

High-Girth Matrices and Polarization

Emmanuel Abbe

Prog. in Applied and Computational Math. and EE Dept.

Princeton University

Email: eabbe@princeton.edu

Yuval Wigderson

Department of Mathematics

Princeton University

Email: yuvalw@princeton.edu

Abstract—The girth of a matrix is the least number of linearly dependent columns, in contrast to the rank which is the largest number of linearly independent columns. This paper considers the construction of *high-girth* matrices, whose probabilistic girth is close to their rank. Random matrices can be used to show the existence of high-girth matrices. This paper uses a recursive construction based on conditional ranks (inspired by polar codes) to obtain a deterministic and efficient construction of high-girth matrices for arbitrary relative ranks. Interestingly, the construction is agnostic to the underlying field and applies to both finite and continuous fields with the same binary matrix. The construction gives in particular the following: (i) over the binary field, high-girth matrices are equivalent to capacity-achieving codes, and our construction turns out to match exactly the BEC polar codes (even at finite block length). It hence gives a different interpretation of BEC polar codes, using the parity-check matrix instead of the generator matrix, and basic linear algebra instead of the mutual information, and generalizes to larger fields; (ii) for the BSC, our construction gives an operational meaning to the Bhattacharyya upper-bound process used in polar codes; (iii) for the reals, it gives an explicit candidate matrix for sparse recovery.

I. INTRODUCTION

Let A be a matrix over a field \mathbb{F} . Assume that A is flat, i.e., it has more columns than rows. The rank of A , denoted by $\text{rank}(A)$, is the maximal number of linearly independent columns. The girth of A , denoted by $\text{girth}(A)$, is the least number of linearly dependent columns. What are the possible tradeoffs between $\text{rank}(A)$ and $\text{girth}(A)$? This depends on the cardinality of the field. It is clear that

$$\text{girth}(A) \leq \text{rank}(A) + 1. \quad (1)$$

Is it possible to have a *perfect-girth matrix* that achieves this upper-bound? If $\mathbb{F} = \mathbb{R}$, drawing the matrix with i.i.d. standard Gaussian entries gives such an example with probability 1. However, if $\mathbb{F} = \mathbb{F}_q$, where q is finite, the problem is different. For $\mathbb{F} = \mathbb{F}_q$, note that

$$\text{girth}(A) = \text{dist}(C_A), \quad (2)$$

where $\text{dist}(C_A)$ is the distance of the q -ary linear code C_A whose parity check matrix is A . In fact, the least number of columns that are linearly dependent in A is equal to the least number of columns whose linear combination can be made 0, which is equal to the least weight of a vector that is mapped to 0 by A , which is the least weight of a codeword, i.e., the code distance since the code is linear.

Hence, over finite fields, the girth is a key parameter for error-correcting codes, and studying the girth/rank tradeoffs for matrices is equivalent to studying the distance/dimension tradeoffs for linear codes. Clearly it is not possible to obtain perfect-girth matrices over $\mathbb{F} = \mathbb{F}_2$, even if we relax the perfect-girth

requirement to be asymptotic, requiring $\text{rank}(A) \sim \text{girth}(A)$ when the number of columns in A tends to infinity.¹ If $\mathbb{F} = \mathbb{F}_2$, the Gilbert-Varshamov bound provides a lower-bound on the maximal girth (conjectured to be tight by some). Namely, for a uniformly drawn matrix A with n columns, with high probability,

$$\text{rank}(A) = nH(\text{girth}(A)/n) + o(n), \quad (3)$$

where H is the binary entropy function.

For $\mathbb{F} = \mathbb{F}_q$, the bound in (1) is a restatement of the Singleton bound for linear codes and expressed in terms of the co-dimension of the code. Asking for a perfect-girth matrix is hence equivalent to asking for an MDS linear code. Such constructions are known when $q = n$ with Reed-Solomon codes. Note that the interest on MDS codes has recently resurged with the applications in distributed data storage, see [5] for a survey.

One may consider instead the case of non-finite fields, typically not covered in coding theory. As shown in Section IV-B, this is relevant for the recovery of sparse signals [6] via compressed measurements. The girth is then sometimes given different names, such as the Kruskal-rank or spark [6]. As stated above, for $\mathbb{F} = \mathbb{R}$, a random Gaussian matrix is perfect-girth with probability one. However, computing the girth of an arbitrary matrix is NP-hard [11] (like computing the distance of a code [12]), making the latter construction non-explicit.

In this paper, we are mainly interested in the following notion of *probabilistic girth*, defined to be the least number of columns that are linearly dependent with high probability, when drawing the columns uniformly at random. Formal definitions are given in the next section. Going from a worst-case to a probabilistic model naturally allows for much better bounds. In particular, defining *high-girth matrices* as matrices whose probabilistic girth and rank are of the same order (up to $o(n)$), a random uniform matrix proves the existence of high-girth matrices even for $\mathbb{F} = \mathbb{F}_2$. However, obtaining an explicit construction is again non-trivial.

In this paper, we obtain explicit and fully deterministic constructions of high-girth matrices over all fields and for any relative ranks. We rely on a polar-code-like construction. Starting with the same squared matrix as for polar or Reed-Muller codes, i.e., the tensor-product/Sierpinski matrix, we then select rows with a different measure based on ranks. For finite fields, we show that high-girth matrices are equivalent to capacity-achieving linear codes for erasure channels, while for errors the speed of convergence of the probabilistic girth requirement matters. In particular, we achieve the Bhattacharyya bound with our explicit construction. For the real field, this

¹We use the notation $a_n \sim b_n$ for $\lim_{n \rightarrow \infty} a_n/b_n = 1$.

allows to construct explicit binary measurement matrices with optimal probabilistic girth.

These results have various other implications. First, our construction gives an operational interpretation to the upper-bound of the Bhattacharyya-process in polar codes. When the channel is not the BEC, the upper-bound of this process used in the polar code literature is in fact the conditional rank process studied in this paper. Note that the fact that this correspondence takes place at any finite block length is not obvious, even with the fact that high-girth matrices are equivalent to capacity-achieving parity-check matrices over the BEC, since the latter only needs to hold asymptotically. Second, this paper gives a high-school level proof (not necessarily trivial but relying only basic linear algebra concepts) of a fully deterministic, efficient, and capacity-achieving code for erasure channels. While capacity-achieving codes for the BEC are well-known by now, most constructions rely still on rather sophisticated tools (e.g. expander codes, polar codes, LDPC codes, spatially-coupled codes), and we felt that an explicit construction relying only on the notions of rank and girth is rather interesting. On the other hand, for $\mathbb{F} = \mathbb{F}_2$, our construction turns out to be equivalent to the polar code for the BEC, so that the difference is mainly about the approach. Nonetheless, it allows to simplify the concepts, not requiring even the notion of mutual information. Finally, the result generalizes to non-binary alphabets, given that our construction does depend on the underlying field.

II. HIGH-GIRTH MATRICES

A. Notation

Let A be a $m \times n$ matrix over a field \mathbb{F} . For any set $S \subseteq [n]$, let $A[S]$ be the submatrix of A obtained by selecting the columns of A indexed by S . For $s \in [0, 1]$, let $A[s]$ be a random submatrix of A obtained by sampling each column independently with probability s . Thus, $A[s] = A[\tilde{S}]$, where \tilde{S} is an i.i.d. $\text{Ber}(s)$ random subset of $[n]$. In expectation, $A[s]$ has sn columns. Throughout the paper, an event E_n takes place with high probability if $\mathbb{P}\{E_n\} \rightarrow 1$ when $n \rightarrow \infty$, where n should be clear from the context.

B. Probabilistic Girth

Definition 1. Let $\{A_n\}$ be a sequence of matrices over a field \mathbb{F} , where A_n has n columns. The *probabilistic girth* of A_n is the supremum of all $s \in [0, 1]$ such that $A_n[s]$ has linearly independent columns with high probability, i.e.,

$$\begin{aligned} \text{girth}_*(\{A_n\}) &:= \\ \sup\{s \in [0, 1] : \mathbb{P}\{A_n[s] \text{ has lin. indep. cols.}\} &= 1 - o(1)\} \end{aligned} \quad (4)$$

Note that a better name might have been the probabilistic *relative* girth, since it is a counterpart of the usual notion of girth in the probabilistic setting with an additional normalization factor of n . We often write $\text{girth}_*(A_n)$ instead of $\text{girth}_*(\{A_n\})$. We will sometimes care about how fast the above probability tends to 1; we then say that A_n has a probabilistic girth with rate $\tau(n)$ if the above definition holds when

$$\mathbb{P}\{A_n[s] \text{ has lin. indep. columns}\} = 1 - \tau(n). \quad (5)$$

Definition 2. We say that A_n is *high-girth* if

$$\text{girth}_*(A_n) = \limsup_{n \rightarrow \infty} \text{rank}(A_n)/n. \quad (6)$$

For $\mu \in [0, 1]$, we say that A_n is μ -high-girth if it is high-girth and $\text{girth}_*(A_n) = \mu$.

Example 1. Consider the following construction, corresponding to Reed-Solomon codes. Let x_1, \dots, x_n be distinct elements of a field \mathbb{F} , and consider the $m \times n$ matrix

$$V = \begin{pmatrix} 1 & 1 & 1 & \dots & 1 \\ x_1 & x_2 & x_3 & \dots & x_n \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ x_1^{m-1} & x_2^{m-1} & x_3^{m-1} & \dots & x_n^{m-1} \end{pmatrix} \quad (7)$$

Then V will satisfy a stronger property than being high-girth, as its actual girth is $m + 1$: every $m \times m$ submatrix will be invertible, since every $m \times m$ submatrix is a Vandermonde matrix whose determinant must be nonzero. However, this example cannot be used to construct high-girth families over a fixed finite field \mathbb{F} . For as soon as n exceeds the size of \mathbb{F} , it will be impossible to pick distinct x_i 's, and we will no longer have high girth.

Example 2. A $\mu n \times n$ uniform random matrix with entries in \mathbb{F}_2 is μ -high-girth with high probability. For in this case, any $\mu n \times \mu n$ submatrix will be a uniform random square matrix, which has full rank with high probability.

III. EXPLICIT CONSTRUCTION OF HIGH-GIRTH MATRICES

A. Sierpinski matrices

Let \mathbb{F} be any field, let n be a power of 2, and let G_n be the matrix over \mathbb{F} defined by

$$G_n = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^{\otimes \log n} \quad (8)$$

Note that the entries of this matrix are only 0's and 1's, hence this can be viewed as a matrix over any field.

Many important codes can be derived from G_n , and they are all based on a simple idea. Namely, we first pick some measure of "goodness" on the rows of G_n . Then, we take the submatrix of G_n obtained by keeping only those rows which are the "best" under this metric, and we finally define a code whose PCM is this matrix. The first important examples are Reed-Muller (RM) codes [8], [9], where goodness is measured by the weight of the rows, and more recently polar codes [2], [3], where goodness is measured by the entropy (or mutual information). In the next section, we define a measure of goodness based on ranks and use it to construct high-girth matrices. A similar construction was proposed in [7] for Hadamard matrices to polarize the Rényi information dimension. We discuss applications to coding and sparse recovery in the next sections.

B. Conditional-rank matrices

With $s \in [0, 1]$ fixed, let $G_n^{(i)}$ denote the submatrix of G_n obtained by taking the first i rows, and let $G_n^{(i)}[s]$ be the random submatrix obtained by sampling each column independently with probability s , as above.

Definition 3. The conditional rank (COR) of row i in G_n is defined by

$$\rho(n, i, s) = \mathbb{E}(\text{rank}_{\mathbb{F}} G_n^{(i)}[s]) - \mathbb{E}(\text{rank}_{\mathbb{F}} G_n^{(i-1)}[s]) \quad (9)$$

where $\text{rank}_{\mathbb{F}}$ denotes the rank computed over the field \mathbb{F} . When $i = 1$, define $\rho(n, i, s) = \mathbb{E}(\text{rank}_{\mathbb{F}} G_n^{(1)}[s])$.

Now, by adding the i th row, we will either keep the rank constant or increase it by 1, and the latter will happen if and

only if the i th row is independent of the previous rows. So we get that

$$\rho(n, i, s) = \mathbb{P}(\text{the } i\text{th row of } G_n[s] \text{ is independent of the previous } i-1 \text{ rows}), \quad (10)$$

where linear independence is also considered over \mathbb{F} . The key property of the conditional ranks is expressed in the following lemma.

Lemma 1. *Define the functions*

$$\ell(x) = 2x - x^2 \quad r(x) = x^2 \quad (11)$$

and define a branching process of depth $\log n$ and offspring 2 (i.e., each node has exactly two descendants) as follows: the base node has value s , and for a node with value x , its left-hand child has value $\ell(x)$ and its right-hand child has value $r(x)$. Then the n leaf-nodes of this branching process are, in order, the values $\rho(n, i, s)$ for $1 \leq i \leq n$.

An important point about this lemma is that the functions ℓ and r do not depend on \mathbb{F} , while the $\rho(n, i, s)$ values do, a priori. Thus, one way to interpret this lemma is that the expected conditional ranks of G_n do not depend on the field \mathbb{F} , even though their definition does. The proof of Lemma 1 is given in Section V.

A key property of the branching process in Lemma 1 is that it is a balanced process, meaning that the average value of the two children of a node with value x is x again:

$$\frac{\ell(x) + r(x)}{2} = \frac{(2x - x^2) + x^2}{2} = x \quad (12)$$

This means that this branching process defines a martingale, by letting a random walk go left or right with probability half. Moreover, since $\rho(n, i, s)$ is a probability, we have that this martingale stays in $[0, 1]$. So by Doob's martingale convergence theorem, we must have that this martingale converges almost surely to its fixed points. In fact, Doob's theorem is not needed here, as one may conclude using only the fact that the increments are orthogonal.² Its fixed points are those x 's for which $\ell(x) = r(x) = x$. The only points satisfying this are 0 and 1, so this martingale polarizes. In fact, much can be said about the speed of polarization of this process, as it is equivalent to the polarization process for BEC channels studied in [4].

Theorem 1 (Application of [4]). *For any n ,*

$$\frac{|\{i \in [n] : \rho(n, i, s) > 1 - 2^{-n^{0.49}}\}|}{n} = s + o(1) \quad (13)$$

$$\frac{|\{i \in [n] : \rho(n, i, s) < 2^{-n^{0.49}}\}|}{n} = (1 - s) + o(1) \quad (14)$$

Hence the theorem tells us is that the above martingale polarizes very quickly: apart from a vanishing fraction, all $\rho(n, i, s)$'s are exponentially close to 0 or 1 as $n \rightarrow \infty$. With this in mind, we define the following.

Definition 4. Let n be a fixed power of 2, and let $s \in [0, 1]$ be fixed. Let $H \subset [n]$ be the set of indices i for which $\rho(n, i, s) > 1 - 2^{-n^{0.49}}$, and let $m = |H|$. By Theorem 1, we know that $m = sn + o(n)$. Let R_n denote the $m \times n$ submatrix of G_n gotten by selecting all the columns of G_n , but only taking

those rows indexed by H . We call R_n the COR matrix of size n with parameter s .

Note that the construction of COR matrices is trivial, like the BEC polar codes, whereas the construction of BSC polar codes requires an efficient approximation algorithm [10].

We will index the rows of R_n by $i \in H$, rather than $j \in [m]$. We sometimes denote R_n by R . The most important property of R_n is expressed in the following theorem.

Theorem 2. *For any $s \in [0, 1]$, $R_n[s]$ has full rank (i.e. rank m) with high probability, as $n \rightarrow \infty$. In fact, $R_n[s]$ has full rank with probability $1 - o(2^{-n^{0.49}})$.*

The proof is a simple consequence of Lemma 1 and Theorem 1, and can be found in the Appendix.

Theorem 2 implies the following.

Theorem 3. *For any $s \in [0, 1]$, R_n is s -high-girth.*

Since the proof of Theorem 2 works independently of the base field \mathbb{F} , the same is true of Theorem 3. Thus, the COR construction is a fully deterministic and works over any field. In fact, it requires only two values (0 and 1) for the matrix, even when $\mathbb{F} = \mathbb{R}$.

IV. APPLICATIONS OF HIGH-GIRTH MATRICES

A. Coding for erasures

Let \mathbb{F} a field and $p \in [0, 1]$. The memoryless erasure channel on \mathbb{F} with erasure probability p , denoted by $MEC(p)$, erases each component of a codeword on \mathbb{F} independently with probability p . Denoting by ε the erasure symbol, the output alphabet is hence $\mathbb{F}_* = \mathbb{F} \cup \{\varepsilon\}$ and the transition probability of receiving $y \in \mathbb{F}_*$ when $x \in \mathbb{F}$ is transmitted is

$$W(y|x) = \begin{cases} p & \text{if } y = \varepsilon, \\ 1 - p & \text{if } y = x. \end{cases} \quad (15)$$

The memoryless extension is defined by $W^n(y^n|x^n) = \prod_{i=1}^n W(y_i|x_i)$ for $x^n \in \mathbb{F}^n, y^n \in \mathbb{F}_*^n$.

Recall that a code of block length n and dimension k over the alphabet \mathbb{F} is a subset of \mathbb{F}^n of cardinality $|\mathbb{F}|^k$. The code is linear if the subset is a subspace of dimension k . In particular, a linear code can be expressed as the image of a generator matrix $G \in \mathbb{F}^{n \times k}$ or as the null space of a parity-check matrix $H \in \mathbb{F}^{(n-k) \times n}$. The rate of a code is defined by k/n . A rate R is achievable over the $MEC(p)$ if the code can correct the erasures with high probability. More specifically, R is achievable if there exists a sequence of codes C_n of blocklength n and dimension k_n having rate R , and decoders $D_n : \mathbb{F}_*^n \rightarrow \mathbb{F}^n$, such that $P_e(C_n) \rightarrow 0$, where for x^n drawn uniformly at random in C_n and y^n the output of x^n over the $MEC(p)$, and

$$P_e(C_n) := \mathbb{P}\{D(y^n) \neq x^n\}. \quad (16)$$

The dependency in D_n is not explicitly stated in P_e as there is no degree of freedom to decode over the MEC (besides guessing the erasure symbol), as shown in the proof of next lemma.

The supremum of the achievable rates is the capacity, given by $1 - p$. We now relate capacity-achieving codes on the MEC and high-girth matrices.

Lemma 2. *A linear code C_n achieves a rate R on the $MEC(p)$ if and only if its parity check matrix has probabilistic girth at least $1 - R$. In particular, a code achieves capacity on the $MEC(p)$ if and only if its parity check matrix is p -high-girth.*

²Private discussion with E. Telatar. See also [1]

In particular, the linear code whose parity-check matrix is a COR matrix of parameter p achieves capacity on the $MEC(p)$.

Remark 1. In the binary case, COR codes give a new interpretation to BEC polar codes: instead of computing the mutual information of the polarized channels via the generator matrix, we can interpret BEC polar codes from the girth of the parity-check matrix. Note that this simplifies the proof that BEC polar codes achieve capacity to a high-school linear algebra — mutual information need not be even introduced. The only part which may not be of a high-school level is the martingale argument, which is in fact not necessary, as already known in the polar code literature (see for example [1, Homework 4], which basic algebra).

Remark 2. As shown in [2], the action of the matrix $G_n = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^{\otimes \log n}$ on a vector can be computed in $O(n \log n)$ time as well, which means that the encoding of the COR code can be done in $O(n \log n)$ time, and the same is true for the code construction (which is not the case for general polar codes). Decoding the COR code can be done by inverting the submatrix of R_n corresponding to the indices that do not have erasure symbols, which can be done by Gaussian elimination in $O(n^3)$ time. Alternatively, COR codes can be decoded as polar codes, i.e., successively, in $O(n \log(n))$ time. Hence, like polar codes for the BEC, COR codes are deterministic, capacity-achieving, and efficiently encodable and decodable for the MEC.

B. Sparse recovery

In the setting of sparse recovery, one wishes to recover a real-valued sparse signal from a lower-dimensional projection [6]. In the worst-case model, a k -sparse signal is a vector with at most k non-zero components, and to recover x that is k -sparse from Ax , it must be that $Ax \neq Ax'$ from all x, x' that are k -sparse (and different). Hence A needs³ to have girth $2k + 1$.

One may instead consider a probabilistic model where a k -sparse signal has a random support, drawn uniformly at random or from an i.i.d. model where each component in $[n]$ belongs to the support with probability $p = k/n$. The goal is then to construct a flat matrix A which allows to recover k -sparse signals with high probability on the drawing of the support. Note that a bad support S is one which is k -sparse and that can be paired with another k -sparse support S' such that that there exists real-valued vectors x, x' supported respectively on S, S' which have the same image through A , i.e.,

$$Ax = Ax' \iff A(x - x') = 0. \quad (17)$$

Note now that this is equivalent to saying that the columns of A indexed by $S \cup S'$ are linearly dependent, since $x - x'$ is supported on $S \cup S'$ which is $2k$ -sparse.

Hence, the probability of error for sparse recovery is given by

$$\mathbb{P}_S\{\exists S' : A[S \cup S'] \text{ has lin. dep. columns}\}. \quad (18)$$

This error probability can be upper-bounded as for errors (see the next section), by estimating the probability that A has a subset of up to $2k$ linearly dependent columns, which relies on the high-girth property of A .

³Note that for noise stability or to obtain a convex relaxation of the decoder, one needs the columns to have in addition singular values close to 1, i.e., the restricted isometry property (RIP).

C. Coding for errors

In this section, we work over the binary field \mathbb{F}_2 . The binary symmetric channel with error probability p , denoted by $BSC(p)$, flips each bit independently with probability p . More formally, the transition probability of receiving $y \in \mathbb{F}_2$ when $x \in \mathbb{F}_2$ is transmitted is given by

$$W(y|x) = \begin{cases} p & \text{if } y \neq x \\ 1 - p & \text{if } y = x \end{cases} \quad (19)$$

The memoryless extension is then defined by $W^n(y^n|x^n) = \prod_{i=1}^n W(y_i|x_i)$, for $x^n, y^n \in \mathbb{F}_2^n$.

Theorem 4. Let $p \in [0, 1/2]$ and $s = s(p) = 2\sqrt{p(1-p)}$ be the Bhattacharyya parameter of the $BSC(p)$. Let $\{C_n\}$ be the COR code with parameter s (the code whose PCM is R_n). Then C_n can reliably communicate over the $BSC(p)$ with high probability, as $n \rightarrow \infty$.

Note that unlike in the erasure scenario, Theorem 4 does not allow us to achieve capacity over the $BSC(p)$. For the capacity of the $BSC(p)$ is $1 - H(p)$, and

$$1 - H(p) \geq 1 - 2\sqrt{p(1-p)} \quad (20)$$

with equality holding only for $p \in \{0, \frac{1}{2}, 1\}$.

This statement, unlike the ones for erasure correction and sparse recovery, was stated only for the COR code, and not for general high-girth codes. Our proof of this statement, which can be found in the Appendix, relies on the actual construction of COR codes and on the upper-bound on the successive probability of error in terms of the COR known from polar codes. One may also attempt to obtain this result solely from the high-girth property, but this requires further dependencies on the high-girth rate of convergence. It is an interesting problem to obtain achievable rates that solely depend on the probabilistic girth for the BSC.

V. SOME PROOFS

Proof of Lemma 1: We induct on $\log n$. The base case is $\log n = 1$, where the calculation is straightforward. The rank of $G_2^{(1)}[s]$ will be 0 if no columns are chosen, and will be 1 if at least 1 column is chosen. Therefore,

$$\rho(2, 1, s) = \mathbb{E}(\text{rank}_{\mathbb{F}}(G_2^{(1)}[s])) \quad (21)$$

$$= 0 \cdot (1 - s)^2 + 1 \cdot 2s(1 - s) + 1 \cdot s^2 \quad (22)$$

$$= 2s - s^2 = \ell(s) \quad (23)$$

Similarly, $\mathbb{E}(\text{rank}_{\mathbb{F}}(G_2^{(2)}[s])) = 2s$, and thus

$$\rho(2, 2, s) = (2s) - (2s - s^2) = s^2 = r(x) \quad (24)$$

Note that all these calculations do not actually depend on \mathbb{F} .

For the inductive step, assume that $\rho(n/2, i, s)$ is the leaf value of the branching process for all $1 \leq i \leq n/2$. To prove the same for $\rho(n, i, s)$, write $G_n = G_{n/2} \otimes G_2$. In other words, we think of G_n as being an $(n/2) \times (n/2)$ matrix whose entries are 2×2 matrices.

We begin with the case when i is odd. By the inductive hypothesis, we wish to prove that

$$\rho(n, i, s) = \ell\left(\rho\left(\frac{n}{2}, \frac{i+1}{2}, s\right)\right) \quad (25)$$

We partition the columns of $G_n^{(i)}[s]$ into two sets: O , which consists of those columns which have an odd index in G_n , and

E , which consists of those with an even index in G_n . Since $G_n = G_{n/2} \otimes G_2$, we see that for i odd, the i th row of G_n is the $((i+1)/2)$ th row of $G_{n/2}$, except that each entry is repeated twice. From this, and from inclusion-exclusion, we see that

$$\rho(n, i, s) = \mathbb{P}(\text{row } i \text{ of } G_n \text{ is independent} \\ \text{of the previous rows}) \quad (26)$$

$$= \mathbb{P}(\text{row } i \text{ of } G_n[O] \text{ is independent} \\ \text{of the previous rows of } G_n[O]) \\ + \mathbb{P}(\text{row } i \text{ of } G_n[E] \text{ is independent} \\ \text{of the previous rows of } G_n[E]) \\ - \mathbb{P}(\text{both of the above}) \quad (27)$$

$$= 2\rho\left(\frac{n}{2}, \frac{i+1}{2}, s\right) - \rho\left(\frac{n}{2}, \frac{i+1}{2}, s\right)^2 \quad (28)$$

$$= \ell\left(\rho\left(\frac{n}{2}, \frac{i+1}{2}, s\right)\right) \quad (29)$$

Next, we consider the case when i is even. In this case, we wish to prove that

$$\rho(n, i, s) = r\left(\rho\left(\frac{n}{2}, \frac{i}{2}, s\right)\right) \quad (30)$$

We proceed analogously. From the equation $G_n = G_{n/2} \otimes G_2$, we see that the i th row of G_n is the $(i/2)$ th row of $G_{n/2}$, except with a 0 interspersed between every two entries. Thus, the i th row will be *dependent* if either it restricted to O or it restricted to E will be dependent; in other words, it'll be independent if and only if both the restriction to O and the restriction to E are independent. Therefore,

$$\rho(n, i, s) = \mathbb{P}(\text{the restriction to } O \text{ and} \\ \text{the restriction to } E \text{ are independent}) \quad (31)$$

$$= \mathbb{P}(\text{the restriction to } O \text{ is independent}) \cdot \\ \mathbb{P}(\text{the restriction to } E \text{ is independent}) \quad (32)$$

$$= \left(\rho\left(\frac{n}{2}, \frac{i}{2}, s\right)\right)^2 = r\left(\rho\left(\frac{n}{2}, \frac{i}{2}, s\right)\right) \quad (33)$$

Note that in all our calculations, we used probability arguments that are valid over any field. Broadly speaking, this works because the above arguments show that the only sorts of linear dependence that can be found in $G_n^{(i)}[s]$ involves coefficients in $\{-1, 0, 1\}$. Since these elements are found in any field, we have that this theorem is true for all fields \mathbb{F} . ■

Proof of Lemma 2: Note that a decoder over the MEC needs to correct the erasures, but there is no bias towards which symbol can have been erased. Hence, a decoder on the MEC is wrong with probability at least half if there are multiple codewords that match the corrupted word. In other words, the probability of error is given by⁴

$$P_e(C_n) = \mathbb{P}_E\{\exists x, y \in C_n, x \neq y, x[E^c] = y[E^c]\} \quad (34)$$

$$= \mathbb{P}_E\{\exists z \in C_n, z \neq 0, z[E^c] = 0\} \quad (35)$$

where E is the erasure pattern of the $MEC(p)$, i.e., a random subset of $[n]$ obtained by picking each element with probability p . Let H_n be the parity-check matrix of C_n , i.e., $C_n = \ker(H_n)$.

Note that E has the property that there exists a codeword $z \in C_n$ such that $z[E^c] = 0$ if and only if the columns indexed by E in H_n are linearly dependent. Indeed, assume first that there exists such a codeword z , where the support of z is contained in E . Since z is in the kernel of H_n , the columns of H_n indexed by the support of z must add up to 0, hence any set of columns that contains the support of z must be linearly dependent. Conversely, if the columns of H_n indexed by E are linearly dependent, then there exists a subset of these columns and a collection of coefficients in \mathbb{F} such that this linear combination is 0, which defines the support of a codeword z . Hence,

$$P_e(C_n) = \mathbb{P}_E\{H_n[E] \text{ has lin. dependent columns}\}. \quad (36)$$

Recalling that the code rate is given by $1 - r$, where r is the relative rank of the parity-check matrix, the conclusions follow. ■

ACKNOWLEDGEMENTS

This work was partially supported by NSF grant CIF-1706648 and the Princeton PACM Director's Fellowship.

REFERENCES

- [1] Emmanuel Abbe. Lecture notes for the course APC 529: Coding Theory and Random Graphs. Available on request, 2013.
- [2] Erdal Arkan. Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels. *Information Theory, IEEE Transactions on*, 55(7):3051–3073, 2009.
- [3] Erdal Arkan. Source polarization. *arXiv preprint arXiv:1001.3087*, 2010.
- [4] Erdal Arkan and Emre Telatar. On the rate of channel polarization. In *Information Theory, 2009. ISIT 2009. IEEE International Symposium on*, pages 1493–1495. IEEE, 2009.
- [5] Alexandros G Dimakis, Kannan Ramchandran, Yunnan Wu, and Changho Suh. A survey on network codes for distributed storage. *Proceedings of the IEEE*, 99(3):476–489, 2011.
- [6] David L Donoho and Michael Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- [7] S. Haghghatshoar and E. Abbe. Polarization of the Rényi information dimension for single and multi terminal analog compression. In *Information Theory Proceedings (ISIT), 2013 IEEE International Symposium on*, pages 779–783, July 2013.
- [8] D.E. Muller. Application of Boolean algebra to switching circuit design and to error detection. *Electronic Computers, Transactions of the I.R.E. Professional Group on*, EC-3(3):6–12, Sept 1954.
- [9] I.S. Reed. A class of multiple-error-correcting codes and the decoding scheme. *Information Theory, Transactions of the IRE Professional Group on*, 4(4):38–49, 1954.
- [10] I. Tal and A. Vardy. How to construct polar codes. *Information Theory, IEEE Transactions on*, 59(10):6562–6582, Oct 2013.
- [11] A.M. Tillmann and M.E. Pfetsch. The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing. *Information Theory, IEEE Transactions on*, 60(2):1248–1259, Feb 2014.
- [12] A. Vardy. The intractability of computing the minimum distance of a code. *Information Theory, IEEE Transactions on*, 43(6):1757–1766, Nov 1997.

⁴If ties are broken at random, an additional factor of $1 - 1/|\mathbb{F}|$ should appear on the right hand side.