

Universal Polar Coding and Sparse Recovery

Emmanuel Abbe

Information Processing Group, EPFL
Email: emmanuel.abbe@epfl.ch

Abstract—This paper investigates universal polar coding schemes. In particular, a notion of ordering (called convolutional path) is introduced between probability distributions to determine when a polar compression (or communication) scheme designed for one distribution can also succeed for another one. The original polar decoding algorithm is also generalized to an algorithm allowing to learn information about the source distribution using the idea of checkers. These tools are used to construct a universal compression algorithm for binary sources, operating at the lowest achievable rate (entropy), with low complexity and with guaranteed small error probability.

In a second part of the paper, the problem of sketching high dimensional discrete signals which are sparse is approached via the polarization technique. It is shown that the number of measurements required for perfect recovery is competitive with the $O(k \log(n/k))$ bound (with optimal constant for binary signals), meanwhile affording a deterministic low complexity measurement matrix.

I. INTRODUCTION

A new technique called ‘polarization’ has recently been introduced by Arıkan in [3] to construct efficient channel coding schemes. The codes resulting from this technique, called polar codes, have several nice attributes: (1) they are linear codes generated by a low-complexity deterministic matrix (2) they can be analyzed mathematically and bounds on the error probability (exponential in the square root of the block length) can be proved (3) they have a low encoding and decoding complexity (4) they allow to reach the Shannon capacity on any discrete memoryless channels (DMC). These codes are indeed the first codes with low decoding complexity that are provably capacity achieving on any DMC.

Since [3], the polarization technique has been generalized to various settings. For example, it has been used in [18] for rate-distortion via duality with test channels, in [19], [16], [20] for wiretap channels and information secrecy, and in [24], [2] for a multi-user problem (multiple accessing).

In this paper, we investigate the problem of robustness of the polar coding schemes with respect to the knowledge of source or channel distribution. The perfect knowledge of this distribution is never available, and it is important that any potentially practical scheme shows some robustness to this knowledge. We hence develop several tools to construct universal polarization schemes.

We then consider the problem of sketching high-dimensional sparse signal using the polarization technique. The hope being to leverage properties (1)-(4) to construct a deterministic low-complexity sketching matrix and an efficient sparse recovery algorithm. Since the method is defined for signals valued in finite sets, it is of interest to lift the

construction to the real setting. Yet in this paper, we focus our attention on the sketching problem for signals that are discrete, motivated by applications dealing with such signals. This occurs for example in network monitoring problems [12], [13]. We will see that, just like one can exploit sparsity in the domain, the sparsity in the magnitude (signals taking values in finite sets) can be exploited to develop an efficient sketching method via the polarization technique.

Some results in this paper have been presented in [1].

A. Channel and source polarization

Arıkan shows in [3] that an arbitrary binary input discrete memoryless channel W can be polarized as follows: n independent uses of W can be transformed into n successive uses of synthesized binary input channels that have (except for a vanishing fraction) a symmetric capacity which tends to either 0 or 1 (with n). In [23], this result is generalized to q -ary input alphabets where q is prime, and in [2] it is extended to q being powers of two (considering q to be a power of two has computational advantages, but the case of powers of prime follows too). We state here the result of [23] for q prime. Notation: $X^n := (X_1, \dots, X_n)$.

Theorem 1. *Let W be a q -ary input discrete memoryless channel with q prime, n a power of 2, and let U^n be i.i.d. uniform random variables on \mathbb{F}_q . Let $X^n = U^n G_n$, where $G_n = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}^{\otimes \log_2(n)}$, and Y^n be the output of n independent uses of W when the input is X^n . Then, for any $\delta \in (0, 1)$,*

$$\frac{1}{n} |\{i : I(U_i; Y^n U^{i-1}) > \delta\}| \xrightarrow{n \rightarrow \infty} I(W), \quad (1)$$

where $I(W)$ is the mutual information of W (with a uniform input distribution).

Theorem 1 can then be used to show the following polarization phenomenon for sources.

Theorem 2. *Let X^n be n i.i.d. random variables with distribution p on \mathbb{F}_q , n a power of 2, and let $U^n = X^n G_n$, where $G_n = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}^{\otimes \log_2(n)}$. Then, for any $\delta \in (0, 1)$,*

$$\frac{1}{n} |\{i : H(U_i | U^{i-1}) > \delta\}| \xrightarrow{n \rightarrow \infty} H(p), \quad (2)$$

where $H(p)$ is the entropy of the distribution p .

We will see in Section III that previous result follows from Theorem 1 via a duality argument. A slightly more general result is presented in [5].

Note that all entropies and mutual informations are computed with logarithms in base q (where q is the input or source alphabet size).

A coding scheme from Theorem 1. The limit in the theorem implies that for n large enough and except for a vanishing fraction of indices i , $I(U_i; Y^n U^{i-1})$ must be close to either 0 or 1. Hence, this suggests a coding scheme: on the indices i for which the channel is good, i.e., $I(U_i; Y^n U^{i-1})$ is close to 1, put uncoded information bits in U_i , and for the other indices, put frozen bits that are predetermined and revealed to the decoder. This defines the vector U^n . Then, the vector X^n is sent over n independent uses of the channel. Note that the rate of this code is given by the logarithm of the number of information bits in U^n divided by n , and by Theorem 1, this can be made arbitrarily close to $I(W)$. Now the receiver knows two things: the location¹ of the indices i containing information and frozen bits, and the value of the frozen bits (on symmetric channels, the frozen bits can be all chosen to be zero). Hence, from the output Y^n of X^n , the receiver starts by decoding the first component U_i which is not frozen. By virtue of Theorem 1, one of the two possible value of U_i will have (w.h.p.) a probability close to one and hence, the decoder has a small probability of decoding U_i incorrectly. This process is then iterated to decode successively the entire vector U^n . An analysis of the scaling² of the bit error probability (decoding wrongly a component in U_i) allows to conclude that w.h.p. errors cannot propagate in this scheme, and hence, this scheme achieves the uniform mutual information of the channel. A remarkable feature of this coding scheme is that the encoding and decoding complexity is shown to be $O(n \log n)$.

A coding scheme from Theorem 2. The limit in the theorem implies that for n large enough and except for a vanishing fraction of indices i , $H(U_i|U^{i-1})$ is close to either 0 or 1. Hence, the transformation G_n extracts the randomness in X^n , which is initially uniformly dissipated over the n components, into specific components indexed by the i 's such that $H(U_i|U^{i-1})$ is close to 1. Lossless compression can then be performed as follows: from a given source output X^n , compute U^n and store the components of U^n which do not have an entropy close to 0. Note that, from Theorem 2, the compression rate can be made arbitrarily close to $H(p)$ (lowest possible rate). For the reconstruction, since the components with entropy close to 0 can be recovered correctly with high probability given the past components, we can proceed successively in an analogue manner as for the channel decoding problem. The speed of polarization is shown to scale similarly as in the channel case, and again, the encoding and decoding complexity of this source coding scheme is only $O(n \log n)$.

B. Goals

In this paper, we are interested in analyzing how sensitive the performance of the previous source/channel coding scheme

¹No analytical formula is known to compute these indices. They are found with algorithms, as in [25]

²To show achievability, the speed of convergence to the polarized channels matters, and it is shown to be roughly $2^{-\sqrt{n}}$ in [4].

is to the knowledge of the source/channel distribution. The knowledge of source/channel distribution is used at two moments for each problem. In the channel coding problem, it is first used to identify the location of the “good channels”, or equivalently, the location of the indices i where the information bits shall be sent. It is then used again in the decoding process, to compute the probabilities that an information bit U_i is equal each element of \mathbb{F}_q (from the polarization phenomenon, we know that one of these probabilities should have a probability close to 1, but one still needs to compute which one it is). Similarly, for source coding, the knowledge of the source distribution is first used to find the components of U^n which must be stored, and then in the reconstruction part, to compute the value of each non-stored components.

We will hence address the problem of constructing polar coding schemes which can compress losslessly sources without requiring perfect knowledge of their distributions, or which can communicate reliably over channels without requiring perfect knowledge of the channel distribution. The application to the channel setting follows then from Section III, where the duality between the source and channel problem is made explicit.

We will then consider the problem of sparse data recovery, using polar codes. From the discussion on the source polarization theorem above, a connection to the sketching problem is apparent: if we sense the signal U^n only in the components i for which $H(U_i|U^{i-1})$ is close to 1, we obtain a sampling of the signal which allows perfect recovery of the full signal, with a significantly reduced number of measurements. There are however several differences between a compressed sensing setting [7], [11] and the source polarization setting; in particular, in the latter setting the source is random with a known distribution and it is valued in a finite field (of arbitrarily large cardinality), whereas it is real and with no prior (besides sparse) in compressed sensing. Hence, a first question is to ask how sparsity, i.e., the property of having many components that are 0, is modeled for such random signals, and how much the choice of a specific sparse probability distribution matters. This part can be investigated using our results on universal source polarization, which establishes the connection between the two parts of this paper.

II. RESULTS

A universal compression algorithm for binary sources is introduced in Section V-A. Theorem 3 shows that this algorithm performs at the lowest achievable rate (entropy), with a $O(n \log_2 n)$ complexity and (roughly) a $O(2^{-\sqrt{n}})$ error probability.

Partial generalizations are discussed for non-binary sources in Section V-B.

In Section VI-B, a low-complexity deterministic sketching matrix is constructed. It is shown in Theorem 4 that for k -sparse signals in \mathbb{F}_a^n , $O(k \log_e n/k)$ measurements taken with the proposed sketching matrix are sufficient to recover perfectly the original vector with a probability at least $1 - O(2^{-\sqrt{n}})$, and a reconstruction algorithm of complexity

$O(a \log_2 a \cdot n \log_2 n)$. An improved version of this Theorem (regarding the dependence in a of the constants) is investigated in Section VI-E.

III. DUALITY BETWEEN SOURCE AND CHANNEL POLARIZATION

In this section, we connect Theorem 1 and Theorem 2. Let p be a distribution on \mathbb{F}_q and consider using Theorem 1 for an additive noise channel, i.e., $Y = X \oplus Z$ for Z distributed under p and independent of X . We then have $Y^n = G_n U^n \oplus Z^n$ and

$$\begin{aligned} I(U_i; Y^n U^{i-1}) &= 1 - H(U_i | Y^n U^{i-1}) \\ &= 1 - H((G_n Y^n \ominus G_n Z^n)_i | Y^n (G_n Y^n \ominus G_n Z^n)^{i-1}) \\ &= 1 - H((G_n Z^n)_i | (G_n Z^n)^{i-1}). \end{aligned} \quad (3)$$

Equality (3) uses the fact that Y^n is independent of Z^n because U^n , and hence $G_n U^n$, are uniformly distributed over \mathbb{F}_q . We also use the fact that $G_n^{-1} = G_n$. Hence, Theorem 1 and (3) imply Theorem 2.

Stated as such, Theorem 2 does not imply Theorem 1, since additive noise channels are not representative of all possible channels. In [5] a slightly more general result than Theorem 2 is stated, where an auxiliary random variable Y (side-information), which is a random variable correlated with X but not intended to be compressed, is introduced in the conditioning of each entropy term. This could be used for the reverse implication.

In this paper, we focus mostly on the source setting, since it is the “simplest” setting, hence the one to start with. Using previous expansions, the results obtained in the source setting will directly admit a counter-part in the channel setting, for the case of additive noise channels.

IV. DEFINING ORDERINGS AND MATHEMATICAL PRELIMINARIES

Definition 1 (Measures). Let a be a prime integer, $\mathbb{F}_a := \{0, 1, \dots, a-1\}$ and $M(a)$ be the set of probability measures on \mathbb{F}_a . For any $k \in \mathbb{F}_a$, let

$$\widehat{M}_k(a) := \{p \in M(a) : p(i) = p(j), \forall i, j \neq k, p(k) \geq \frac{a-1}{a}\}$$

and $\widehat{M}(a) := \cup_{k \in \mathbb{F}_a} \widehat{M}_k(a)$. We refer to the measure in $\widehat{M}(a)$ as the spike measures.

Definition 2 (Matrices). We denote by $\text{Doub}(a)$ the set of doubly stochastic matrices of size $a \times a$, and by $\text{Circ}(a)$ the set of circulant stochastic matrices of size $a \times a$.

Definition 3 (Orders). We define

$$p_1 \prec_h p_2 \quad \text{if} \quad h(p_1) \geq h(p_2), \quad (4)$$

$$p_1 \prec_d p_2 \quad \text{if} \quad p_1 = D p_2 \text{ for } D \in \text{Doub}(a), \quad (5)$$

$$p_1 \prec_c p_2 \quad \text{if} \quad p_1 = C p_2 \text{ for } C \in \text{Circ}(a). \quad (6)$$

Note that \prec_d is the majorization order and $p_1 \prec_c p_2$ is equivalent to $p_1 = c \star p_2$ for $c \in M(a)$, where \star denotes the circular convolution on \mathbb{F}_a .

Note that we use the term “order” in a broad sense here (not a mathematical order).

Lemma 1 (Orders hierarchy).

$$p_1 \prec_c p_2 \quad \Rightarrow \quad p_1 \prec_d p_2 \quad \Rightarrow \quad p_1 \prec_h p_2. \quad (7)$$

Proof: The first implication follows from the fact that $\text{Circ}(a) \subset \text{Doub}(a)$ and the second implication follows from the Schur-concavity of the entropy [21]. ■

One can easily find examples showing that there is no reverse implications in Lemma 1. In this paper, we are interested in the \prec_c order, and previous Lemma gives a first idea on how this order compares to the majorization order. But we will only work with \prec_c in this paper. Also note that the set of measures which are worst than a given $p \in M(a)$ with respect to \prec_c is given by the convex hull of the orbit of p through cycles, whereas it is given by the convex hull of the orbit of p through permutations when considering \prec_d . Note that if $p \in \widehat{M}(a)$, these two sets are the same.

Definition 4. For $p \in M(a)$, we define the Fourier transform of p by

$$\mathcal{F}(p)(\omega) = \sum_{k=0}^{a-1} p(k) e^{-2\pi i k \omega / a}, \quad \omega \in \mathbb{F}_a \quad (8)$$

and the inverse Fourier transform of $h : \mathbb{F}_a \rightarrow \mathbb{C}$ by

$$\mathcal{F}^{-1}(h)(k) = \frac{1}{a} \sum_{w=0}^{a-1} h(w) e^{2\pi i k w / a}, \quad k \in \mathbb{F}_a. \quad (9)$$

Remark 1.

- $\mathcal{F}(p \star q) = \mathcal{F}(p) \mathcal{F}(q)$ for any $p, q \in M(a)$.
- If $p \in \widehat{M}_k(a)$ with $p(k) = 1 - P$, we have that $\mathcal{F}(p)$ is given by $\mathcal{F}(p)(0) = 1$ and

$$\mathcal{F}(p)(\omega) = \left(1 - \frac{aP}{a-1}\right) e^{-2\pi i k \omega / a}, \quad \omega \neq 0. \quad (10)$$

- From previous remark, note that $(\widehat{M}(a), \star)$ is a semi-group.

Definition 5. For $p \in M(a)$, let $\text{DOM}_c(p)$ be the set of probability measures which dominate p with respect to \prec_c , i.e., $\text{DOM}_c(p) = \{q \in M(a) : p \prec_c q\}$.

Remark 2. Note that it is easier to describe the set of measures that are dominated by a fixed measure p than the reverse. However, we can write $\text{DOM}_c(p) = \{q \in M(a) : \mathcal{F}^{-1}(\mathcal{F}(p)/\mathcal{F}(q)) \geq 0\}$, and we can use the FFT algorithm to compute $\text{DOM}_c(p)$ efficiently.

Lemma 2. For any $a \geq 1$,

$$p_1, p_2 \in \widehat{M}(a), \quad p_1 \prec_h p_2 \quad \Rightarrow \quad p_1 \prec_c p_2. \quad (11)$$

Proof: Assume that $p_1 \in \widehat{M}_k(a)$ and $p_2 \in \widehat{M}_l(a)$ for $k, l \in \mathbb{F}_a$, and $p_1 \prec_h p_2$. Then, denoting $1 - P_1 = p_1(k)$ and $1 - P_2 = p_2(l)$,

$$\mathcal{F}(p_1)(\omega) / \mathcal{F}(p_2)(\omega) = \frac{1 - \frac{aP_1}{a-1}}{1 - \frac{aP_2}{a-1}} e^{-2\pi i (k \ominus a l) / a}. \quad (12)$$

Hence, if $(1 - \frac{aP_1}{a-1})/(1 - \frac{aP_2}{a-1}) \in \text{Im}(f)$, where $f : P \in [0, 1] \mapsto (1 - aP/(a-1))$, we have that (12) is the Fourier transform of an element in $\widehat{M}(a)$. This is easily verified since $\text{Im}(f) = [0, 1]$ and since by assumption $1 - P_1 \leq 1 - P_2$. ■

Since $\widehat{M}(2) = M(2)$, we have the following corollary.

Corollary 1.

$$p_1, p_2 \in M(2), \quad p_1 \prec_h p_2 \quad \Rightarrow \quad p_1 \prec_c p_2. \quad (13)$$

We now introduce one more ordering notion.

Definition 6. We define for $p_1, p_2 \in M(a)$

$p_1 \prec_{cp} p_2$ iff $p_1 = p_2 \star \nu$ where ν is an infinitely divisible probability distribution.

Definition 7. A probability distribution $p \in M(a)$ is infinitely divisible if for any $k \geq 1$, there exists $p_k \in M(a)$ such that

$$p = \star_{i=1}^k p_k,$$

or equivalently, if $\mathcal{F}(\mathcal{F}(p)^{1/k}) \geq 0$.

Note that checking the infinitely divisibility condition for a large enough k implies the result for smaller k 's (by grouping the p_k 's). Hence, denoting $\varepsilon = 1/k$, we need to check that $\mathcal{F}(p)^\varepsilon$ has a valid inverse Fourier transform when ε tends to 0. Let $z = \mathcal{F}(p)$ and denote the component of z by $z_j = r_j e^{i\theta_j}$. Then

$$\begin{aligned} z_j^\varepsilon &= r_j^\varepsilon e^{i\varepsilon\theta_j} = (1 + \varepsilon \log_e r_j)(1 + i\varepsilon\theta_j) + o(\varepsilon) \\ &= 1 + \varepsilon(\log_e r_j + i\theta_j) + o(\varepsilon). \end{aligned}$$

Hence, by the linearity of \mathcal{F}^{-1} ,

$$\mathcal{F}^{-1}(z^\varepsilon) = (1, 0, \dots, 0) + \varepsilon \mathcal{F}^{-1}((\log_e r_j + i\theta_j)_{j=0}^{a-1}) + o(\varepsilon)$$

and to ensure $\mathcal{F}^{-1}(z^\varepsilon) \geq 0$ for any $\varepsilon > 0$, we need to ensure that

$$\begin{aligned} y(1), \dots, y(a-1) &\geq 0 \\ \text{where } y &= \mathcal{F}^{-1}((\log_e r_j + i\theta_j)_{j=0}^{a-1}). \end{aligned} \quad (14)$$

Note that the dependency in k has been removed in previous condition.

To summarize: we have defined a notion of ‘convolution ordering’, with \prec_c , where one can reach a distribution from another one with a circular convolution, and a notion of ‘convolutional path ordering’, with \prec_{cp} , where one can reach the second distribution with small convolutional steps.

V. UNIVERSALITY IN POLARIZATION

As mentioned in the introduction, there are two parts which require the knowledge of the source distribution in the source polar coding scheme: one in the compression and one in the reconstruction part. We present in this section two lemmas to be used for each of these parts in universal results. We start with the compression part.

Definition 8 (Polar storage sets). Let $\delta \in (0, 1)$, n a power of 2 and $p \in M(a)$,

$$\mathcal{S}_{\delta,n}(p) := \{i \in [n] : H(U_i|U^{i-1}) \geq \delta\}$$

where $U^n = X^n G_n$, $G_n = \begin{bmatrix} 1 & 0 \\ & 1 \end{bmatrix}^{\otimes \log_2(n)}$, $X^n \stackrel{\text{iid}}{\sim} p$. We use the notation

$$\mathcal{S}(p_1) \supseteq \mathcal{S}(p_2) \text{ if } \mathcal{S}_{\delta,n}(p_1) \supseteq \mathcal{S}_{\delta,n}(p_2) \quad \forall \delta \in (0, 1), n.$$

(We will sometimes call the components of U^n on \mathcal{S} the information bits.) The reason why we are interested in nested storage sets is clear: if one stores the components of a source distributed under p_1 , it will also store the information components of any source p_2 with $\mathcal{S}(p_1) \supseteq \mathcal{S}(p_2)$ (it will consume more rate than required for compressing a source under p_2 specifically, but it will allow lossless compression for both). However, for the reconstruction, it is not clear whereas the nested structure is sufficient to induce a universal decoding process. But let us postpone for now the reconstruction problem and focus on the nested structure only.

Lemma 3. For any $a \geq 1$,

$$p_1 \prec_c p_2 \quad \Rightarrow \quad \mathcal{S}(p_2) \subseteq \mathcal{S}(p_1). \quad (15)$$

Proof: By assumption, there exists $c \in M(a)$ such that $p_1 = p_2 \star c$. Let $X^n \stackrel{\text{iid}}{\sim} p_2$, $Z^n \stackrel{\text{iid}}{\sim} c$ independent of X^n and $\tilde{X}^n = X^n \oplus Z^n \stackrel{\text{iid}}{\sim} p_1$. Define $U^n = G_n X^n$, $\tilde{U}^n = G_n \tilde{X}^n$ and $W^n = G_n Z^n$, hence $\tilde{U}^n = U^n \oplus W^n$. We have

$$H(\tilde{U}_i|\tilde{U}^{i-1}) \geq H(\tilde{U}_i|\tilde{U}^{i-1}, W^n) \quad (16)$$

$$= H(U_i|U^{i-1}, W^n) \quad (17)$$

$$= H(U_i|U^{i-1}) \quad (18)$$

where the last equality follows from the fact that U^n is independent of W^n since X^n is independent of Z^n . ■

Note that the ordering in (18) indeed holds for all indices. We now investigate the reconstruction problem. We first recall the decoding algorithm used in [3]

Definition 9. [polar-dec algorithm [3], [5]]

Inputs: $p \in M(a)$, $n \in \mathbb{Z}_+$, $\mathcal{S} \subseteq [n]$ and $u[\mathcal{S}] \in \mathbb{F}_a^{|\mathcal{S}|}$.

Output: polar-dec($p, u[\mathcal{S}], n$) $\in \mathbb{F}_a^n$.

The algorithm proceeds as follows:

(0) Initialize $\mathcal{M} = \mathcal{S}$;

(1) Find the smallest integer i in \mathcal{M}^c and compute

$$u_i = \arg \max_{x \in \mathbb{F}_a} \mathbb{P}_p\{U_i = x|u[\mathcal{M}]\};$$

(2) Update $\mathcal{M} = \mathcal{M} \cup \{i\}$ and go back to (1) until $\mathcal{M} = [n]$;

(3) Output $x^n = u^n G_n$ where $u^n = u[\mathcal{M}]$.

The term $\mathbb{P}_p\{U_i = x|u[\mathcal{M}]\}$ is the probability that $U_i = x$ when $U[\mathcal{M}]$ is observed, where $U^n = X^n G_n$ and $X^n \stackrel{\text{iid}}{\sim} p$. It is shown in [3], [5] that the computational cost for each of these probabilities, as well as the overall algorithm, is bounded as $O(n \log_2 n)$ (more precisely $O(a^2 n \log_2 n)$ for the dependence in a and $O(a \log_2(a) \cdot n \log_2 n)$ if [2] is used and a is power of 2). We refer to [3], [5] for the recursive procedure to compute these probabilities, which uses a ‘‘divide and conquer’’ procedure based on the Kronecker structure of G_n .

Definition 10. Let $p_1, p_2 \in M(a)$, $\delta \in (0, 1)$, $n \geq 1$, $X^n \stackrel{\text{iid}}{\sim} p_2$, and $\hat{X}^n = \text{polar-dec}(p_1, U[\mathcal{S}_{\delta,n}(p_1)], n)$, where $U^n =$

$X^n G_n$. We define

$$P_e(p_1|p_2) = \mathbb{P}\{X^n \neq \hat{X}^n\}.$$

Lemma 4. For any $a \geq 1$, $\delta \in (0, 1/2)$, $n \geq 1$,

$$p_1 \prec_{cp} p_2 \Rightarrow P_e(p_1|p_2) \leq P_e(p_2|p_2).$$

Proof: Fix n and $\delta < 1/2$. If p_1 is the uniform distribution, $\mathcal{S}_{\delta,n}(p_1) = \{1, \dots, n\}$ and the claim is clear: since we store all components, the left-hand side error probability is 0. Hence, assume that p_1 is not the uniform distribution.

Let us assume that $a = 2$, the proof for $a > 2$ is similar. Let $p_2 \in \mathcal{M}(2)$ and $q = p_2 \star 1_\rho$, where $1_\rho = [1 - \rho, \rho]$. Since $q \prec_c p_2$, we have $\mathcal{S}_{\delta,n}(q) \supseteq \mathcal{S}_{\delta,n}(p_2)$ and for the components i to be decoded

$$\delta > H(U_i|U^{i-1}) \geq H(V_i|V^{i-1}),$$

where the U_i 's (resp. V_i 's) are i.i.d. under q (resp. p_2). For W_1, \dots, W_n i.i.d. under p and $w, w_1, \dots, w_{i-1} \in \{0, 1\}$, define the mapping

$$f_{w|w^{i-1}} : p \mapsto P(W_i = w | W^{i-1} = w^{i-1}). \quad (19)$$

Note that $f_{w|w^{i-1}}$ is continuous over $\mathcal{M}(2)$ (with the topology induced by \mathbb{R}^2) for any $w, w_1, \dots, w_{i-1} \in \{0, 1\}$.

Also note that $H(U_i|U^{i-1}) < \delta$ implies that there exist $\xi(\delta)$ with $\xi(\delta) \xrightarrow{\delta \rightarrow 0} 0$, such that for any $u^{i-1} \in \{0, 1\}^n$,

$$P(U_i = 0 | U^{i-1} = u^{i-1}) \wedge P(U_i = 1 | U^{i-1} = u^{i-1}) < \xi(\varepsilon). \quad (20)$$

Hence, using (20) and the continuity of $f_{w|w^{i-1}}$, we have for ρ small enough and any i in the complement of $\mathcal{S}_{\delta,n}(q)$,

$$\arg \max_{u \in \{0,1\}} P(U_i = u | U^{i-1} = u^{i-1}) \quad (21)$$

$$= \arg \max_{u \in \{0,1\}} P(V_i = u | V^{i-1} = u^{i-1}). \quad (22)$$

For $p_1 \prec_{cp} p_2$, we have that $p_1 = p_2 \star_{i=1}^k 1_\delta$ for any k where δ depends on k and is decreasing with k increasing. We now want to iterate previous argument, but we have to use the continuity of (19) at different distribution q 's, namely $q = p_2 \star_{i=1}^l 1_\delta$ for $l = 1, \dots, k$. Since this is a compact set of q 's (the entire path from p_2 to p_1), we can pick k large enough such that the continuity argument remains effective along the entire path, and (22) is proved by (20). It is important to assume that p_1 is not uniform, so as to keep ρ bounded below from 0.

Finally, from (22), we have that the algorithm `polar-dec` used with the mismatched distribution still leads to the same output than when used with the matched one. Since in the mismatched scenario we observe more components than needed, strictly speaking we have an inequality in the error probability as in the lemma's statement.

For $a \geq 3$, the proof is identical, except that we are moving along $p_2 \star_{i=1}^k \nu$ where ν is close to a delta function over \mathbb{F}_a , and (20), resp. (22), holds when the minimum, resp. maximum, include all elements of \mathbb{F}_a . ■

This result tells us that, if we do the compression and the reconstruction using the distribution p_1 , we can compress and reconstruct losslessly any source distribution which are better than p_1 in terms of \prec_{cp} .

If one were to use a compression scheme ignoring any complexity considerations, then, simply by knowing that the source distribution has an entropy at most R , it would be possible to compress and reconstruct the source losslessly, using the method of types for example, at rate R . And the set $\{p \in \mathcal{M}(a) : h(p) \leq R\}$ are essentially the largest sets which can be compressed losslessly at a fixed rate. It is ambitious to ask for such a ‘‘broad universality’’ with polar codes, since these are structured codes with complexity attributes, in contrast to the codes derived with the method of types. We may have to give up some extra rate to achieve this goal, or we may universally compress only certain subsets of source distributions. We now investigate these points.

A. Results for binary sources

For binary sources, it is possible to achieve a broad universal result with polar coding.

Notation: For a given $0 \leq R < 1$, let $p_0(R), p_1(R)$ be the two binary probability distributions such that $H(p_0(R)) = H(p_1(R)) = R$.

Definition 11 (Universal polar compression algorithm).

A. Compression:

Inputs: $R \in [0, 1]$ (the rate of compression), δ (the target error probability), $x \in \mathbb{F}_2^n$ (the data).

Output: $v \in \mathbb{F}_2^{nR+o(n)}$ (the stored data).

The compression algorithm proceeds as follows:

1. Compute $u = G_n x$
2. Store $u[\mathcal{S}_{\delta,n}(p_0(R))]$ and $u[n]$

B. Reconstruction:

Inputs: $n, R, u[\mathcal{S}_{\delta,n}(p_0(R))]$ and $u[n]$

Outputs: `polar-dec-adapt`($p_0(R), p_1(R), u[\mathcal{S}_{\delta,n}(p_0(R))]$, $u[n], n$)

Definition 12 (`polar-dec-adapt` algorithm).

Inputs: $p_1, \dots, p_k \in \mathcal{M}(a)$, $n \in \mathbb{Z}_+$, $\mathcal{S} \subseteq [n]$, $\mathcal{T} \subseteq \mathcal{S}^c$ (called the set of checkers) and $u[\mathcal{T} \cup \mathcal{S}] \in \mathbb{F}_a^{|\mathcal{T}|+|\mathcal{S}|}$.

Output: `polar-dec-adapt`($p_1, \dots, p_k, u[\mathcal{S}], u[\mathcal{T}], n$) $\in \mathbb{F}_a^n$.

The algorithm proceeds as follows:

- (1) For $j = 1, \dots, k$, run $u_{(j)}^n = \text{polar-dec}(p_j, u[\mathcal{S}], n)$
- (2) Find $t = \arg \min_{j=1, \dots, k} d_H(u_{(j)}[\mathcal{T}], u[\mathcal{T}])$ (pick one at random for ties)
- (3) Output `polar-dec`($p_t, u[\mathcal{S}], n$). (Variant: output `polar-dec`($p_t, u[\mathcal{S} \cup \mathcal{T}], n$).

Theorem 3. [Universal polar compression] Let $X^n = [X_1, \dots, X_n]$ be i.i.d. Bernoulli with $H(X_1) \leq R$. The universal polar compression algorithm allows to compress X^n at rate R , with error probability $O(2^{-n^\beta})$, for any $\beta < 1/2$, and compression/reconstruction complexity $O(n \log_2 n)$.

(Note: Using the duality argument of Section III, this theorem admits an analogue for universal coding over binary

symmetric channels.)

For the proof of this theorem, we show that:

1. Any binary source which is known to have entropy at most R can be compressed universally with polar codes by storing the information bits on $\mathcal{S}(p^*)$, where p^* is one of the two distributions with entropy R .
2. If it is known on which symbol the source distribution puts more mass, the source can also be losslessly reconstructed with `polar-dec` using a checker.
3. If it is not known on which symbol the source distribution puts more mass, the source can also be losslessly reconstructed using the modified decoding algorithm `polar-dec-adapt`.

Proof of Theorem 3: Let $D(R) \subseteq \mathcal{M}(2)$ be the set of binary distributions with entropy at most R , and as before, denote by $p_0(R)$ and $p_1(R)$ the two distributions of entropy equal to R (assume $R < 1$, the result is otherwise trivial). Note that, by Corollary 1, for $i = 0, 1$

$$p_i(R) \prec_c D(R),$$

and by Lemma 3

$$\mathcal{S}_{\delta,n}(p_i(R)) \supseteq \mathcal{S}_{\delta,n}(p), \quad \forall p \in D(R), \delta, n.$$

Hence, by storing the components on $\mathcal{S}_{\delta,n}(p_0(R))$, we are not losing any information bits. We have to set $\delta = \delta_n = 2^{-n^\alpha}$ with $\alpha < 1/2$ large enough to reach the desired β in the Theorem.

Let $D_i(R)$, $i = 0, 1$, be the two regions of $D(R)$ containing distributions putting more mass on 0, respectively 1, assuming consistent indexing with $p_0(R)$ and $p_1(R)$. Note that for $i = 0, 1$

$$p_i(R) \prec_{cp} D_i(R).$$

Hence, if we know that the source distribution belongs to $D_0(R)$, we can conclude from Lemma 4 that `polar-dec`(p^* , $u[\mathcal{S}_{\delta,n}(p^*)]$, n) leads to an exact recovery, with error probability at most equal to the error probability of the source polar scheme designed with perfect knowledge of the source distribution, which is from [5], $O(2^{-n^\beta})$, for any $\beta < 1/2$. From the same paper, we conclude that the compression and reconstruction complexity $O(n \log_2 n)$.

If we do not know whether the true distribution is in $D_0(R)$ or $D_1(R)$, we can learn it as follows. Assume that the type of X^n is close to its Bernoulli distribution; this is not the case with an exponentially small probability. Notice that the observed data $U[\mathcal{S}_{\delta,n}(p_0(R))]$ corresponds to an equally likely string under both a distribution in $D_0(R)$ and $D_1(R)$, since the distribution on $\mathcal{S}_{\delta,n}(p_0(R))$ when X^n is drawn under $p_0(R)$ or $p_1(R)$ is uniform. Say w.l.o.g. that the true distribution, p_* , is in $D_0(R)$. If we use $p_0(R)$ for `polar-dec`, we will get the right output (modulo the error probability). If we use $p_1(R)$, we will recover \tilde{X}^n which is typical under \tilde{p}_* , the measure obtained by p_* by exchanging the probability mass at 0 and 1. To see this, note that for a typical X^n under p_* , $X^n + 1^n$ is typical under \tilde{p}_* (where 1^n is the n -dimensional vector filled with 1's). Moreover,

$$1^n G_n = [0^{n-1}, 1],$$

and the last component of U^n cannot be in $\mathcal{S}_{\delta,n}(p)$ (the last component of U^n is the one with least conditional entropy) unless p is the uniform distribution. Hence X^n and $X^n + 1^n$ are both typical upon observing $U[\mathcal{S}_{\delta,n}(p_0(R))]$, and we must have been able to recover correctly X^n or $X^n + 1^n$ when knowing if the true distribution was in $D_0(R)$ or $D_1(R)$ and decoding with respectively $p_0(R)$ or $p_1(R)$. Hence, by storing the value of $U[n]$ (even if it has low entropy) and running the algorithm `polar-dec-adapt` with both $p_0(R)$ and $p_1(R)$, we can check which one of the two models provides the correct estimate for $U[n]$ and learn whether p_* is in $D_0(R)$ or $D_1(R)$. Indeed, there is no need to run twice the algorithm, it is sufficient to run it once and use the value of $U[n]$ which had been stored. In any case, we will make an error, if `polar-dec` fails, which happens from [5] with probability $O(2^{-n^\beta})$, for any $\beta < 1/2$, and the complexity of this scheme is $O(n \log_2 n)$. ■

B. Results for a -ary sources

Definition 13. For $D \subset \mathcal{M}(a)$, let

$$p_c(D) := \arg \min_{p \in \mathcal{M}(a): p \prec_c D} H(p), \quad (23)$$

$$\hat{p}_c(D) := \arg \min_{p \in \hat{\mathcal{M}}(a): p \prec_c D} H(p). \quad (24)$$

In any of the above minimization, if the minimizer is not unique, pick one arbitrarily.

We clearly have that $\hat{p}_c(D) \prec_h p_c(D)$; however, it is trivial to find $\hat{p}_c(D)$ while finding $p_c(D)$ is more difficult.

Let us assume for instance that the exact source distribution is unknown for the compression part, but is known for the reconstruction part. If the source distribution is known to belong to a set $D \subset \mathcal{M}(a)$, one way to construct the storage set is to pick $\mathcal{S}(p_c(D))$. Then, from Lemma 3 and Corollary 1, this retains the information bits of any source in D . Of course, this may consume more rate than needed with an optimal source code, in other words, if we define $H_{\max}(D) := \max_{p \in D} H(p)$, we have in general

$$H(p_c(D)) \geq H_{\max}(D). \quad (25)$$

The inequality can be strict since $p_1 \prec_h p_2$ does not imply in general $p_1 \prec_c p_2$, and there are examples where equality holds in the above, in which case a compression designed for $p_c(D)$ requires the minimal rate to compress any source in D .

Remark 3. Let $D \subset \mathcal{M}(a)$ be such that $\arg \max_{p \in D} H(p)$ is unique (denoted $p_h(D)$) and satisfies $p_h(D) \prec_c D$. Then a source polar code designed for $p_h(D)$ can compress any source in D at the lowest achievable rate $\max_{p \in D} H(p)$ without losing any information bits.

(Note that $p_h(D) \prec_c D$ implies that $p_c(D) = p_h(D)$.) The set $\text{DOM}_c(p)$, plotted in Figure 1, satisfies (by definition) the condition of Remark 3 for any p . Comparing Figure 1 with the plot of $\text{DOM}_h(p) := \{q \in \mathcal{M}(a) : q \prec_h p\}$ also shows that there are sets for which (25) holds with a strict inequality. One can also check how much rate is lost by




Fig. 1. Plots of $\text{DOM}_h([0.2, 0.2, 0.6])$ (red region) included in $\text{DOM}_c([0.2, 0.4, 0.4])$ (blue region).

compressing for a distribution that is dominated in terms of \prec_c as opposed to \prec_h , for example the gap between the rate needed to compress $B_R := \{p \in \mathcal{M}(a) : H(p) \leq R\}$ using $p_c(B_R)$ and the minimal rate R needed with the method of types, i.e., $H(p_c(B_R)) - R$. This gap can be computed using Remark 2, in the case of Figure 1 for example, it is 0.095 for $R=0.865$ and $a = 3$.

Also note that $p_1 \not\prec_c p_2$ may not imply $\mathcal{S}(p_2) \not\subseteq \mathcal{S}(p_1)$, and there may be other ways to construct storage sets which contain the information bits for several distributions (than using \prec_c). We investigate this point in Section VI-E and now move to the decoding part for a -ary sources.

Definition 14. For $D \subset \mathcal{M}(a)$, let

$$p_{cp}(D) := \arg \min_{p \in \mathcal{M}(a): p \prec_c D} H(p). \quad (26)$$

If the minimizer is not unique, pick one arbitrarily.

The following follows by definitions.

Lemma 5. A source distribution known to belong to a set $D \subset \mathcal{M}(a)$ can be compressed and reconstructed losslessly at rate $H(p_{cp}(D))$, using a polar code designed for the distribution $p_{cp}(D)$.

C. Non-universality of a -ary source polar codes

In this section, we show that in general, polar codes cannot achieve the lowest rate for lossless compression of compound sources when $a \geq 3$, no matter how the storage sets are constructed (i.e., not necessarily via \prec_c). A similar result has been derived in [15] for channel coding, however, it is not possible to leverage the counter-example found in [15] to the source case (since the channel polarization results have a source counter-part only for additive noise channels, and the counter-example in [15] does not use only with additive noise channels). In this section, we assume that for the decoding part, we have the aid of a genie that provides the exact source distribution.

We consider two source distributions p and q on \mathbb{F}_a , and we are interested in finding the rates at which one can compress these two sources without losing the information bits of any of them. We denote by $C_{\text{pol}}(p, q)$ the infimum of these rates, and we provide different bounds on this quantity. Clearly

$$C(p, q) := H(p) \vee H(q) \leq C_{\text{pol}}(p, q).$$

From previous section, we have the upper bound $C_{\text{pol}}(p, q) \leq H(p_c(p, q))$, where $p_c(p, q)$ is as defined in (23) for the set $D = \{p, q\}$. In our definition, $C_{\text{pol}}(p, q)$ is given by the limit inferior of

$$\frac{1}{n} |\mathcal{S}_\delta(p) \cup \mathcal{S}_\delta(q)|.$$

Let $n = 2^\ell$, $U^n = G_n X^n$ where X^n is i.i.d. under p , and $V^n = G_n Y^n$ where Y^n is i.i.d. under q . Let us also denote by P (resp. Q) the additive noise channel whose noise distribution is p (resp. q). We then have from Section III

$$H(U_i | U^{i-1}) = 1 - I(P_i), \quad H(V_i | V^{i-1}) = 1 - I(Q_i) \quad (27)$$

where P_i (resp. Q_i) are the channels corresponding to P^σ for $\sigma \in \{-, +\}^\ell$, as defined in [3] with the tree construction. Moreover, if we define for $\delta \in (0, 1)$ $\mathcal{G}_\delta(P) = \{i \in \{1, \dots, n\} : I(P_i) \geq \delta\}$, we have

$$\mathcal{S}_\delta(p) \cup \mathcal{S}_\delta(q) = (\mathcal{G}_\delta(P) \cap \mathcal{G}_\delta(Q))^c. \quad (28)$$

This shows that the compound capacity for source or channel coding are related and we can use the result of Section III and Theorem 5 in [15] to get the following bounds.

Lemma 6.

$$C_{\text{pol}}(p, q) \leq \frac{1}{2^\ell} \sum_{\sigma \in \{-, +\}^\ell} I(\text{BEC}(Z(P^\sigma)) \vee I(Z(Q^\sigma))) \quad (29)$$

$$C_{\text{pol}}(p, q) \geq \frac{1}{2^\ell} \sum_{\sigma \in \{-, +\}^\ell} H(p^\sigma) \vee H(q^\sigma) \quad (30)$$

where P (resp. Q) is the additive noise channel with noise distribution p (resp. q). Moreover each bound is monotonically approaching $C_{\text{pol}}(p, q)$.

Note that the upper bound is straightforward, and the notation $H(p^\sigma)$ refers to $H(U_i | U^{i-1})$ for the index i corresponding to σ . It is interesting to note that if BECs can be used to compute previous bounds, we cannot use the counter-example of [15] to show that polar codes do not achieve compound capacity in source coding, since BECs do not correspond to a valid source distribution via the duality of Section III. However, we can use the duality and BECs to construct storage sets which are included in $\mathcal{S}_\delta(p) \cup \mathcal{S}_\delta(q)$, in a different manner than done in previous section. Let us give an example with $\ell = 1$. For two source distributions p and q , consider finding the BECs with parameter $Z(P)$ and $Z(Q)$ (P and Q as defined above). Then, as in [15], the good indices for P and Q satisfy

$$\mathcal{G}(P) \cap \mathcal{G}(Q) \supset \mathcal{G}(\text{BEC}(Z(P))) \cap \mathcal{G}(\text{BEC}(Z(Q))) \quad (31)$$

$$\equiv \mathcal{G}(\text{BEC}(Z(P) \vee Z(Q))) \quad (32)$$

and from (28), $\mathcal{G}(\text{BEC}(Z(P) \vee Z(Q)))$ gives a storage set to compress p and q without losing information bits. This provides an interesting and different approach to constructing universal polar codes, although it may not be practical and has the drawback of requiring the source distribution for the reconstruction (as opposed to the \prec_{cp} ordering). In a work in progress, we propose the use of spike measures $\widehat{M}(a)$ to replace the “worst BECs” directly with “worst source distributions”. The common feature between the spike measures and BECs is that they are both families that have a nested structures for the storage/good index sets and that span the whole range of entropy/mutual information between 0 and 1. Also note that as opposed to the channel polarization case, degradedness in source polarization is less restrictive, since there are less degrees of freedom for source distributions than channels.

Now, to show that polar codes do not achieve the compound capacity in source coding, we can still use the lower bound of Lemma 6, but we need to pick two source distributions on ternary source alphabets.

Proposition 1. *Polar codes do not achieve the compound capacity for source coding when the source alphabet has strictly more than 2 elements.*

Counter-example: Let $p = [0.08, 0.36, 0.56]$, $q = [0.11, 0.62, 0.27]$, such that $H(p) = 0.8143$, $H(q) = 0.8126$ and $C = H(p) \vee H(q) = 0.8143$. The LHS of Lemma 6 for $\ell = 1$ evaluates at 0.8174 which is strictly larger than C .

VI. SKETCHING AND SPARSE RECOVERY

In compressed sensing (CS), a k -sparse signal of high dimensionality n can be recovered with overwhelming probability from a small number of random measurements $m = O(k \log(n/k))$ with a convex optimization method [7], [11]. If the use of random measurement matrices simplifies the mathematical analysis, a drawback is that they have a heavy structure and it there is no efficient way to check if a given matrix realization satisfies the desired property (RIP) for the reconstruction (although one can show that this happens with high probability). Other drawbacks of random sensing matrices are discussed in [17], [6]. It has hence become a challenging problem to construct explicit matrices that, yet, can perform competitively (in terms of measurement rate) with the random ones. Different deterministic matrices have been proposed in the literature, but in [8], [22], [10] the number of rows is at least quadratic in k and in [9] one needs $\Omega(n)$ bits to specify a matrix entry. In [17], binary matrices with m brought down to $k2^{O(\log \log n)^E} = kn^{o(1)}$, with $E > 1$, are proposed and in [6], a rather general condition for constructing deterministic matrices satisfying a statistical restricted isometric property (STRIP) is given.

In this section, we are interested in designing an explicit measurement matrix using the polarization technique. The motivation being that the matrix used in previous section for polar source compression is deterministic and easily constructed. Of course, in the basic source compression problem,

the compression matrix is designed adaptively to the source distribution, whereas the CS results are universal as long as the signal is sparse. Hence, we would like to construct an explicit matrix with the polarization technique that is also universal. The tools of previous section for universal polarization will hence be used.

Note that there are a few more distinctions between CS and the problems of previous sections. First, the source in our case is a random process, whereas in the original works on compressed sensing, the signal is deterministic. The case of random signals has been considered in several subsequent works for compressed sensing, such as in [6]. Another important difference, is that the source in our setting is valued in \mathbb{F}_a , as opposed to \mathbb{R} for arbitrary sparse signals. One way to address this problem is via quantization, which requires a careful treatment. It is related to the fact that in the CS setting, measurements can be done with arbitrary precision whereas in our setting they are quantified in bits. On the other hand, we focus here on applications where the signal is valued in a discrete set to start with, such as in certain network monitoring problems [12], [13]. For example, if one wishes to track the number of packets flowing between the different IP addresses of a network (e.g. to detect unusual behaviors), the state vector can be of dimension up to 2^{32} . Since it is not feasible to maintain such a huge dimensional vector, one wishes to use a much smaller sketch vector that is still carrying all the significant information of the state vector, by exploiting the fact that the state vector is sparse. We hence keep such applications as our motivation and focus mainly on the sketching (sensing) and sparse recovery of such discrete signals. A possible lifting of the results in this paper to the real field setting is investigated in a work in progress.

Before attacking the problem of constructing a universal deterministic sketching matrix via the polarization technique, we consider a specific example by choosing a particular distribution for the signal to get started.

A. Assuming knowledge of the signal distribution

In this section, we assume that the distribution of the signal is known. The case of unknown distributions is discussed in Section VI-B. Assume that X_1, \dots, X_n are i.i.d. under the following spike distribution

$$p_\varepsilon := (1 - \varepsilon, \varepsilon/(a-1), \dots, \varepsilon/(a-1)) \in \widehat{M}_0(a). \quad (33)$$

Note that for n i.i.d. samples drawn under p_ε , the number of non-zero components is in expectation $n\varepsilon$.

Definition 15 (Polar sketching matrix for a single distribution). Let $\delta \in (0, 1)$ and $\phi_{\delta,n}(p_\varepsilon) = I_S \cdot G_n$ be the matrix obtained by deleting the rows of G_n which are not indexed by $S = \mathcal{S}_{\delta,n}(p_\varepsilon)$ (cf. Definition 8).

Rephrasing the source polarization result, we obtain the following.

Lemma 7. *Let n be a power of 2, X be an n -dimensional vector drawn i.i.d. under p_ε , and let $\phi = \phi_{\delta,n}(p_\varepsilon)$ be the*

polar sketching matrix defined for p_ε and $\delta_n = 2^{-n^\beta}$ with $\beta \in (0, 1/2)$ (cf. Definition 15). For any $\alpha \in (0, 1/2)$, there exists $\beta \in (0, 1/2)$ such that the number of rows of ϕ is given by

$$m = n(1 - \varepsilon) \log_a \left(\frac{1}{1 - \varepsilon} \right) + n\varepsilon \log_a \left(\frac{a - 1}{\varepsilon} \right) + O(2^{-n^\alpha})$$

and using the polar decoding algorithm for p_ε , we can recover X from $Y = \phi X$ with probability $O(2^{-n^\alpha})$ and with a complexity bonded as $O(a^2 n \log_2 n)$ (and if a is a power of 2, the complexity can be reduced to $O(a \log_2 a \cdot n \log_2 n)$ by using the approach in [2]).

Discussion: Note that m is simply the cardinality of \mathcal{S} , which is approximately $nH(p_\varepsilon)$. Defining $n\varepsilon = k$, we have

$$m = k \log_a \frac{n}{k} + o\left(\log_a \frac{n}{k}\right) \quad (34)$$

$$= \frac{1}{\log_e a} k \log_e \frac{n}{k} + o\left(\log_e \frac{n}{k}\right). \quad (35)$$

This expression is similar to the $O(k \log_e \frac{n}{k})$ expression encountered in the CS literature ([7]) for the number of measurements. It is even a tighter form since the constant is less than 2; of course, for the reasons discussed at the beginning of Section VI, the comparison is inappropriate, since (in particular) we are modifying the assumption on the signal: it is drawn from a specific known distribution. The reason why the number of measurements decreases when a increases may seem strange; however notice that a measurement for signals in \mathbb{F}_a is made with a precision of a bits. Hence, to compare the number of measurements for different values of a , one should use the same unit for the measurement. Let us check how the number of measurements scale with a . Rewriting (34) with the dependency in a , we have

$$m = k(1 + \log_a(a - 1)) + k \log_a \frac{n}{k} = 2k + o_a(1). \quad (36)$$

Hence, if we allow infinite precision for the measurements, for large a we only need $2k$ measurements, but of course, the complexity blows up. If we express all measurements in nats, we have

$$m = k(1 + \log_e(a - 1)) + k \log_e \frac{n}{k} \quad (37)$$

and m , as a function of a , grows like $k \log_e a$.

Of course, Lemma 7 requires knowledge of the signal distribution whereas CS results are universal. There is no reason to assume that (33) is distribution of the signal. The exact knowledge of the source distribution is in general unrealistic, and as discussed in previous section, even with an estimate of the distribution, it is crucial to show at least some robustness with respect to possible mismatched distributions. For applications, it may actually be interesting to have adaptive results, but this is also changing the rules of the game. We now investigate the universality problem.




Fig. 2. The simplex with $\text{Spa}(3, \varepsilon)$ (lower triangle, in red) for $\varepsilon \approx 1/5$ and the spike measures at 0, namely $\hat{M}_0(3)$ (middle line, in blue).

B. Universal prior

A possible way of defining k -sparse random sources, is to ask that the source distribution leads to an expected number of at most k non-zero values. Specifically, let a be a prime number and let $\mathbb{F}_a = \{0, \dots, a - 1\}$. Let $X^n = (X_1, \dots, X_n)$ be i.i.d. samples from a distribution μ , with $\mu(0) = 1 - \varepsilon$. Then, the number $K(X^n)$ of components of X^n which are not equal to 0 is in expectation

$$\mathbb{E}K(X^n) = n\varepsilon. \quad (38)$$

Let

$$\text{Spa}(a, \varepsilon) := \{\mu \in \mathcal{M}(a) : \mu(0) \geq 1 - \varepsilon\}, \quad (39)$$

and consider samples $X^n = (X_1, \dots, X_n)$ that are i.i.d. from a distribution in $\text{Spa}(a, \varepsilon)$. From previous remark, the number of components in X^n that are not equal to 0 is bounded by $n\varepsilon$. If we are interested in ε as a measure of sparsity, then $K(X^n)/n$ concentrates exponentially fast around ε . The set $\text{Spa}(a, \varepsilon)$ is pictured in Figure 3 for $a = 3$.

The results that we will derive do not depend on the fact that 0 is the special value for the distributions in $\text{Spa}(a, \varepsilon)$, in other words, we could equally well consider sources that are sparse with respect to an arbitrary $i \in \mathbb{F}_a$. For simplicity, we stick with $i = 0$ for now, although considering arbitrary i 's may be useful when dealing with the problem of quantizing a signal to \mathbb{F}_a . Also note that $\text{Spa}(a, \varepsilon)$ contains sources which can be supported on any subset of \mathbb{F}_a (e.g., a may be large but this set still contains sparse binary sources). It may be reasonable to assume that there is no such variation in the probability mass assigned to the non-special values, this will be discussed later.

Remark 4. From an information-theoretic point of view, we can ask the question of finding the smallest rate at which one could compress a source whose distribution is in $\text{Spa}(a, \varepsilon)$ without any further knowledge on the distribution (and irrespectively of the compression scheme employed). As discussed

in Section V, the answer to this question is given by the maximal entropy that can be reached with a distribution in $\text{Spa}(a, \varepsilon)$. It turns out that the distribution with maximal entropy is precisely (33), as in Section VI-A. Hence, theoretically, the strong performances presented in Section VI-A can still hold in the universal setting. However, the scheme used to achieve such Shannon limiting performance may be highly complex, whereas the whole point here, is to consider explicit schemes of low complexity.

Theorem 4. *Let X^n , with n a power of 2, be an n -sample drawn i.i.d. from a distribution which has at most ε mass on the non-zero entries of \mathbb{F}_a , and let $\phi(a, \varepsilon)$ be the $m \times n$ polar sketching matrix constructed deterministically for a and ε (cf. Definition 16). We have*

$$m = C(a, \varepsilon) \cdot k \log_e \frac{n}{k} + O(k), \quad k = n\varepsilon,$$

with

$$\lim_{\varepsilon \rightarrow 0} C(a, \varepsilon) = \frac{a-1}{\log_e a}$$

and with probability $1 - O(2^{-n^\beta})$, $\beta \in (0, 1/2)$, X^n can be exactly reconstructed from ϕX^n using the polar decoding algorithm (cf. Remark 5) with a complexity of $O(a^2 n \log_2 n)$ (or $O(a \log_2 a \cdot n \log_2 n)$ if a is a power of 2 and [2] is used).

Remark 5.

1. The polar decoding algorithm (cf. Definition 9) must be evaluated as

$$\text{polar-dec}(p_{cp}(\text{Spa}(a, \varepsilon)), \phi x^n[\mathcal{S}_{\delta, n}(p_{cp}(\text{Spa}(a, \varepsilon))), n])$$

for $\delta = \delta_n = 2^{-n^\alpha}$ with $\alpha < 1/2$ large enough to reach the desired β in the Theorem, and where $p_{cp}(\text{Spa}(a, \varepsilon))$ is the distribution of minimal entropy that is dominated by the entire set $\text{Spa}(a, \varepsilon)$ for \prec_{cp} as in (26), and \mathcal{S} as in Definition (8).

2. The multiplication ϕX^n is carried out over \mathbb{F}_a .

3. The same result holds if the distribution of X^n has at most ε mass on an arbitrary $i \in \mathbb{F}_a$.

4. In Section VI-E, we discuss improvements of the constant C .

Definition 16. Given a set D of probability measures on \mathbb{F}_a , we construct a sketching matrix $\phi(D)$ of dimension n as follows:

- (i) Find $p_{cp}(D)$ as defined in (26)
- (ii) Find $\mathcal{S} = \mathcal{S}_{\delta, n}(p_{cp}(D))$ as in Definition 8 for $0 < \delta < 1$
- (iii) Define $\phi = I_{\mathcal{S}} G_n$, where $G_n = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}^{\otimes \log_2 n}$ and where $I_{\mathcal{S}}$ is the matrix whose columns indexed by \mathcal{S} form the identity matrix and whose other columns are filled in with zeros. Note that ϕ is an $m \times n$ matrix, where $m = |\mathcal{S}|$.

In particular, we define $\phi(a, \varepsilon) := \phi(\text{Spa}(a, \varepsilon))$ and to have the optimal error decay we pick $\delta = 2^{-n^\alpha}$ with $\alpha < 1/2$.

Implementation of ϕ .

- 1. Step (i) can be easily computed, cf. Remark 2 and the proof below.
- 2. Step (ii) requires a comment: finding \mathcal{S} with an analytic

formula is a hard open problem in polar codes. However, it is mostly a mathematical challenge, since one can run simulations to determine \mathcal{S} with a good accuracy, or find arbitrarily tight bounds on the entropy terms in \mathcal{S} in polynomial time [25].

3. The construction of G_n is straightforward because of its Kronecker structure, which also allows an efficient decoding algorithm running in $O(n \log_2 n)$.

C. Interpretation of Theorem 4

In view of

$$\lim_{\varepsilon \rightarrow 0} C(a, \varepsilon) = \frac{a-1}{\log_e(a)},$$

for a fixed small ε , the quantization level a should be at most $1/\varepsilon$ in order to have a dimensionality reduction. Hence, if a signal sparser in its domain than its magnitude, where we define the magnitude-sparsity of a signal taking a possible values by $1/(a-1)$, and the domain-sparsity as before by $\varepsilon = k/n$, then the approach of Theorem 4 gives interesting results.

For a small, this is interesting for most ε . In particular for $a = 2$, the sparsity in magnitude is maximal, namely 1, and for any ε we have an optimal dimensionality reduction

$$m = 1.44 \cdot k \log_e(n/k),$$

where the optimality refers here not only to the order $k \log_e(n/k)$ but also to the constant 1.44 (recall that the measurements are taken in bits). By Shannon, one cannot further improve this bound (even with schemes of high complexity).

For a large, e.g. $a = 257$, we get reasonable dimensionality reduction for very sparse data, for example, if $\varepsilon = 10^{-3}$ and $n = 10^6$, we get a reduction of 68% for the number of measurements (compared to n). But for $a = 257$ and $\varepsilon = 0.1$, there is almost no dimensionality reduction. However, we will see in next section that this is due to the analysis employed in the proof of Theorem 4 rather than the use of the polar matrix.

D. Proof of Theorem 4

Proof: The number of measurements m is given by $nH(p_{cp}(\text{Spa}(a, \varepsilon))) + o(n)$. Note that by symmetry, $p_{cp}(\text{Spa}(a, \varepsilon))$ is a spike measure (i.e., an element of $\widehat{M}_0(a)$). We have

$$p_{cp}(\text{Spa}(a, \varepsilon)) = (1 - \eta(\varepsilon), \eta(\varepsilon)/(q-1), \dots, \eta(\varepsilon)/(q-1))$$

where $\eta(\varepsilon)$ is the smallest positive η ensuring

$$(1 - \eta, \eta/(q-1), \dots, \eta/(q-1)) \prec_{cp} p$$

for any $p \in \text{Spa}(a, \varepsilon)$. Moreover, it is sufficient to check

$$(1 - \eta, \eta/(q-1), \dots, \eta/(q-1)) \prec_{cp} (1 - \varepsilon, \varepsilon, 0, \dots, 0),$$

i.e.,

$$\mathcal{F}^{-1} \left(\frac{\mathcal{F}(1 - \eta, \eta/(q-1), \dots, \eta/(q-1))}{\mathcal{F}(1 - \varepsilon, \varepsilon, 0, \dots, 0)} \right)^{1/k} \geq 0$$

for any $k \geq 1$. Using (42), the dependence in k can also be removed. Defining

$$z = \frac{\mathcal{F}(1 - \eta, \eta/(q-1), \dots, \eta/(q-1))}{\mathcal{F}(1 - \varepsilon, \varepsilon, 0, \dots, 0)} \quad (40)$$

$$= \frac{(1, 1 - \eta \frac{q}{q-1}, \dots, 1 - \eta \frac{q}{q-1})}{(1 - \varepsilon + \varepsilon e^{-2\pi i t/a})_{t=0}^{a-1}}, \quad (41)$$

and denoting the component of z by $z_j = r_j e^{i\theta_j}$, we need to ensure

$$y(1), \dots, y(a-1) \geq 0 \quad (42)$$

where $y = \mathcal{F}^{-1}((\log_e r_j + i\theta_j)_{j=0}^{a-1})$.

Numerically, one can then easily find $\eta(\varepsilon)$ by means of the FFT algorithm. In Figure 4, we have plotted $\varepsilon \mapsto \eta(\varepsilon)$ for different values of a . Note that one can also find analytically $\eta(\varepsilon)$ using the following approach. Assume $a = 3$. Let us first find $p_c(\text{Spa}(3, \varepsilon))$. Here also, we have $p_c(\text{Spa}(3, \varepsilon)) = (1 - \bar{\eta}(\varepsilon), \bar{\eta}(\varepsilon)/2, \bar{\eta}(\varepsilon)/2)$ where we need to find $\bar{\eta}(\varepsilon)$. Note that all distributions that are worst than $(1 - \varepsilon, \varepsilon, 0)$ for \prec_c are given by the convex hull of the orbit of $(1 - \varepsilon, \varepsilon, 0)$ though cycles, that is $\text{hull}((1 - \varepsilon, \varepsilon, 0), (0, 1 - \varepsilon, \varepsilon), (\varepsilon, 0, 1 - \varepsilon))$. Hence, the projection p_{cp} of $(1 - \varepsilon, \varepsilon, 0)$, i.e., the distribution in this convex hull which belongs to the spike measures and has minimal entropy is found by taking the intersection between the line connecting $(1 - \varepsilon, \varepsilon, 0)$ to $(\varepsilon, 0, 1 - \varepsilon)$ and the line of spike measures parametrized by $(1 - d, d/2, d/2)$. An elementary computation yields

$$\bar{\eta}(\varepsilon) = 1 - 2\varepsilon(1 - \varepsilon) = 1 - 2\varepsilon + o(\varepsilon). \quad (43)$$

Note that the scaling $1 - 2\varepsilon + o(\varepsilon)$ is clear, since for small ε the line connecting $(1 - \varepsilon, \varepsilon, 0)$ to $(\varepsilon, 0, 1 - \varepsilon)$ is almost parallel to the line connecting $(1, 0, 0)$ to $(0, 0, 1)$. Indeed, one can easily generalize this for $a > 3$ to

$$\bar{\eta}_a(\varepsilon) = 1 - (a-1)\varepsilon + o(\varepsilon). \quad (44)$$

To find the projection p_{cp} of $(1 - \varepsilon, \varepsilon, 0)$, we need to move from $(1 - \varepsilon, \varepsilon, 0)$ towards spike measures with tiny convolutional steps. But once we have made a small step in the direction $(1 - \varepsilon, \varepsilon, 0) - (\varepsilon, 0, 1 - \varepsilon)$ to reach (x, y, z) , we need to move next in the rotated picture, i.e., in the direction $(x, y, z) - (y, z, x)$, as illustrated in Figure ???. Defining $f(x) = x + \gamma(\Pi - I)x$, where $\Pi = \circ(0, 1, 0)$, we are interested in $f^k(x)$ where $x = (1 - \varepsilon, \varepsilon, 0)$. Hence, we look for A^k where $A = I + \gamma(\Pi - I)$. Since Π is circulant, so is A and the eigenvector of A are the Fourier (DFT) basis elements and the eigenvalues are $1 + \gamma(\lambda_i - 1)$, where λ_i are the corresponding 3 roots of unity. Therefore, the eigenvalues of A^k are $[1 + \gamma(\lambda_i - 1)]^k$, and keeping $\gamma k = c$, we obtain

$$\lim_{k \rightarrow \infty} [1 + c/k(\lambda_i - 1)]^k = \exp(c(\lambda_i - 1)).$$

Hence,

$$\tau(c) = F_3 \text{diag}(\exp(c(\lambda_i - 1))) F_3^* (1 - \varepsilon, \varepsilon, 0)^t,$$

for $c \geq 0$, and where F_3 is the Fourier (DFT) matrix of dimension 3, parametrizes the path starting at $(1 - \varepsilon, \varepsilon, 0)$




Fig. 3. The simplex with the p_c projection (first point in red) and the p_{cp} projection (second point in blue) of $[1 - \varepsilon, 0, \varepsilon]$.

and obtained with incremental convolutional steps which are “targeting” spike measures. Equating the second and first components of τ_c , i.e., solving $\tau(c)_2 = \tau(c)_3$ gives a closed form expression for c and for $\eta(\varepsilon) = 2\tau(c)_2$, and we get as for $\bar{\eta}(\varepsilon)$,

$$\eta(\varepsilon) = 2\varepsilon + o(\varepsilon). \quad (45)$$

That is, for small ε , the penalty endured by considering the p_{cp} projection rather than the p_c one (for Spa) is negligible. This is not surprising, since for ε small, the path from $(1 - \varepsilon, \varepsilon, 0)$ to spike measures is anyway small (as required for the p_{cp} projection). With similar arguments, we conclude that for any a ,

$$\eta_a(\varepsilon) = (a-1)\varepsilon + o(\varepsilon). \quad (46)$$

Finally, we need to evaluate $H(p_{\eta(\varepsilon)})$ where

$$p_{\eta(\varepsilon)} = (1 - \eta(\varepsilon), \eta(\varepsilon)/(a-1), \dots, \eta(\varepsilon)/(a-1)).$$

Note we can compare the cost for universality of a low complexity scheme (obtained with the p_{cp} analysis and the polar matrix) with respect to the limiting performance (Shannon): instead of $nH(p_\varepsilon)$ we need $nH(p_{(a-1)\varepsilon})$ measurements, when ε is small. For $a = 2$, these two are identical, and this is consistent with Theorem 3. For arbitrary a , we get

$$H(p_{\eta(\varepsilon)}) = \frac{a-1}{\log_e a} \varepsilon \log_e \frac{1}{\varepsilon} + O(\varepsilon) \quad (47)$$

as opposed to

$$H(p_\varepsilon) = \frac{1}{\log_e a} \varepsilon \log_e \frac{1}{\varepsilon} + O(\varepsilon). \quad (48)$$

■




Fig. 4. Plots of $\varepsilon \mapsto \eta(\varepsilon)$ for different values of $a = q$. Note that $(a-1)\varepsilon$ is an upper bound to $\eta(\varepsilon)$, tight for small ε as shown in the proof of Theorem 4. Hence for ε not too small $\eta(\varepsilon)$ provides a better measurement rate than what is obtained with the crude bound of Theorem 4.




Fig. 5. Plots of $\varepsilon \mapsto \eta^*(\varepsilon)$ for different values of $a = q$

E. Improving Theorem 4

In Remark 4, we concluded that the minimal number of measurements (in a -ary bits) needed to recover a source from $\text{Spa}(a, \varepsilon)$ is given by $nH(p_\varepsilon)$ where p_ε is the spike measure with mass $1 - \varepsilon$ at 0 and $\varepsilon/(a-1)$ elsewhere. To approach this performance with a polarization scheme of low complexity is challenging. We have shown that, for the universal problem of recovering sources from any distribution in $\text{Spa}(a, \varepsilon)$, one could still use the adaptive setting but designed for a specific distribution; namely $p_{cp}(\text{Spa}(a, \varepsilon))$. This distribution is dominated by the entire set $\text{Spa}(a, \varepsilon)$ with respect to \prec_{cp} , and we showed that because of this property, a scheme designed for $p_{cp}(\text{Spa}(a, \varepsilon))$ guarantees successful recovery for any distributions in $\text{Spa}(a, \varepsilon)$. In a sense, $p_{cp}(\text{Spa}(a, \varepsilon))$ is the worst case scenario. Of course, there may be other ways (than using \prec_{cp}) to order distributions and find a worst case distribution which has lower entropy (one may attempt to replace \prec_c with \prec_d). There may even be other ways of tackling the universal problem than looking for an ordering and a worst case distribution. The advantage

of using a worst case approach, is that we can then inherit the complexity attributes and convergence rate property from the adaptive setting. Ordering also allows to give a ‘hierarchy’ between the different distributions, and helps designing robust schemes (by backing off from the estimated distribution to guarantee performances).

One point is that it may not be needed to consider the entire set $\text{Spa}(a, \varepsilon)$ for a given problem. For example, the set of distributions that are dominated by $p_{cp}(\text{Spa}(a, \varepsilon))$ (w.r. to \prec_{cp}), is already quite large and contains most distributions of sparsity ε . It does not contain the distributions of sparsity ε which have small supports, but if these can be ruled out, then we can bring back the constant $C(a, \varepsilon)$ to $1/\log_e(a)$, which is the Shannon limiting performance.

We now discuss another approach to construct a universal sketching and reconstruction method for $\text{Spa}(a, \varepsilon)$. A before, there are two parts to discuss. First the sketching, i.e., to know which rows of G_n can be deleted without losing information about X^n . Then the reconstruction, i.e., to know how to run a decoding algorithm that ignores the exact distribution.

1) *Sketching*: Here is a brut-force approach to achieve universal sketching.

Definition 17 (brut-univ-sketching algorithm).

Inputs: ε (the sparsity degree), a (the size of the signal alphabet), n (the dimension).

Outputs: $\eta^*(\varepsilon)$.

We present two variants of the algorithm.

Variant A:

For η from ε until $\eta(\varepsilon)$ (with a given step size);

if $\mathcal{S}(p_\eta) \supseteq \mathcal{S}(q)$ for any q in the convex hull of $(1 - \varepsilon, \varepsilon, 0, \dots, 0)$, $(1 - \varepsilon, 0, \varepsilon, 0, \dots, 0)$, ..., $(1 - \varepsilon, 0, \dots, 0, \varepsilon)$;

output η ;

otherwise increase the step size.

Variant B:

For η from ε until $\eta(\varepsilon)$ (with a given step size);

if $\mathcal{S}(p_\eta) \supseteq \mathcal{S}(q^{(\varepsilon)})$ where $q^{(\varepsilon)} = (1 - \varepsilon, \varepsilon, 0, \dots, 0)$;

output η ;

otherwise increase the step size.

Note: one could also consider a dichotomic approach for the search of $\eta^*(\varepsilon)$, and by symmetry, one can restrict the search of q to only one portion of the convex hull. One also has to specify the precision ξ for the computations of the sets $\mathcal{S} = \mathcal{S}_{\xi, n}$, we omitted it in the algorithm to simplify the notation.

Variant B has low complexity, since it conducts a search in a one-dimensional space (for η) and since the computation of \mathcal{S} can be done at low computational costs. Variant A requires a larger search for q , which can be constraining for a large a .

Result: Variant A of brut-univ-sketching provides $\eta^*(\varepsilon)$ such that

$$\mathcal{S}(p_{\eta^*(\varepsilon)}) \supseteq \mathcal{S}(p), \quad \forall p \in \text{Spa}(a, \varepsilon).$$

Conjecture: Variant B of brut-univ-sketching leads to the same output than Variant A.

In Figure 5, we show $\eta^*(\varepsilon)$ (obtained with Variant B of `brut-univ-sketching`) and the Shannon limit consisting of the diagonal, and reached for $a = 2$. As observed, the improvement is significant with respect to $\eta(\varepsilon)$ (obtained with the p_{cp} projection). Indeed, this brings the number of measurement very close to the optimal performance. Emre Telatar is gratefully acknowledged for his help in producing these plots.

For the decoding part, there are no guarantees that decoding with $p_{\eta^*(\varepsilon)}$ allows a correct recovery. One can use the algorithm `polar-dec-adapt` to learn the distribution, but one needs to first add checkers in the set of stored components. Checkers are components that need not to be stored (because they have low conditional entropy) but that we still store to help the decoder get information about the source distributions. As long as the number of checkers is $o(n)$, the measurement rate is not affected.

2) *Reconstruction*: We now proceed to use `polar-dec-adapt` to decoder the sensed components of previous part. We first proceed to a patching of $\text{Spa}(a, \varepsilon)$. Consider a uniform discretization of the convex hull of $(1 - \varepsilon, \varepsilon, 0, \dots, 0), (1 - \varepsilon, 0, \varepsilon, 0, \dots, 0), \dots, (1 - \varepsilon, 0, \dots, 0, \varepsilon)$. Enumerate a uniform discretization of this convex hull as $D_k := \{p_1, \dots, p_d\}$. Call $\text{Spa}_d(a, \varepsilon)$ the sets of distributions that dominates any of the elements in D_k with respect to \prec_{cp} . We then have

$$\text{Spa}_d(a, \varepsilon) \rightarrow \text{Spa}(a, \varepsilon),$$

meaning that the set $\text{Spa}_d(a, \varepsilon)$ is dense in $\text{Spa}(a, \varepsilon)$. For a targeted ε , we then pick a ε' slightly larger than ε and a d large enough such that

$$\text{Spa}_d(a, \varepsilon') \supseteq \text{Spa}(a, \varepsilon).$$

We then use `polar-dec-adapt` with the output of `brut-univ-sketching` and $o(n)$ checkers to learn which of the distributions p_k is a good ‘model’ for the sensed data. The term model is used because by construction of $\text{Spa}_d(a, \varepsilon')$, the sensed data on the checkers must look typical with at least one of the p_i ’s, although there might be more than one, and although none of these p_i ’s may be the true distribution of X^n (but they will be dominated with respect to \prec_{cp} by the true distribution, which is good enough to ensure correct decoding). One has to pick $\varepsilon' - \varepsilon$ small enough and d large enough to ensure a small increase in the number of measurements (one needs to study the scaling of these parameters for a more precise statement). The overall complexity of this decoding scheme scales multiplicatively with d . Hence, as long as d is not of the order of $n \log_2 n$, the overall complexity remains low.

In a work in progress, we also consider another approach for universal decoding via an algebraic characterization of the possible likelihood ratios computed in `polar-dec`.

A. Universal polar coding

We summarize here two ideas introduced in this paper to construct universal polar coding schemes

- 1) Convolutional path ordering: to tell when a polar coding scheme designed for one distribution can succeed for another one
- 2) Checkers: to learn some information about the distribution by storing components that did not need to be stored

In particular, we developed an algorithm which allows to compress universally binary sources at the lowest achievable rate, with low complexity and with guaranteed low error probability.

B. Sparse recovery and sketching

We applied the tools developed for universal polar coding to the problem of sketching sparse signals, constructing a deterministic sketching matrix by deleting appropriate rows of the polar matrix G_n . We summarize here some conclusions and extensions on this approach.

- 1) *An sketching method tuned to discrete signals.*

Compressed sensing exploits the sparsity of signals in their *domain* to acquire them efficiently. If for the application of interest, the signal is also sparse in its *magnitude*, that is, if it takes values in a set of small cardinality, this can also be exploited as shown in this paper. For example, if the signal is binary, we developed a sketching method with a deterministic low complexity matrix, an optimal number of measurements (for the scaling and the constant) and a low complexity recovery algorithm with a proved exponentially small (in \sqrt{n}) probability of error. We extended this results to a -ary vectors, noticing a better fit for small a and proposing an improved approach for larger a (Section VI-E). We also underline that, for a given application, the method proposed here can be used adaptively by designing an appropriate probabilistic model for the signal. This can improve the measurement rate.

- 2) *Lifting this work to the reals?*

Most works in the CS literature constructing explicit sensing matrices are based on algebraic constructions [17], [10], [6]. In these works, matrix acting on the reals can then be obtained. Of course, we also made the point (previous item) that for certain application, it may be more natural to work with the discrete setting directly. Yet, an interesting extension would be to study a lifting of our results to the real case. A possible approach would be via a quantization procedure, where problems of robustness to noise must be investigated. Another possible problem would be to attempt detecting the signal support only (which is a binary signal).

REFERENCES

- [1] E. Abbe, *Universal source polarization and sparse recovery*, Information theory workshop, Dublin, August 2010.
- [2] E. Abbe and E. Telatar, *Polar codes for the m-user MAC*, in Proc. 2010 International Zurich Seminar on Communications, Zurich, 2010. Available at arXiv:1002.0777v2 [cs.IT].
- [3] E. Arıkan, *Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels*, IEEE Trans. Inform. Theory, vol. IT-55, pp. 3051–3073, July 2009.
- [4] E. Arıkan and E. Telatar, *On the rate of channel polarization*, in Proc. 2009 IEEE Int. Symp. Inform. Theory, Seoul, pp. 1493–1495, 2009.
- [5] E. Arıkan, *Source polarization*, in Proc. IEEE Int. Symp. Inform. Theory, Austin, 2010.
- [6] R. Calderbank, S. Howard and S. Jafarpour, *Construction of a Large Class of Deterministic Sensing Matrices That Satisfy a Statistical Isometry Property* IEEE Journal of Selected Topics in Signal Processing, vol. 4, no. 2, pp. 358 - 374 , April 2010.
- [7] E. Candes and T. Tao, *Near-Optimal Signal Recovery From Random Projections: Universal Encoding Strategies*, IEEE Trans. Inform. Theory, vol. 52, no. 12, pp. 5406-5425, December 2006.
- [8] G. Cormode and S. Muthukrishnan, *Combinatorial algorithms for compressed sensing*, SIROCCO, 2006
- [9] R. DeVore, *Optimal computation*, Proceedings of the International Congress of Math- ematicians, 2006.
- [10] R. DeVore, *Deterministic constructions of compressed sensing matrices*, Journal of Complexity, Volume 23, Issue 4-6, August 2007.
- [11] D. Donoho, *Compressed Sensing*, IEEE Trans. Inform. Theory, vol. 52, no. 4, pp. 1289-1306, April 2006.
- [12] C. Estan and G. Varghese, *New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice*, ACM Transactions on Computer Systems, 2003.
- [13] A. Gilbert, Y. Kotidis, S. Muthukrishnan, and M. Strauss, *Quicksand: Quick summary and analysis of network data*, DIMACS Technical Report 2001-43, 2001.
- [14] J. Haupt, W.U. Bajwa, M. Rabbat, and R. Nowak, *Compressed sensing for networked data*, IEEE Signal Processing Mag., vol. 25, no. 2, pp. 92-101, March 2008.
- [15] S. H. Hassani, S. B. Korada, R. Urbanke, *The Compound Capacity of Polar Codes*, arXiv:0907.3291v1 [cs.IT], July 2009.
- [16] E. Hof, S. Shamai, *Secrecy-Achieving Polar-Coding for Binary-Input Memoryless Symmetric Wire-Tap Channels*, available at arXiv:1005.2759v2 [cs.IT], May 2010.
- [17] P. Indyk, *Explicit constructions for compressed sensing of sparse signals*, Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms, 2008.
- [18] S. B. Korada, R. Urbanke, *Polar Codes are Optimal for Lossy Source Coding*, Information Theory Workshop, Taormina, October 2009.
- [19] O. O. Koyluoglu, H. El Gamal, *Polar Coding for Secure Transmission and Key Agreement*, available at arXiv:1003.1422v1 [cs.IT], March 2010.
- [20] H. Mahdavifar, A. Vardy, *Achieving the Secrecy Capacity of Wiretap Channels Using Polar Codes*, IEEE International Symposium on Information Theory, Austin, June 2010.
- [21] A.W. Marshall and I. Olkin, *Inequalities: Theory of Majorization and Its Applications*, Academic Press, 1980.
- [22] S. Muthukrishnan, *Compressed sensing algorithms for functions*, Allerton Conference on Communication, Control and Computing, 2006.
- [23] E. Şaşıođlu, E. Telatar, E. Arıkan, *Polarization for arbitrary discrete memoryless channels*, Information Theory Workshop, Taormina, October 2009.
- [24] E. Şaşıođlu, E. Telatar, E. Yeh, *Quasi-polarization for the two user binary input multiple access channel*, IEEE Information Theory Workshop, Cairo, January 2010. Available at arXiv:1006.4255v1 [cs.IT].
- [25] I. Tal, A. Vardy, *How to Construct Polar Codes*, Information theory workshop, Dublin, August 2010.