$t-1$} ($\bar{U}_{[\cdot, j]}$)
 - 4: **for all** $i \in [k]^{t-1}$ **do**
 - 5: $Z_{[i, \cdot]} \leftarrow (Z_i^{(1)}, Z_i^{(2)}, \dots, Z_i^{(k)}) \cdot M^{-1}$
 - 6: **Return** Z
-

⁶We use the notation $S^c = [n] \setminus S$.

A.2 The Successive-Cancellation Decoder

Here we describe a successive-cancellation decoder. Note that this decoder is not efficient, but the fast decoder described later will have the exact same error probability as this decoder.

For given channel outputs \mathbf{Y} , let \mathbf{Z} be the posterior distribution on channel inputs given outputs \mathbf{Y} . Each $\mathbf{Z}_i \in \Delta(\mathbb{F}_q)$ is the conditional distribution $\mathbf{Z}_i | \mathbf{Y}_i$ defined by the channel $\mathcal{C}_{Y|Z}$ and the received output \mathbf{Y}_i .

Now, for a set $S \subseteq [n]$ and a fixing $\boldsymbol{\alpha} \in \mathbb{F}_q^{S^c}$, we define the decoder on input \mathbf{Z} as follows.

Algorithm 3 Successive-Cancellation Decoder

Constants: $M \in \mathbb{F}_q^{k \times k}$, $S \subseteq [n]$, $\boldsymbol{\alpha} \in \mathbb{F}_q^{S^c}$

Input: $\mathbf{Z} \in \Delta(\mathbb{F}_q^k)^n$

Output: $\hat{\mathbf{U}} \in \mathbb{F}_q^n$

```

1: procedure SC-DECODER( $\mathbf{Z}$ )
2:   Compute joint distribution  $\mathbf{U} \in \Delta(\mathbb{F}_q^n) : \mathbf{U} \leftarrow \mathbf{Z}M^{\otimes t}$ 
3:   for all  $i \in [n]$  do
4:     If  $i \in S$  then
5:        $\hat{\mathbf{U}}_i \leftarrow \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr_{\mathbf{U}}(\mathbf{U}_i = x)$ 
6:     else
7:        $\hat{\mathbf{U}}_i \leftarrow \boldsymbol{\alpha}_i$ 
8:     Update distribution  $\mathbf{U} \leftarrow (\mathbf{U} | \mathbf{U}_i = \hat{\mathbf{U}}_i)$ 
9:   Return  $\hat{\mathbf{U}}$ 

```

Note that several of the above steps, including computing the joint distribution of \mathbf{U} and marginal distributions of \mathbf{U}_i , are not computationally efficient.

A.2.1 Decoding Analysis

We will first reason about the “genie-aided” case, when the fixing $\boldsymbol{\alpha} \in \mathbb{F}_q^{|S^c|}$ of non-message bits is chosen uniformly at random, and revealed to both the encoder and decoder. Then, we will argue that it is sufficient to use a deterministic fixing $\boldsymbol{\alpha} = \boldsymbol{\alpha}_0$.

We now argue that over a uniform choice of message \mathbf{U}_S , and a uniform fixing $\boldsymbol{\alpha}$ of non-message bits, the probability of decoding failure is bounded as follows.

Claim A.1. *For $V \sim \mathbb{F}_q^S$, and $\boldsymbol{\alpha} \sim \mathbb{F}_q^{S^c}$, take $Z := \text{POLAR-ENCODER}(V; \boldsymbol{\alpha})$ and Y sampled according to the channel $Y := \mathcal{C}(Z)$. Let $\mathbf{U} \in \mathbb{F}_q^n$ be specified by V on coordinates $i \in S$, and specified by $\boldsymbol{\alpha}$ on coordinates $i \notin S$. With this notation, we have*

$$\Pr[\text{SC-DECODER}(\mathbf{Y}; \boldsymbol{\alpha}) \neq \mathbf{U}] \leq \sum_{i \in S} H(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}).$$

Proof. Note that \mathbf{U} is uniform over \mathbb{F}_q^n . Now, we have:

$$\begin{aligned} \Pr[\text{SC-DECODER}(\mathbf{Y}; \boldsymbol{\alpha}) \neq \mathbf{U}] &= \Pr[\exists i \hat{\mathbf{U}}_i \neq \mathbf{U}_i] \\ &= \sum_{i \leq n} \Pr[\hat{\mathbf{U}}_i \neq \mathbf{U}_i | \hat{\mathbf{U}}_{<i} = \mathbf{U}_{<i}]. \end{aligned}$$

Clearly for $i \notin S$ we have $\Pr[\hat{U}_i \neq U_i] = 0$, since both are defined to be equal to α_i on those coordinates. It is enough to show that for $i \in S$ we have

$$\Pr(\hat{U}_i \neq U_i \mid \mathbf{U}_{<i} = \hat{\mathbf{U}}_{<i}) \leq H(\mathbf{U}_i \mid \mathbf{U}_{<i}, \mathbf{Y}).$$

This follows directly from Lemma 2.2, as \hat{U}_i is defined exactly as a maximum likelihood estimator of U_i given channel outputs \mathbf{Y} and conditioning on $\mathbf{U}_{<i}$. \blacksquare

Claim A.2. *If X_t satisfies (τ, ε) -polarization, then there exist a set $S \subset [n]$ of size $(\text{Capacity}(\mathcal{C}_{Y|Z}) - \varepsilon - 2\tau)n$, such that*

$$\sum_{i \in S} H(\mathbf{U}_i \mid \mathbf{U}_{<i}, \mathbf{Y}) \leq \tau n \log q.$$

Proof. First, observe that for uniform choice of $i \in [n]$, $H(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y})$ is distributed identically as X_t in the Arikan Martingale. Because, by definition of the encoder, we have channel inputs $\mathbf{Z} = \mathbf{U} \cdot (M^{\otimes t})^{-1}$ or equivalently

$$\mathbf{U} = \mathbf{Z} \cdot M^{\otimes t}.$$

And $\mathbf{Y} = \mathcal{C}(\mathbf{Z})$. Thus, by Lemma A.6, $\bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y})$ for a random $i \in [n]$ is distributed identically as X_t .

Now, for symmetric channels, the uniform distribution achieves capacity (See eg, Theorem 7.2.1 in [4]). Thus, for uniform channel input \mathbf{Z} ,

$$\text{Capacity}(\mathcal{C}_{Y|Z}) = \bar{H}(\mathbf{Z}) - \bar{H}(\mathbf{Z} | \mathbf{Y}) = 1 - \bar{H}(\mathbf{Z} | \mathbf{Y}).$$

Let S be the set of all indices i such that $\bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}) < \tau$. By definition, we have

$$\sum_{i \in S} \bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}) \leq \tau n$$

as desired.

Now observe that polarization of martingale X_t directly implies that we have at most εn indices i satisfying $\bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}) \in (\tau, 1 - \tau)$. Let S' be a set of indices for which $\bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}) > 1 - \tau$. We have

$$\begin{aligned} n(1 - \text{Capacity}(\mathcal{C}_{Y|Z})) &= \bar{H}(\mathbf{U}(M^{-1})^{\otimes t} | \mathbf{Y}) \\ &= \bar{H}(\mathbf{U}_1, \dots, \mathbf{U}_n | \mathbf{Y}) && \text{(Since } (M^{-1})^{\otimes t} \text{ is full rank)} \\ &= \sum_{i \in [n]} \bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}) && \text{(Chain rule)} \\ &\geq \sum_{i \in S'} \bar{H}(\mathbf{U}_i | \mathbf{U}_{<i}, \mathbf{Y}) \\ &\geq (1 - \tau)|S'|, \end{aligned}$$

which implies (for $\tau < \frac{1}{2}$) that

$$|S'| \leq n(1 - \text{Capacity}(\mathcal{C}_{Y|Z}) + 2\tau),$$

and finally

$$|S| \geq n - |S'| - \varepsilon n \geq n(\text{Capacity}(\mathcal{C}_{Y|Z}) - \varepsilon - 2\tau).$$

\blacksquare

We can now combine the above to prove a version of Theorem 1.9 for the (inefficient) successive-cancellation decoder:

Theorem A.3. *If a matrix $M \in \mathbb{F}_q^{k \times k}$ strongly polarizes for a symmetric channel $\mathcal{C}_{Y|Z}$ then for every $c < \infty$ there exists $t_0(x) = O(\log x)$ such that for every $\varepsilon > 0$, for every $t \geq t_0(1/\varepsilon)$, the rows of $(M^{-1})^{\otimes t}$ together with an affine shift generate an affine code of codeword length $n = k^t$, and of rate $\text{Capacity}(\mathcal{C}_{Y|Z}) - \varepsilon$. Furthermore, the successive-cancellation decoder succeeds with probability at least $1 - n^{-c}$.*

Proof. Fix some constant c , and take $\gamma < k^{-c-1} \log^{-1} q$. By the definition of strong polarization property, we know that for some constants β, η , martingale X_t is $(\gamma^t, \beta \cdot \eta^t)$ -polarizing, hence by Claim A.2, there exist a set $S \subset [n]$ of size $(\text{Capacity}(\mathcal{C}_{Y|Z}) - \beta \cdot \eta^t - 2\gamma^t)n$, such that

$$\begin{aligned} \sum_{i \in S} H(\mathbf{U}_i \mid \mathbf{U}_{<i}, \mathbf{Y}) &\leq \gamma^t n \log q \\ &\leq n^{-c}. \end{aligned}$$

Thus, by Claim A.1,

$$\begin{aligned} \Pr[\text{SC-DECODER}(\mathbf{Y}; \boldsymbol{\alpha})_S \neq U] &\leq \sum_{i \in S} H(\mathbf{U}_i \mid \mathbf{U}_{<i}, \mathbf{Y}) \\ &\leq n^{-c} \end{aligned}$$

and so the SC-DECODER fails with probability (arbitrary) inverse-polynomial. Note that this failure probability is an average over random choice of fixing $\boldsymbol{\alpha}$, but this implies there is some deterministic fixing $\boldsymbol{\alpha} = \boldsymbol{\alpha}_0$ with failure probability at least as good. Further, by linearity of the encoding, such a deterministic fixing yields an affine code.

Finally, for $t \geq \Omega_{\eta, \beta}(\log(1/\varepsilon))$, we have $\beta\eta^t + 2\gamma^t \leq \varepsilon$, and hence the size $|S| \geq (\text{Capacity}(\mathcal{C}_{Y|Z}) - \varepsilon)n$. Now the rate of the code as defined above is $|S|/n \geq \text{Capacity}(\mathcal{C}_{Y|Z}) - \varepsilon$, as desired. \blacksquare

A.2.2 Fast Decoder

In this section we will define the recursive FAST-DECODER algorithm. The observation that polar codes admit this recursive fast-decoder was made in the original work of Arikan [2].

FAST-DECODER will take on input descriptions of the posterior distributions on channel inputs $\{\mathbf{Z}_i\}_{i \in [k]^s}$ for some s , where each individual $\mathbf{Z}_i \in \Delta(\mathbb{F}_q)$ is a distribution over \mathbb{F}_q , as well as $\boldsymbol{\alpha} \in (\mathbb{F}_q \cup \{\perp\})^{[k]^s}$ where $\boldsymbol{\alpha}_i \neq \perp$ for $i \notin S$, are fixed values corresponding to non-message positions. The output of FAST-DECODER is a vector $\hat{\mathbf{Z}} \in (\mathbb{F}_q^k)^{\otimes s}$ — the guess for the actual channel inputs. To recover the message, it is enough to apply $\hat{\mathbf{U}} := \hat{\mathbf{Z}}M^{\otimes t}$, and restrict it to unknown positions S .

Below, for $\mathbf{W}_i \in \Delta(\mathbb{F}_q^k)$ — a description of joint probability distribution over \mathbb{F}_q^k , we will write $\pi_j(\mathbf{W}_i) \in \Delta(\mathbb{F}_q)$ as a j -th marginal of \mathbf{W}_i , i.e. projection on the j -th coordinate.

Algorithm 4 Fast Decoder

Constants: $M \in \mathbb{F}_q^{k \times k}$

Input: $\mathbf{Z} = \{\mathbf{Z}_i \in \Delta(\mathbb{F}_q)\}_{i \in [k]^s}$, $\boldsymbol{\alpha} \in (\mathbb{F}_q \cup \{\perp\})^{[k]^s}$

Output: $\hat{\mathbf{Z}} \in (\mathbb{F}_q^k)^{\otimes s}$

```

1: procedure FAST-DECODERs( $\mathbf{Z}; \boldsymbol{\alpha}$ )
2:   If  $s = 0$  then
3:     If  $\boldsymbol{\alpha} = \perp$  then
4:       Return  $\hat{\mathbf{Z}} = \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{Z} = x]$ 
5:     else
6:       Return  $\hat{\mathbf{Z}} = \boldsymbol{\alpha}$ 
7:   else
8:     for all  $i \in [k]^{s-1}$  do
9:       Compute joint distribution  $\mathbf{W}_i \in \Delta(\mathbb{F}_q^k)$ , given by  $\mathbf{W}_i \leftarrow \mathbf{Z}_{[\cdot, i]}M$ 
10:      for all  $j \in [k]$  do
11:         $\mathbf{Z}'^{(j)} \leftarrow \{\pi_j(\mathbf{W}_i)\}_{i \in [k]^{s-1}}$ 
12:         $\hat{\mathbf{V}}^{(j)} \leftarrow \text{FAST-DECODER}_{s-1}(\mathbf{Z}'^{(j)}; \boldsymbol{\alpha}_{[j, \cdot]})$ 
13:        for all  $i \in [k]^{s-1}$  do
14:          Update distribution  $\mathbf{W}_i \leftarrow (\mathbf{W}_i | \pi_j(\mathbf{W}_i) = \hat{\mathbf{V}}_i^{(j)})$ 
15:        for all  $i \in [k]^{s-1}$  do
16:           $\hat{\mathbf{Z}}_{[\cdot, i]} \leftarrow \mathbf{W}_i M^{-1}$ 
17:      Return  $\hat{\mathbf{Z}}$ 

```

Note that in Line 16 above, \mathbf{W}_i is technically a distribution, but by this point \mathbf{W}_i is deterministic, since all its coordinates have been set previously via Line 14. We abuse notation by using \mathbf{W}_i in Line 16 to denote a fixed vector in \mathbb{F}_q^k .

The FAST-DECODER as described above runs in time $\mathcal{O}(n \log n)$, where $n = k^t$ is blocklength.

Moreover, it can be seen directly that FAST-DECODER is equivalent to the SC-DECODER on the same input:

Lemma A.4. *For all product distributions \mathbf{Z} on inputs (where each $\mathbf{Z}_i \in \Delta(\mathbb{F}_q)$ is a distribution over \mathbb{F}_q), for all sets $S \subseteq [k]^s$ and all fixings $\boldsymbol{\alpha} \in (\mathbb{F}_q \cup \{\perp\})^{[k]^s}$ of the set S , the following holds:*

$$\text{FAST-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}) \cdot M^{\otimes s} = \text{SC-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}|_S).$$

Proof of Lemma A.4. Given $\mathbf{Z}, S, \boldsymbol{\alpha}$ as in the statement, let \mathbf{U} be the joint distribution defined by $\mathbf{U} := \mathbf{Z}M^{\otimes s}$. We will use this joint distribution on (\mathbf{U}, \mathbf{Z}) throughout the proof.

First, notice the following about the operation of the SC-DECODER: SC-DECODER($\mathbf{Z}; \boldsymbol{\alpha}|_S$) by definition iteratively computes estimates $\hat{\mathbf{U}}^{\text{SC}}$ such that for all $i \in [k]^s$:

$$\hat{\mathbf{U}}_i^{\text{SC}} = \begin{cases} \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}_i = x | \mathbf{U}_{\prec i} = \hat{\mathbf{U}}_{\prec i}^{\text{SC}}] & \text{for } i \notin S \\ \boldsymbol{\alpha}_i & \text{for } i \in S. \end{cases} \quad (20)$$

We now argue that the FAST-DECODER computes the identical estimates. That is,

$$\text{FAST-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}) \cdot M^{\otimes s} = \hat{\mathbf{U}}^{\text{SC}} = \text{SC-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}|_S)$$

We prove this by induction on s . $\mathbf{i} \in [k]^s$ per the lexicographic order \prec . Note that this equality clearly holds for the “frozen” indices $\mathbf{i} \in S$, so we focus on proving the claim for indices $\mathbf{i} \notin S$.

For $s = 0$, the claim is true by definition (both FAST-DECODER and SC-DECODER, on input a distribution $Z \in \Delta(\mathbb{F}_q)$, output $\operatorname{argmax}_x \Pr[Z = x]$).

For $s > 0$: Let $\hat{\mathbf{U}}^{\text{F}} := \text{FAST-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}) \cdot M^{\otimes s}$. We argue by a further induction on indices that $\hat{\mathbf{U}}^{\text{F}} = \hat{\mathbf{U}}^{\text{SC}}$. Suppose for induction that for some $j \in [k]$, we have $\hat{\mathbf{U}}_{[\prec j, \cdot]}^{\text{F}} = \hat{\mathbf{U}}_{[\prec j, \cdot]}^{\text{SC}}$. We will now show that $\hat{\mathbf{U}}_{[j, \cdot]}^{\text{F}} = \hat{\mathbf{U}}_{[j, \cdot]}^{\text{SC}}$.

Let $\mathbf{Z}'^{(j)}$ be the random variable as defined by FAST-DECODER($\mathbf{Z}; \boldsymbol{\alpha}$) in Line 11, at the fixed iteration $j \in [k]$. And for all $\ell \in [k]$, let $\hat{\mathbf{V}}^{(\ell)}$ denote the variable defined in Line 12, when the loop variable j is equal to ℓ .

By definition of the FAST-DECODER, and the structure of the tensor-product, we have

$$\forall \ell \in [k] : \hat{\mathbf{V}}^{(\ell)} \cdot M^{\otimes s-1} = \hat{\mathbf{U}}_{[\ell, \cdot]}^{\text{F}} \quad (21)$$

Indeed, consider the value of \mathbf{W}_i defined in the FAST-DECODER, when it is accessed in Line 16 of the algorithm. Let $\overline{\mathbf{W}} \in (\mathbb{F}_q^k)^{\otimes s}$ be defined by $\forall \mathbf{i} \in [k]^{s-1}, j \in [k] : \overline{\mathbf{W}}_{[j, \mathbf{i}]} = \pi_j(\mathbf{W}_i)$, recalling that \mathbf{W}_i is deterministic. Let $\hat{\mathbf{Z}}$ be the return value in Line 15. Now, the assignment in Line 14 sets

$$\hat{\mathbf{Z}} = \overline{\mathbf{W}} \cdot (M^{-1} \otimes I_k^{\otimes s-1})$$

where I_k is the $k \times k$ Identity. Thus, we have

$$\begin{aligned} \hat{\mathbf{U}}_{[\ell, \cdot]}^{\text{F}} &:= [\text{FAST-DECODER}(\mathbf{Z}; \boldsymbol{\alpha}) \cdot M^{\otimes s}]_{[\ell, \cdot]} \\ &= [\hat{\mathbf{Z}} \cdot M^{\otimes s}]_{[\ell, \cdot]} \\ &= [\overline{\mathbf{W}} \cdot (M^{-1} \otimes I_k^{\otimes s-1}) \cdot M^{\otimes s}]_{[\ell, \cdot]} \\ &= [\overline{\mathbf{W}} \cdot (I_k \otimes M^{\otimes s-1})]_{[\ell, \cdot]} \\ &= \overline{\mathbf{W}}_{[\ell, \cdot]} \cdot M^{\otimes s-1} \\ &= \{\pi_\ell(\mathbf{W}_i)\}_{i \in [k]^{s-1}} \cdot M^{\otimes s-1} \\ &= \hat{\mathbf{V}}^{(\ell)} \cdot M^{\otimes s-1}, \end{aligned} \quad (\text{by Line 12 in FAST-DECODER})$$

which establishes (21).

Combined with our inductive assumption, this gives

$$\forall \ell < j : \hat{\mathbf{V}}^{(\ell)} \cdot M^{\otimes s-1} = \hat{\mathbf{U}}_{[\ell, \cdot]}^{\text{F}} = \hat{\mathbf{U}}_{[\ell, \cdot]}^{\text{SC}}. \quad (22)$$

With this setup, by definition of the FAST-DECODER we have:

$$\mathbf{Z}'^{(j)} := \{\pi_j(\mathbf{W}_i)\}_{i \in [k]^{s-1}} \quad (23)$$

and

$$\hat{\mathbf{V}}^{(j)} := \text{FAST-DECODER}(\mathbf{Z}'^{(j)}; \boldsymbol{\alpha}_{[j, \cdot]}) \quad (24)$$

where (by Line 8 in FAST-DECODER)

$$\mathbf{W}_i \equiv \mathbf{Z}_{[\cdot, i]} M \mid \{ \forall \ell < j : \pi_\ell(\mathbf{Z}_{[\cdot, i]} M) = \hat{\mathbf{V}}_i^{(\ell)} \} \quad (25)$$

The main observation is the following.

Claim A.5. *The following distributions are identical:*

$$\{\mathbf{U}_{[j,\cdot]} \mid \mathbf{U}_{[<j,\cdot]} = \hat{\mathbf{U}}_{[<j,\cdot]}^{\text{SC}}\} \equiv \mathbf{Z}'^{(j)} \cdot M^{\otimes s-1}$$

Proof. By the structure of the tensor product (using our joint distribution $\mathbf{U} := \mathbf{Z}M^{\otimes s}$), we have $\forall \ell \in [k] : \mathbf{U}_{[\ell,\cdot]} = \{\pi_\ell(\mathbf{Z}_{[\cdot,i]}M)\}_{i \in [k]^{s-1}} \cdot M^{\otimes s-1}$, or equivalently,

$$\forall \ell \in [k] : \mathbf{U}_{[\ell,\cdot]} \cdot (M^{-1})^{\otimes s-1} = \{\pi_\ell(\mathbf{Z}_{[\cdot,i]}M)\}_{i \in [k]^{s-1}}. \quad (26)$$

Now, we can re-write the conditional distributions (where the indexing over $\{\dots\}_i$ is always over $\{\dots\}_{i \in [k]^{s-1}}$),

$$\begin{aligned} & \{\mathbf{U}_{[j,\cdot]} \mid \mathbf{U}_{[<j,\cdot]} = \hat{\mathbf{U}}_{[<j,\cdot]}^{\text{SC}}\} \\ & \equiv \{\mathbf{U}_{[j,\cdot]} \mid \{\forall \ell < j : \mathbf{U}_{[\ell,\cdot]} = \hat{\mathbf{U}}_{[\ell,\cdot]}^{\text{SC}}\}\} \\ & \equiv \{\mathbf{U}_{[j,\cdot]} \mid \{\forall \ell < j : \mathbf{U}_{[\ell,\cdot]}(M^{-1})^{\otimes s-1} = \hat{\mathbf{U}}_{[\ell,\cdot]}^{\text{SC}}(M^{-1})^{\otimes s-1}\}\} \quad (\text{as } (M^{-1})^{\otimes s-1} \text{ has full rank}) \\ & \equiv \{\mathbf{U}_{[j,\cdot]} \mid \{\forall \ell < j : \mathbf{U}_{[\ell,\cdot]}(M^{-1})^{\otimes s-1} = \hat{\mathbf{V}}^{(\ell)}\}\} \quad (\text{Equation (22)}) \\ & \equiv \{\mathbf{U}_{[j,\cdot]} \mid \{\forall \ell < j : \{\pi_\ell(\mathbf{Z}_{[\cdot,i]}M)\}_i = \hat{\mathbf{V}}^{(\ell)}\}\} \quad (\text{Equation (26)}) \\ & \equiv \{\{\pi_j(\mathbf{Z}_{[\cdot,i]}M)\}_i \cdot M^{\otimes s-1} \mid \{\forall \ell < j : \{\pi_\ell(\mathbf{Z}_{[\cdot,i]}M)\}_i = \hat{\mathbf{V}}^{(\ell)}\}\} \quad (\text{Equation (26)}) \\ & \equiv \{\{\pi_j(\mathbf{W}_i)\}_i \cdot M^{\otimes s-1}\} \quad (\star) \\ & \equiv \{\mathbf{Z}'^{(j)} \cdot M^{\otimes s-1}\}. \quad (\text{by definition of } \mathbf{Z}'^{(j)}) \end{aligned}$$

Line (\star) follows by noting that since \mathbf{Z} is a product distribution, the joint distributions $\mathbf{Z}_{[\cdot,i]}$ for each i are conditionally independent given the event $\{\forall \ell < j : \{\pi_\ell(\mathbf{Z}_{[\cdot,i]}M)\}_i = \hat{\mathbf{V}}^{(\ell)}\}$. Thus, these distributions may be equivalently sampled via \mathbf{W}_i , which agrees on the marginals by definition (Equation (25)). This concludes the proof of Claim A.5. \blacksquare

We now have:

$$\hat{\mathbf{U}}_{[j,\cdot]}^{\text{F}} = \hat{\mathbf{V}}^{(j)} \cdot M^{\otimes s-1} \quad (\text{Equation (21)})$$

$$= \text{FAST-DECODER}(\mathbf{Z}'^{(j)}; \boldsymbol{\alpha}_{[j,\cdot]}) \cdot M^{\otimes s-1}. \quad (\text{Equation (24)})$$

We may now apply the inductive hypothesis for $(s-1)$. Letting

$$\mathbf{U}' := \mathbf{Z}'^{(j)} \cdot M^{\otimes s-1} \quad (27)$$

and

$$\hat{\mathbf{U}}'^{\text{F}} := \text{FAST-DECODER}(\mathbf{Z}'^{(j)}; \boldsymbol{\alpha}_{[j,\cdot]}) \cdot M^{\otimes s-1},$$

the inductive hypothesis guarantees that $\hat{\mathbf{U}}'^{\text{F}} = \text{SC-DECODER}(\mathbf{Z}'^{(j)}; \boldsymbol{\alpha}_{[j,\cdot]}|_S)$, and thus

$$\forall \mathbf{i} \in [k]^{s-1} : \hat{\mathbf{U}}_{\mathbf{i}}'^{\text{F}} = \begin{cases} \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}'_{\mathbf{i}} = x \mid \mathbf{U}'_{\prec \mathbf{i}} = \hat{\mathbf{U}}_{\prec \mathbf{i}}'^{\text{F}}] & \text{for } [j, \mathbf{i}] \notin S \\ \boldsymbol{\alpha}_{[j, \mathbf{i}]} & \text{for } [j, \mathbf{i}] \in S. \end{cases}$$

Thus, for all indices $[j, \mathbf{i}] \notin S$, we have:

$$\begin{aligned}
\hat{U}_{[j, \mathbf{i}]}^{\text{F}} &= \hat{U}_{\mathbf{i}}^{\prime \text{F}} \\
&= \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}'_{\mathbf{i}} = x | \mathbf{U}'_{\prec \mathbf{i}} = \hat{U}_{\prec \mathbf{i}}^{\text{F}}] \\
&= \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}'_{\mathbf{i}} = x | \mathbf{U}'_{\prec \mathbf{i}} = \hat{U}_{[j, \prec \mathbf{i}]}^{\text{F}}] && \text{(as } \hat{U}_{[j, \cdot]}^{\text{F}} = \hat{U}^{\prime \text{F}}) \\
&= \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[(\mathbf{Z}'^{(j)} M^{\otimes s-1})_{\mathbf{i}} = x | (\mathbf{Z}'^{(j)} M^{\otimes s-1})_{\prec \mathbf{i}} = \hat{U}_{[j, \prec \mathbf{i}]}^{\text{F}}] && \text{(Equation (27))} \\
&= \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}_{[j, \mathbf{i}]} = x | \mathbf{U}_{[j, \prec \mathbf{i}]} = \hat{U}_{[j, \prec \mathbf{i}]}^{\text{F}} \wedge \mathbf{U}_{[\prec j, \cdot]} = \hat{U}_{[\prec j, \cdot]}^{\text{SC}}]. && \text{(Claim A.5)}
\end{aligned}$$

Since this relation holds for all indices $[j, \mathbf{i}] \notin S$ (and for $[j, \mathbf{i}] \in S$ we trivially have $\hat{U}_{[j, \mathbf{i}]}^{\text{F}} = \alpha_{[j, \mathbf{i}]} = \hat{U}_{[j, \mathbf{i}]}^{\text{SC}}$), we can unwrap the above relation by induction on \mathbf{i} to find that:

$$\begin{aligned}
\hat{U}_{[j, \mathbf{i}]}^{\text{F}} &= \operatorname{argmax}_{x \in \mathbb{F}_q} \Pr[\mathbf{U}_{[j, \mathbf{i}]} = x | \mathbf{U}_{[j, \prec \mathbf{i}]} = \hat{U}_{[j, \prec \mathbf{i}]}^{\text{SC}} \wedge \mathbf{U}_{[\prec j, \cdot]} = \hat{U}_{[\prec j, \cdot]}^{\text{SC}}] \\
&= \hat{U}_{[j, \mathbf{i}]}^{\text{SC}}.
\end{aligned}$$

Thus we have shown that $\hat{U}_{[j, \cdot]}^{\text{F}} = \hat{U}_{[j, \cdot]}^{\text{SC}}$, completing the inductive step. This concludes the proof of Lemma A.4. \blacksquare

This equivalence (Lemma A.4), together with Theorem A.3 on the correctness of the SC-DECODER, suffices to prove Theorem 1.9.

A.2.3 Arıkan Martingale and Polar Coding

Here we build a correspondence between the definition of the Arıkan Martingale and the process of polar coding.

Let $\mathbf{Z} \in \mathbb{F}_q^{k^t}$, $\mathbf{Y} \in \mathbb{F}_q^{k^t}$, and $\mathbf{U} \in \mathbb{F}_q^{k^t}$. We think of \mathbf{Z} as the channel inputs, \mathbf{Y} as the channel outputs, and \mathbf{U} as the encoding inputs.

Lemma A.6. *For a matrix $M \in \mathbb{F}_q^{k \times k}$ and symmetric channel $\mathcal{C}_{Y|Z}$, let $\{X_t\}$ be the associated Arıkan Martingale. For a given t , let $L = M^{\otimes t}$ be the polarization transform, and let $n = k^t$ be the blocklength. Let the channel inputs Z_i be i.i.d. uniform in \mathbb{F}_q , and channel outputs $\mathbf{Y}_i \sim \mathcal{C}(Z_i)$.*

Then, for a uniformly random index $i \in [n]$, the normalized entropy $\bar{H}((\mathbf{Z}L)_i | \mathbf{Y}, (\mathbf{Z}L)_{\prec i})$ is distributed identically as X_t .

Proof. Throughout this proof, we will switch to considering vectors in $\mathbb{F}_q^{k^t}$ as tensors in $(\mathbb{F}_q^k)^{\otimes t}$, for convenience — this correspondence is induced by lexicographic ordering \prec on tuples $[k]^t$. Also, we will write $P(\mathbf{Z})$ to mean the operator P acting on \mathbf{Z} . In this notation, we wish to show that for a uniformly random multiindex $\mathbf{i} \in [k]^t$, the entropy $H((M^{\otimes t}(\mathbf{Z}))_{\mathbf{i}} | \mathbf{Y}, (M^{\otimes t}(\mathbf{Z}))_{\prec \mathbf{i}}) \sim X_t$.

We will show by induction that for all t , there is some permutation of coordinates⁷ $\sigma' : [k]^t \rightarrow [k]^t$ such that the joint distributions

$$\{(\mathbf{A}', \mathbf{B}')\}_{(\mathbf{A}', \mathbf{B}') \sim D_t} \equiv \{(M^{\otimes t}(\mathbf{Z}), \sigma'(\mathcal{C}(\mathbf{Z})))\}_{\mathbf{U} \sim (\mathbb{F}_q^k)^{\otimes t}}$$

⁷This is in fact just a reversal of the co-ordinates, i.e. $\sigma'((i_1, i_2, \dots, i_t)) = (i_t, \dots, i_2, i_1)$.

where $(\mathbf{A}', \mathbf{B}') \sim D_t$ are the distributions defined in the t -th step of the Arıkan martingale, and $\mathbf{Z} \sim (\mathbb{F}_q^k)^{\otimes t}$ is sampled with iid uniform coordinates. This is sufficient, because a permutation of the channel outputs does not affect the relevant entropies. That is,

$$H(\mathbf{A}'_i \mid \mathbf{B}', \mathbf{A}'_{\setminus i}) = H(\mathbf{A}'_i \mid \sigma'(\mathbf{B}'), \mathbf{A}'_{\setminus i}).$$

First, the base case $t = 0$ follows by definition of the distribution D_0 in the Arıkan martingale.

For the inductive step, assume the claim holds for $t - 1$. Let σ be the permutation guaranteed for $t - 1$. For each $j \in [k]$, sample an independent uniform $\mathbf{Z}^{(j)} \sim (\mathbb{F}_q^k)^{\otimes t-1}$ and define

$$(\mathbf{A}^{(j)}, \mathbf{B}^{(j)}) := (M^{\otimes t-1}(\mathbf{Z}^{(j)}), \sigma(\mathcal{C}(\mathbf{Z}^{(j)}))). \quad (28)$$

By the inductive hypothesis, $(\mathbf{A}^{(j)}, \mathbf{B}^{(j)}) \sim D_{t-1}$, for each $j \in [k]$.

As in the Arıkan martingale, define $(\mathbf{A}', \mathbf{B}')$ deriving from $\{(\mathbf{A}^{(j)}, \mathbf{B}^{(j)})\}_{j \in [k]}$ as

$$\mathbf{A}'_{[i, \cdot]} := M((\mathbf{A}_i^{(1)}, \dots, \mathbf{A}_i^{(k)})) \quad \text{and} \quad \mathbf{B}'_{[j, \cdot]} := \mathbf{B}^{(j)}. \quad (29)$$

Note that \mathbf{B}' can equivalently be written (unwrapped) as

$$\mathbf{B}' := (\mathbf{B}^{(1)}, \mathbf{B}^{(2)}, \dots, \mathbf{B}^{(k)})$$

By definition of the Arıkan martingale, we have $(\mathbf{A}', \mathbf{B}') \sim D_t$.

Finally, define $\mathbf{Z} \in (\mathbb{F}_q^k)^{\otimes t}$ by

$$\mathbf{Z}_{[\cdot, j]} := \mathbf{Z}^{(j)}.$$

To finish the proof, we will show that $(\mathbf{A}', \mathbf{B}') = (M^{\otimes t}(\mathbf{Z}), \sigma'(\mathcal{C}(\mathbf{Z})))$ for some permutation σ' .

The main claim is the following.

Claim A.7. *For every instantiation of the underlying randomness \mathbf{Z} , we have*

$$\mathbf{A}' = M^{\otimes t}(\mathbf{Z}).$$

Proof of Claim A.7. Expanding the recursive definition of the tensor product, Equation (2), we have:

$$[M^{\otimes t}(\mathbf{Z})]_{[i, \cdot]} = M((\mathbf{W}_i^{(1)}, \mathbf{W}_i^{(2)}, \dots, \mathbf{W}_i^{(k)}))$$

where

$$\mathbf{W}^{(j)} := M^{\otimes t-1}(\mathbf{Z}_{[\cdot, j]}) = M^{\otimes t-1}(\mathbf{Z}^{(j)}) = \mathbf{A}^{(j)}.$$

Where the last equality is by the inductive assumption. Thus,

$$\begin{aligned} [M^{\otimes t}(\mathbf{Z})]_{[i, \cdot]} &= M((\mathbf{A}_i^{(1)}, \dots, \mathbf{A}_i^{(k)})) \\ &= \mathbf{A}'_{[i, \cdot]}. \end{aligned} \quad (\text{By definition, Equation (29)})$$

And so $M^{\otimes t}(\mathbf{Z}) = \mathbf{A}'$ as desired. ■

Continuing the proof of Lemma A.6, we now have

$$\begin{aligned}
(\mathbf{A}' , \mathbf{B}') &= (\mathbf{A}' , (\mathbf{B}^{(1)}, \mathbf{B}^{(2)}, \dots, \mathbf{B}^{(k)})) && \text{(By definition, Equation (29))} \\
&= (\mathbf{A}' , (\sigma(\mathcal{C}(\mathbf{Z}^{(1)})), \sigma(\mathcal{C}(\mathbf{Z}^{(2)})), \dots, \sigma(\mathcal{C}(\mathbf{Z}^{(k)})))) \\
&&& \text{(Definition of sampling, Equation (28))} \\
&= (\mathbf{A}' , \sigma'(\mathcal{C}(\mathbf{Z}))) && (\star) \\
&= (M^{\otimes t}(\mathbf{Z}) , \sigma'(\mathcal{C}(\mathbf{Z}))).
\end{aligned}$$

In the above, the equality in line (\star) , follows from the fact in both cases, entries of the tensor are $\{\mathcal{C}(\mathbf{Z}_i)\}_i$, and they only differ by some permutation of coordinates.

Since we established $(\mathbf{A}', \mathbf{B}') \sim D_t$ by definition, this completes the proof. ■