

Efficient Encoding of Low-Density Parity-Check Codes

Thomas J. Richardson and Rüdiger L. Urbanke

Abstract—Low-density parity-check (LDPC) codes can be considered serious competitors to turbo codes in terms of performance and complexity and they are based on a similar philosophy: constrained random code ensembles and iterative decoding algorithms.

In this paper, we consider the encoding problem for LDPC codes. More generally, we consider the *encoding* problem for codes specified by sparse parity-check matrices. We show how to exploit the sparseness of the parity-check matrix to obtain efficient encoders. For the $(3, 6)$ -regular LDPC code, for example, the complexity of encoding is essentially quadratic in the block length. However, we show that the associated coefficient can be made quite small, so that encoding codes even of length $n \simeq 100\,000$ is still quite practical. More importantly, we will show that “optimized” codes actually admit linear time encoding.

Index Terms—Binary erasure channel, decoding, encoding, parity check, random graphs, sparse matrices, turbo codes.

I. INTRODUCTION

LOW-DENSITY parity-check (LDPC) codes were originally invented and investigated by Gallager [1]. The crucial innovation was Gallager’s introduction of iterative decoding algorithms (or message-passing decoders) which he showed to be capable of achieving a significant fraction of channel capacity at low complexity. Except for the papers by Zyablov and Pinsker [2], Margulis [3], and Tanner [4] the field then lay dormant for the next 30 years. Interest in LDPC codes was rekindled in the wake of the discovery of turbo codes and LDPC codes were independently rediscovered by both MacKay and Neal [5] and Wiberg [6].¹ The past few years have brought many new developments in this area. First, in several papers Luby, Mitzenmacher, Shokrollahi, Spielman, and Stemann introduced new tools for the investigation of message-passing decoders for the binary-erasure channel (BEC) and the binary-symmetric channel (BSC) (under hard-decision message-passing decoding) [9], [10], and they extended Gallager’s definition of LDPC codes to include *irregular* codes (see also [5]). The same authors also exhibited sequences of codes which, asymptotically in the block length, provably achieve

capacity on a BEC. It was then shown in [11] that similar analytic tools can be used to study the asymptotic behavior of a very broad class of message-passing algorithms for a wide class of channels and it was demonstrated in [12] that LDPC codes can come extremely close to capacity on many channels.

In many ways, LDPC codes can be considered serious competitors to turbo codes. In particular, LDPC codes exhibit an asymptotically better performance than turbo codes and they admit a wide range of tradeoffs between performance and decoding complexity. One major criticism concerning LDPC codes has been their apparent high *encoding* complexity. Whereas turbo codes can be encoded in linear time, a straightforward encoder implementation for an LDPC code has complexity quadratic in the block length. Several authors have addressed this issue.

- 1) It was suggested in [13] and [9] to use cascaded rather than bipartite graphs. By choosing the number of stages and the relative size of each stage carefully one can construct codes which are encodable and decodable in linear time. One drawback of this approach lies in the fact that each stage (which acts like a subcode) has a length which is, in general, considerably smaller than the length of the overall code. This results, in general, in a performance loss compared to a standard LDPC code with the same overall length.
- 2) In [14] it was suggested to force the parity-check matrix to have (almost) lower triangular form, i.e., the ensemble of codes is restricted not only by the degree constraints but also by the constraint that the parity-check matrix have lower triangular shape. This restriction guarantees a linear time encoding complexity but, in general, it also results in some loss of performance.

It is the aim of this paper to show that, even without cascade constructions or restrictions on the shape of the parity-check matrix, the encoding complexity is quite manageable in most cases and provably linear in many cases. More precisely, for a $(3, 6)$ -regular code of length n the encoding complexity seems indeed to be of order n^2 but the actual number of operations required is no more than $0.017^2 n^2 + O(n)$, and, because of the extremely small constant factor, even large block lengths admit practically feasible encoders. We will also show that “optimized” irregular codes have a *linear* encoding complexity and that the required amount of preprocessing is of order at most $n^{3/2}$.

The proof of these facts is achieved in several stages. We first show in Section II that the encoding complexity is upper-bounded by $n + g^2$, where g , the *gap*, measures in some way to be made precise shortly, the “distance” of the given parity-check matrix to a lower triangular matrix. In Section III, we then discuss several greedy algorithms to triangulate matrices and we

Manuscript received December 15, 1999; revised October 10, 2000. This work was performed while both authors were at Bell Labs, Lucent Technologies, Murray Hill, NJ 07974 USA.

T. J. Richardson was with Bell Labs, Lucent Technologies, Murray Hill, NJ 07974 USA. He is now with Flarion Technologies, Bedminster, NJ 07921 USA (e-mail: richardson@flarion.com).

R. L. Urbanke was with Bell Labs, Lucent Technologies, Murray Hill, NJ 07974 USA. He is now with EPFL, LTHC-DSC, CH-1015 Lausanne, Switzerland (e-mail: rudiger.urbanke@epfl.ch).

Communicated by D. A. Spielman, Guest Editor.
Publisher Item Identifier S 0018-9448(01)00739-8.

¹Similar concepts have also appeared in the physics literature [7], [8].

show that for these algorithms, when applied to elements of a given ensemble, the gap concentrates around its expected value with high probability. As mentioned above, for the $(3, 6)$ -regular code the best greedy algorithm which we discuss results in an expected gap of $0.017n$. Finally, in Section IV, we prove that for all known “optimized” codes the expected gap is actually of order less than $\sqrt{n}$, resulting in the promised linear encoding complexity. In practice, the gap is usually a small constant. The $\sqrt{n}$ bound can be improved but it would require a significantly more complex presentation.

We finish this section with a brief review of some basic notation and properties concerning LDPC codes. For a more thorough discussion we refer the reader to [1], [11], [12].

LDPC codes are linear codes. Hence, they can be expressed as the null space of a *parity-check matrix* H , i.e., x is a codeword if and only if

$$Hx^T = 0^T.$$

The modifier “low-density” applies to H ; the matrix H should be sparse. For example, if H has dimension $\frac{n}{2} \times n$, where n is even, then we might require H to have three 1’s per column and six 1’s per row. Conditioned on these constraints, we choose H at random as discussed in more detail below. We refer to the associated code as a $(3, 6)$ -regular LDPC code. The sparseness of H enables efficient (suboptimal) decoding, while the randomness ensures (in the probabilistic sense) a good code [1].

Example 1. [Parity-Check Matrix of $(3, 6)$ -Regular Code of Length 12]: The following matrix H will serve as an example.

$$H = \begin{pmatrix} 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 \\ 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 1 & 0 \end{pmatrix}. \quad (1)$$

□

In the theory of LDPC codes it is customary and useful not to focus on particular codes but to consider ensembles of codes. These ensembles are usually defined in terms of ensembles of *bipartite graphs* [13], [15]. For example, the bipartite graph which represents the code defined in Example 1 is shown in Fig. 1. The *left* set of nodes represents the *variables* whereas the *right* set of nodes represents the *constraints*. An ensemble of bipartite graphs is defined in terms of a pair of *degree distributions*. A degree distribution $\gamma(x) = \sum_i \gamma_i x^{i-1}$ is simply a polynomial with nonnegative real coefficients satisfying $\gamma(1) = 1$. Typically, γ_i denotes the fraction of edges in a graph which are incident to a node (variable or constraint node as the case may be) of degree i . In the sequel, we will use the shorthand $\int \gamma$ to denote

$$\sum_{i \geq 1} \frac{\gamma_i}{i} = \int_0^1 \gamma(x) dx.$$

This quantity gives the inverse of the average node degree. Associated to a degree distribution pair (λ, ρ) is the *rate* $r(\lambda, \rho)$ defined as

$$r(\lambda, \rho) := 1 - \frac{\int \rho}{\int \lambda}. \quad (2)$$




Fig. 1. Graphical representation of a $(3, 6)$ -regular LDPC code of length 12. The left nodes represent the variable nodes whereas the right nodes represent the check nodes.

For example, for the degree distribution pair (x^2, x^5) , which corresponds to the $(3, 6)$ -regular LDPC code, the rate is $\frac{1}{2}$.

Given a pair (λ, ρ) of degree distributions and a natural number n , we define an *ensemble* of bipartite graphs $\mathcal{C}^n(\lambda, \rho)$ in the following way. All graphs in the ensemble $\mathcal{C}^n(\lambda, \rho)$ will have *left* nodes which are associated to λ and *right* nodes which are associated to ρ . More precisely, assume that

$$\lambda(x) := \sum_{i \geq 1} \lambda_i x^{i-1}$$

and

$$\rho(x) := \sum_{i \geq 1} \rho_i x^{i-1}.$$

We can convert these degree distributions into *node perspective* by defining

$$\tilde{\lambda}_i := \frac{\lambda_i}{i \int \lambda} \quad \text{and} \quad \tilde{\rho}_i := \frac{\rho_i}{i \int \rho}.$$

Each graph in $\mathcal{C}^n(\lambda, \rho)$ has $n\tilde{\lambda}_i$ left nodes of degree i and $(1 - r(\lambda, \rho))n\tilde{\rho}_i$ right nodes of degree i . The order of these nodes is arbitrary but fixed. Here, to simplify notation, we assume that (λ, ρ) and n are chosen in such a way that all these quantities are integer. A node of degree i has i *sockets* from which the i edges emanate and these sockets are *ordered*. Thus, in total there are

$$\begin{aligned} s &:= \sum_{i \geq 1} i n \tilde{\lambda}_i = \sum_{i \geq 1} n \frac{\lambda_i}{\int \lambda} = \frac{n}{\int \lambda} = \frac{(1 - r(\lambda, \rho))n}{\int \rho} \\ &= (1 - r(\lambda, \rho)) \sum_{i \geq 1} n \frac{\rho_i}{\int \rho} = (1 - r(\lambda, \rho)) \sum_{i \geq 1} i n \tilde{\rho}_i \end{aligned}$$

ordered sockets on the left as well as on the right. Let σ be a permutation on $[s] := \{1, \dots, s\}$. We can associate a graph to such a permutation by connecting the i th socket on the left to the $\sigma(i)$ th socket on the right. Letting σ run over the set of permutations on $[s]$ generates a set of graphs. Endowed with the uniform probability distribution this is the ensemble $\mathcal{C}^n(\lambda, \rho)$. Therefore, if in the future we choose a graph at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$ then the underlying probability distribution is the uniform one.

It remains to associate a code to every element of $\mathcal{C}^n(\lambda, \rho)$. We will do so by associating a parity-check matrix to each graph. At first glance, it seems natural to define the parity-check matrix associated to a given element in $\mathcal{C}^n(\lambda, \rho)$ as that $\{0, 1\}$ -matrix which has a nonzero entry at row i and column j if and only if (iff) the i th right node is connected to the j th left node. Unfortunately, the possible presence of multiple edges between pairs of nodes requires a more careful definition. Since the encoding is done over the field $\text{GF}(2)$, we define the parity-check matrix H as the $\{0, 1\}$ matrix which has a nonzero entry at row i and column j iff the i th right node is connected to the j th left node an *odd* number of times. As we will see, the encoding is accomplished in two steps, a *preprocessing* step, which is an offline calculation performed once only for the given code, and the actual encoding step which is the only data-dependent part. For the preprocessing step it is more natural to work with matrices which contain the multiplicities of edges and, therefore, we define the *extended parity-check matrix* $\hat{H}$ as that matrix which has an entry d at row i and column j iff the i th right node is connected to the j th left node by d edges. Clearly, H is equal to $\hat{H}$ modulo 2. In the sequel, we will also refer to these two matrices as the *adjacency matrix* and the *extended adjacency matrix* of the bipartite graph. Since for every graph there is an associated code, we will use these two terms interchangeably so we will, e.g., refer to codes as elements of $\mathcal{C}^n(\lambda, \rho)$.

Most often, LDPC codes are used in conjunction with *message-passing decoders*. Recall that there is a received message associated to each variable node which is the result of passing the corresponding bit of the codeword through the given channel. The decoding algorithm proceeds in *rounds*. At each round, a message is sent from each variable node to each neighboring check node, indicating some estimate of the associated bit's value. In turn, each check node collects its incoming messages and, based on this information, sends messages back to the incident variable nodes. Care must be taken to send out only *extrinsic* information, i.e., the outgoing message along a given edge must not depend on the incoming message along the same edge. As we will see, the preprocessing step for the encoding is closely related to the message-passing decoder for the BEC. We will therefore review this particular decoder in more detail.

Assume we are given a code in $\mathcal{C}^n(\lambda, \rho)$ and assume that we use this code to transmit over a BEC with an erasure probability of α . Therefore, an expected fraction α of the variable nodes will be *erasures* and the remaining fraction $(1 - \alpha)$ will be *known*. We first formulate the iterative decoder not as a message-passing decoder but in a language which is more suitable for our current purpose, see [9].

Decoder for the Binary Erasure Channel:

0. [Initialization]
1. [Stop or Extend] If there is no known variable node and no check node of degree one then output the (partial) codeword and stop. Otherwise, all known variable nodes and all their adjacent edges are deleted.
2. [Declare Variables as Known] Any variable node which is connected to a degree one check node is declared to be known. Goto 1.

This decoder can equivalently be formulated as a message-passing decoder. Messages are from the set $\{0, 1\}$ with a 0 indicating that the corresponding bit has not been determined yet (along the given edge). We will call a 0 message an *erasure message*. At a variable node, the outgoing message along an edge e is the erasure message if the received message associated to this node is an erasure and if all incoming messages (excluding the incoming message along edge e) are erasure messages, otherwise, the outgoing message is a 1. At a check node, the outgoing message along an edge e is the erasure message if at least one of the incoming messages (excluding the incoming message along edge e) is the erasure message, and a 1 otherwise. If we declare that an originally erased variable node becomes *known* as soon as it has at least one incoming message which is not an erasure then one can check that at any time the set of known variable nodes is indeed identical under both descriptions.

It was shown in [16] that (asymptotically in n) the expected fraction α_ℓ of erasure messages after the ℓ th decoding round is given by

$$\alpha_\ell = \alpha\lambda(1 - \rho(1 - \alpha_{\ell-1})) \quad (3)$$

where $\alpha_0 = \alpha$. Let $\alpha^*(\lambda, \rho)$, called the *threshold* of the degree distribution pair, be defined as

$$\alpha^*(\lambda, \rho) := \sup \left\{ 0 \leq \alpha \leq 1 : \alpha_\ell(\alpha) \xrightarrow{\ell \rightarrow \infty} 0 \text{ where } \alpha_\ell(\alpha) := \alpha\lambda(1 - \rho(1 - \alpha_{\ell-1})); \alpha_0 = \alpha \right\}. \quad (4)$$

Note first that the function $f(x, y) := y\lambda(1 - \rho(1 - x))$ is increasing in both its arguments for $x, y \in [0, 1]$. It follows by finite induction that if $\alpha_\ell(\alpha) \xrightarrow{\ell \rightarrow \infty} 0$ then $\alpha_\ell(\alpha') \xrightarrow{\ell \rightarrow \infty} 0$ for any $\alpha' \leq \alpha$. If we choose $\alpha < \alpha^*(\lambda, \rho)$, then the asymptotic expected fraction of erasure messages converges to zero. Consequently, the decoder will be successful with high probability in this case. If, on the other hand, we choose $\alpha > \alpha^*(\lambda, \rho)$ then, with high probability, the decoding process will not succeed. We will see shortly that, correctly interpreted, this decoding procedure constitutes the basis for all preprocessing algorithms that we consider in this paper.

Example 2. [(3, 6)-Regular Code]: Let

$$(\lambda(x), \rho(x)) = (x^2, x^5).$$

Then $r(x^2, x^5) = \frac{1}{2}$. The exact threshold $\alpha^*(x^2, x^5)$ was determined in [17] and can be expressed as follows. Let σ be given by

$$\sigma = \frac{1}{36} + \frac{\sqrt{\frac{25}{324} - a + b}}{2} + \frac{\sqrt{\frac{25}{162} + a - b + \frac{685}{2916\sqrt{\frac{25}{324} - a + b}}}}{2}$$

where

$$a = \frac{22}{27} 5^{\frac{2}{3}} \left(\frac{2}{-85 + 3\sqrt{24465}} \right)^{\frac{1}{3}}$$

and

$$b = \frac{1}{27} \left(\frac{5}{2} \left(-85 + 3\sqrt{24465} \right) \right)^{\frac{1}{3}}.$$

Then

$$\alpha^*(x^2, x^5) := \frac{1 - \sigma}{(1 - \sigma^5)^2} \sim 0.42944$$

□ **Ex. 1022.3**




Fig. 2. An equivalent parity-check matrix in lower triangular form.




Fig. 3. The parity-check matrix in approximate lower triangular form.

II. EFFICIENT ENCODERS BASED ON APPROXIMATE LOWER TRIANGULATIONS

In this section, we shall develop an algorithm for constructing efficient encoders for LDPC codes. The efficiency of the encoder arises from the sparseness of the parity-check matrix H and the algorithm can be applied to any (sparse) H . Although our example is binary, the algorithm applies generally to matrices H whose entries belong to a field F . We assume throughout that the rows of H are linearly independent. If the rows are linearly dependent, then the algorithm which constructs the encoder will detect the dependency and either one can choose a different matrix H or one can eliminate the redundant rows from H in the encoding process.

Assume we are given an $m \times n$ parity-check matrix H over F . By definition, the associated code consists of the set of n -tuples x over F such that

$$Hx^T = 0^T.$$

Probably the most straightforward way of constructing an encoder for such a code is the following. By means of Gaussian elimination bring H into an equivalent lower triangular form as shown in Fig. 2. Split the vector x into a systematic part s , $s \in F^{n-m}$, and a parity part p , $p \in F^m$, such that $x = (s, p)$. Construct a systematic encoder as follows: i) Fill s with the $(n - m)$ desired information symbols. ii) Determine the m parity-check symbols using *back-substitution*. More precisely, for $l \in [m]$ calculate

$$p_l = \sum_{j=1}^{n-m} H_{l,j} s_j + \sum_{j=1}^{l-1} H_{l,j+n-m} p_j.$$

What is the complexity of such an encoding scheme? Bringing the matrix H into the desired form requires $O(n^3)$ operations of preprocessing. The actual encoding then requires $O(n^2)$ operations since, in general, after the preprocessing the matrix will no longer be sparse. More precisely, we expect that we need about $n^2 \frac{r(1-r)}{2}$ XOR operations to accomplish this encoding, where r is the rate of the code.

Given that the original parity-check matrix H is sparse, one might wonder if encoding can be accomplished in $O(n)$. As we will show, typically for codes which allow transmission at rates close to capacity, linear time encoding is indeed possible. And for those codes for which our encoding scheme still leads to quadratic encoding complexity the constant factor in front of the

n^2 term is typically very small so that the encoding complexity stays manageable up to very large block lengths.

Our proposed encoder is motivated by the above example. Assume that by performing row and column permutations only we can bring the parity-check matrix into the form indicated in Fig. 3. We say that H is in *approximate lower triangular form*. Note that since this transformation was accomplished solely by permutations, the matrix is still sparse. More precisely, assume that we bring the matrix in the form

$$H = \begin{pmatrix} A & B & T \\ C & D & E \end{pmatrix} \quad (5)$$

where A is $(m-g) \times (n-m)$, B is $(m-g) \times g$, T is $(m-g) \times (m-g)$, C is $g \times (n-m)$, D is $g \times g$, and, finally, E is $g \times (m-g)$. Further, all these matrices are sparse² and T is lower triangular with ones along the diagonal. Multiplying this matrix from the left by

$$\begin{pmatrix} I & 0 \\ -ET^{-1} & I \end{pmatrix} \quad (6)$$

we get

$$\begin{pmatrix} A & B & T \\ -ET^{-1}A + C & -ET^{-1}B + D & 0 \end{pmatrix}. \quad (7)$$

Let $x = (s, p_1, p_2)$ where s denotes the systematic part, p_1 and p_2 combined denote the parity part, p_1 has length g , and p_2 has length $(m - g)$. The defining equation $Hx^T = 0^T$ splits naturally into two equations, namely

$$As^T + Bp_1^T + Tp_2^T = 0 \quad (8)$$

and

$$(-ET^{-1}A + C)s^T + (-ET^{-1}B + D)p_1^T = 0. \quad (9)$$

Define $\phi := -ET^{-1}B + D$ and assume for the moment that ϕ is nonsingular. We will discuss the general case shortly. Then from (9) we conclude that

$$p_1^T = -\phi^{-1}(-ET^{-1}A + C)s^T.$$

Hence, once the $g \times (n - m)$ matrix $-\phi^{-1}(-ET^{-1}A + C)$ has been precomputed, the determination of p_1 can be accomplished in complexity $O(g \times (n - m))$ simply by performing

²More precisely, each matrix contains at most $O(n)$ elements.

TABLE I
EFFICIENT COMPUTATION OF $p_1^T = -\phi^{-1}(-ET^{-1}A + C)s^T$

Operation	Comment	Complexity
As^T	Multiplication by sparse matrix	$O(n)$
$T^{-1}[As^T]$	$T^{-1}[As^T] = y^T \Leftrightarrow [As^T] = Ty^T$	$O(n)$
$-E[T^{-1}As^T]$	Multiplication by sparse matrix	$O(n)$
Cs^T	Multiplication by sparse matrix	$O(n)$
$[-ET^{-1}As^T] + [Cs^T]$	Addition	$O(n)$
$-\phi^{-1}[-ET^{-1}As^T + Cs^T]$	Multiplication by dense $g \times g$ matrix	$O(g^2)$

TABLE II
EFFICIENT COMPUTATION OF $p_2^T = -T^{-1}(As^T + Bp_1^T)$

Operation	Comment	Complexity
As^T	Multiplication by sparse matrix	$O(n)$
Bp_1^T	Multiplication by sparse matrix	$O(n)$
$[As^T] + [Bp_1^T]$	Addition	$O(n)$
$-T^{-1}[As^T + Bp_1^T]$	$-T^{-1}[As^T + Bp_1^T] = y^T \Leftrightarrow -[As^T + Bp_1^T] = Ty^T$	$O(n)$

a multiplication with this (generically dense) matrix. This complexity can be further reduced as shown in Table I. Rather than precomputing $-\phi^{-1}(-ET^{-1}A + C)$ and then multiplying with s^T we can determine p_1 by breaking the computation into several smaller steps, each of which is efficiently computable.

To this end, we first determine As^T , which has complexity $O(n)$ since A is sparse. Next, we multiply the result by T^{-1} . Since $T^{-1}[As^T] = y^T$ is equivalent to the system $[As^T] = Ty^T$ this can also be accomplished in $O(n)$ by back-substitution, since T is lower triangular and also sparse. The remaining steps are fairly straightforward. It follows that the overall complexity of determining p_1 is $O(n + g^2)$. In a similar manner, noting from (8) that $p_2^T = -T^{-1}(As^T + Bp_1^T)$, we can accomplish the determination of p_2 in complexity $O(n)$ as shown step by step in Table II.

A summary of the proposed encoding procedure is given in Table III. It entails two steps. A *preprocessing* step and the actual *encoding* step. In the preprocessing step, we first perform row and column permutations to bring the parity-check matrix into approximate lower triangular form with as small a gap g as possible. We will see, in subsequent sections, how this can be accomplished efficiently. We also need to check whether $\phi := -ET^{-1}B + D$ is nonsingular. Rather than premultiplying by the matrix $\begin{pmatrix} I & 0 \\ -ET^{-1} & I \end{pmatrix}$, this task can be accomplished efficiently by Gaussian elimination. If, after clearing the matrix E the resulting matrix ϕ is seen to be singular we can simply perform further column permutations to remove this singularity. This is always possible when H is not rank deficient, as assumed. The actual encoding then entails the steps listed in Tables I and II.

We will now demonstrate this procedure by means of our running example.

Example 3. [Parity Check Matrix of (3,6)-Regular Code of Length 12]: For this example if we simply reorder the columns such that, according to the original order, we have the ordering 1, 2, 3, 4, 5, 6, 7, 10, 11, 12, 8, 9, then we put the parity-check

matrix into an approximate lower triangular form with $g = 2$

$$\left(\begin{array}{c|c|c} A & B & T \\ \hline C & D & E \end{array} \right) = \left(\begin{array}{cccccc|cc|cccc} 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 \\ \hline 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 \end{array} \right). \quad (10)$$

We now use Gaussian elimination to clear E . This results in

$$\left(\begin{array}{cccccc|cc|cccc} 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 1 \\ \hline 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \end{array} \right).$$

We see that $\phi := -ET^{-1}B + D = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ is singular. This singularity can be removed if we exchange e.g., column 5 with column 8 which gives $\phi = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$. In terms of the original order the final column order is then 1, 2, 3, 4, 10, 6, 7, 5, 11, 12, 8, 9, and the resulting equivalent parity-check matrix is

$$\left(\begin{array}{c|c|c} A & B & T \\ \hline C & D & E \end{array} \right) = \left(\begin{array}{cccccc|cc|cccc} 1 & 1 & 1 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ \hline 0 & 1 & 0 & 1 & 0 & 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 0 & 1 \end{array} \right). \quad (11)$$

Ex. 1022.5

TABLE III
SUMMARY OF THE PROPOSED ENCODING PROCEDURE. IT ENTAILS TWO STEPS: A PREPROCESSING STEP AND THE ACTUAL ENCODING STEP

Preprocessing: Input: Non-singular parity-check matrix H . Output: An equivalent parity check matrix of the form $\begin{pmatrix} A & B & T \\ C & D & E \end{pmatrix}$ such that $-ET^{-1}B + D$ is non-singular.

1. [Triangulation] Perform row and column permutations to bring the parity check matrix H into approximate lower triangular form

$$H = \begin{pmatrix} A & B & T \\ C & D & E \end{pmatrix}$$

with as small a gap g as possible. We will see in subsequent sections how this can be accomplished efficiently.

2. [Check Rank] Use Gaussian elimination to effectively perform the pre-multiplication

$$\begin{pmatrix} I & 0 \\ ET^{-1} & I \end{pmatrix} \begin{pmatrix} A & B & T \\ C & D & E \end{pmatrix} = \begin{pmatrix} A & B & T \\ -ET^{-1}A + C & -ET^{-1}B + D & 0 \end{pmatrix} \quad (12)$$

in order to check that $-ET^{-1}B + D$ is non-singular, performing further column permutations if necessary to ensure this property. (Singularity of H can be detected at this point.)

Encoding: Input: Parity-check matrix of the form $\begin{pmatrix} A & B & T \\ C & D & E \end{pmatrix}$ such that $-ET^{-1}B + D$ is non-singular and a vector $s \in \mathbb{F}^{n-m}$. Output: The vector $x = (s, p_1, p_2)$, $p_1 \in \mathbb{F}^g$, $p_2 \in \mathbb{F}^{m-g}$, such that $Hx^T = 0^T$.

1. Determine p_1 as shown in Table 1.
2. Determine p_2 as shown in Table 2.

Assume now we choose $s = (1, 0, 0, 0, 0, 0)$. To determine p_1 we follow the steps listed in Table I. We get

$$\begin{aligned} As^T &= (1, 1, 0, 1)^T \\ T^{-1}[As^T] &= (1, 1, 0, 0)^T \\ -E[T^{-1}As^T] &= (0, 1)^T \\ Cs^T &= (0, 0)^T \\ [-ET^{-1}As^T] + [Cs^T] &= (0, 1)^T \end{aligned}$$

and

$$\phi^{-1}[-ET^{-1}As^T + Cs^T] = (0, 1)^T = p_1^T.$$

In a similar manner, we execute the steps listed in Table II to determine p_2 . We get

$$\begin{aligned} Bp_1^T &= (0, 1, 0, 0)^T \\ [As^T] + [Bp_1^T] &= (1, 0, 0, 1)^T \end{aligned}$$

and

$$T^{-1}[As^T + Bp_1^T] = (1, 0, 1, 0)^T = p_2^T.$$

Therefore the codeword is equal to

$$(s, p_1, p_2) = (1, 0, 0, 0, 0, 0, 0, 1, 1, 0, 1, 0).$$

A quick check verifies that $Hx^T = 0^T$, as required.

III. APPROXIMATE UPPER TRIANGULATION VIA GREEDY ALGORITHMS

We saw in the previous section that the encoding complexity is of order $n + g^2$, where g is the gap of the approximate tri-

angulation. Hence, for a given parity-check matrix we are interested in finding an approximate lower triangulation with as small a gap as possible. Given that we are interested in large block lengths, there is little hope of finding the optimal row and column permutation which results in the minimum gap. So we will limit ourselves to *greedy* algorithms. As discussed in the previous section, the following greedy algorithms work on the extended adjacency matrices since these are, except for the ordering of the sockets, in one-to-one correspondence with the underlying graphs.

To describe the algorithms we first need to extend some of our previous definitions. Recall that for a given pair (μ, ν) of degree distributions we associate to it two important parameters. The first parameter $r(\mu, \nu)$ is the *rate* of the degree distribution pair and is defined in (2). Note that

$$1 - r(\mu, \nu) = \frac{1}{1 - r(\nu, \mu)}. \quad (13)$$

The second parameter $\alpha^*(\mu, \nu)$ is called the *threshold* of the degree distribution pair and is defined in (4). If $r(\mu, \nu) \geq 0$, as we have tacitly assumed so far, then we can think of (μ, ν) as the degree distribution pair of an ensemble of LDPC codes of rate $r(\mu, \nu)$. Further, as discussed in Section I, in this case it was shown in [9] that $\alpha^*(\mu, \nu)$ is the threshold of this ensemble when transmitting over the BEC assuming a belief propagation decoder. In general, $r(\mu, \nu)$ may be negative and, hence, the degree distribution pair does not correspond to an ensemble of LDPC codes. Nevertheless, the definitions are still meaningful.

Example 4: Let $(\mu(x), \nu(x)) = (x^5, x^2)$. In this case, we have $r(x^5, x^2) = -1$ and, using the techniques described in [17], the threshold can be determined to be

$$\alpha^*(x^5, x^2) = \frac{3^{18}}{2^{1755}} \sim 0.945851. \quad \square$$

In a similar way, the definition of the ensemble $\mathcal{C}^l(\mu, \nu)$ as well as the association of (extended) adjacency matrices to elements of $\mathcal{C}^l(\mu, \nu)$ carry over to the case $r(\mu, \nu) \leq 0$. Assume now that, for a given ensemble $\mathcal{C}^l(\mu, \nu)$, we create a new ensemble by simply exchanging the roles of left and right nodes. This new ensemble is equivalent to the ensemble

$$\mathcal{C}^{(1-r(\mu, \nu))l}(\nu, \mu) = \mathcal{C}^{\frac{1}{1-r(\mu, \nu)}l}(\nu, \mu)$$

where we have used (13). For the associated (extended) adjacency matrices this simply amounts to transposition.

Assume we are given a matrix A of dimension $(1-r)l \times l$ with elements in $\mathbb{N}$, where r is some real-valued parameter with $r \leq 1$. We will say that a row and a column are *connected* if the corresponding entry in A is nonzero. Furthermore, we will say that a row (column) has *degree* i if its row (column) sum equals i . Assume now that we want to bring A into approximate lower triangular form. The class of greedy algorithms that we will consider is based on the following simple procedure. Given the $(1-r)l \times l$ matrix A and a fixed integer k , $k \leq (1-r)l$, permute, if possible, the rows and columns in such a way that the first row has its last nonzero entry at position $(l-k+1)$. If this first step was successful then fix the first row and permute, if possible, the remaining rows and all columns in such a way that the second row has its last nonzero entry at position $(l-k+2)$. In general, assuming that the first $i-1$ steps were successful, permute at the i th step, if possible, the last $(1-r)l - (i-1)$ rows and all columns in such a way that the i th row has its last nonzero entry at position $(l-k+i)$. If this procedure does not terminate before the k th step then we accomplished an approximate lower triangulation of the matrix A . We will say that A is in approximate lower triangular form with *row gap* $(1-r)l-k$ and *column gap* $l-k$, as shown in Fig. 4.

A. Greedy Algorithm A

We will now give a precise description of the greedy algorithm A. The core of the algorithm is the *diagonal extension step*.

Diagonal Extension Step: Assume we are given a matrix A and a subset of the columns which are classified as *known*. In all cases of interest to us, either none of these known columns are connected to rows of degree one or all of them are. Assume the latter case. Let $c_1, \dots, c_k$ denote the known columns and let $r_1, \dots, r_k$ be degree-one rows such that c_i is connected to r_i .³ Reorder, if necessary, the rows and columns of A such that $r_1, \dots, r_k$ form the leading k rows of A and such that $c_1, \dots, c_k$ form the leading k columns of A as shown in Fig. 5, where $\tilde{A}$ denotes the submatrix of A which results from deleting the rows and columns indexed by $r_1, \dots, r_k$ and $c_1, \dots, c_k$. Note that after this reordering the top-left $k \times k$ submatrix of A has diagonal form and that the top k rows of A have only this one nonzero entry.

³ r_i may not be determined uniquely.




Fig. 4. Approximate lower triangulation of the matrix A with row gap $(1-r)l-k$ and column gap $l-k$ achieved by a greedy algorithm.




Fig. 5. Given the matrix A let $c_1, \dots, c_k$ denote those columns which are connected to rows of degree one and let $r_1, \dots, r_k$ be degree-one rows such that c_i is connected to r_i . Reorder the rows and columns in such a way that $r_1, \dots, r_k$ form the first k rows and such that $c_1, \dots, c_k$ form the first k columns. Note that the top-left $k \times k$ submatrix has diagonal form and that the first k rows have only this one nonzero entry.

By a diagonal extension step we will mean the following. As input, we are given the matrix A and a set of known columns. The algorithm performs some row and column permutations and specifies a *residual matrix* $\tilde{A}$. More precisely, if none of the known columns are connected to rows of degree one then perform a column permutation so that all the known columns form the leading columns of the matrix A . Furthermore, delete these known columns from the original matrix and declare the resulting matrix to be $\tilde{A}$. If, on the other hand, all known columns are connected to rows of degree one then perform a row and column permutation to bring A into the form depicted in Fig. 5. Furthermore, delete the known columns $c_1, \dots, c_k$ and the rows $r_1, \dots, r_k$ from the original matrix and declare the resulting matrix to be $\tilde{A}$.

In terms of this diagonal extension step, greedy algorithm A has a fairly succinct description.

Greedy Algorithm A:

0. [Initialization] Given a matrix A declare each column independently to be *known* with probability $1 - \alpha$ or, otherwise, to be an *erasure*. Let $\tilde{A} := A$.
1. [Stop or Extend] If $\tilde{A}$ contains neither a known column nor a row of degree one then output the present matrix. Otherwise, perform a diagonal extension step.
2. [Declare Variables as Known] Any column in $\tilde{A}$ which is connected to a degree one row is declared to be known. Goto 1.




Fig. 6. (a) The given matrix A . (b) After the first application of step one, the $(1-\alpha)l$ known columns are reordered to form the first $(1-\alpha)l$ columns of the matrix A . (c) After the second application of step one, the k new known columns and their associated rows are reordered to form a diagonal of length k . (d) If the procedure does not terminate prematurely then the diagonal is extended to have length αl and, therefore, the row gap is equal to $(1-r-\alpha)l$ and the column gap is equal to $(1-\alpha)l$.

To see that greedy algorithm A indeed gives rise to an approximate triangulation assume that we start with the $(1-r)l \times l$ matrix A as shown in Fig. 6(a). In the initialization step, an expected fraction $(1-\alpha)$ of all columns are classified as known and the rest is classified as erasures. The first time the algorithm performs step one these $(1-\alpha)l$ known columns are reordered to form the leading columns of the matrix A as shown in Fig. 6(b). Assuming that the residual matrix has rows of degree one, the columns connected to these degree-one rows are identified in the second step. Let these columns be $c_1, \dots, c_k$ and let $r_1, \dots, r_k$ be degree-one rows such that c_i is connected to r_i . During the second application of step one these new known columns and their associated rows are ordered along a diagonal as shown in Fig. 6(c). Furthermore, in each additional iteration this diagonal is extended further. If this procedure does not stop prematurely then the resulting diagonal has expected length αl and, therefore, the row gap has expected size $(1-r-\alpha)l$ and the column gap has expected size $(1-\alpha)l$ as shown in Fig. 6(d). If, on the other hand, the procedure terminates before all columns are exhausted then we get an approximate triangulation by simply reordering the remaining columns to the left. Assuming that the remaining fraction of columns is equal to ϵl then the resulting expected row gap is equal to $(1-r-\alpha+\epsilon)l$ and the resulting expected column gap is equal to $(1-\alpha+\epsilon)l$.

Lemma 1 [Performance of Greedy Algorithm A]: Let (μ, ν) be a given degree pair and choose $\alpha < \alpha^*(\mu, \nu)$. Pick a graph at random from the ensemble $\mathcal{C}^l(\mu, \nu)$ and let $\hat{A}$ be its extended adjacency matrix. Apply greedy algorithm A to the extended adjacency matrix $\hat{A}$. Then (asymptotically in l) the row gap is concentrated around the value $(1-r-\alpha)l$ and the column gap is concentrated around the value $(1-\alpha)l$. Letting $\alpha \uparrow \alpha^*$, we see that the minimum row gap achievable with greedy algorithm A is equal to $(1-r-\alpha^*)l$ and that the minimum column gap is equal to $(1-\alpha^*)l$.

Proof: Assume we are given a graph and an associated extended adjacency matrix $\hat{A}$ from the ensemble $\mathcal{C}^l(\mu, \nu)$.

Assume first that $r(\lambda, \rho) \geq 0$ so that $\mathcal{C}^l(\mu, \nu)$ represents an ensemble of LDPC codes of rate $r(\lambda, \rho)$. For the same code/graph consider the process of transmission over an erasure channel with erasure probability α followed by decoding using the message-passing decoder described in Section I. Compare this procedure to the procedure of the greedy algorithm A. Assume that the bits erased by the channel correspond to exactly those columns which in the initial step are classified as erasures. Under this assumption, one can see that those columns which are declared known in the ℓ th round of greedy algorithm A correspond exactly to those variable nodes which are declared known in the ℓ th round of the decoding algorithm. Hence, there is a one-to-one correspondence between these two algorithms.

As discussed in Section I, if $\alpha < \alpha^*(\mu, \nu)$ then (asymptotically in l) with high probability the decoding process will be successful. Because of the one-to-one correspondence we conclude that in this case (asymptotically in l) greedy algorithm A will extend the diagonal to (essentially) its full length αl with high probability so that the row and column gaps are as stated in the Lemma.

In the case that $r(\mu, \nu) < 0$ we cannot associate an ensemble of codes to the degree distribution pair (μ, ν) . Nevertheless, recursion (3) still correctly describes the expected progress of greedy algorithm A. It is also easy to see that the concentration around this expected value still occurs. It follows that the same analysis is still valid in this case. $\square$

1) Greedy Algorithm AH: By greedy algorithm AH we mean the direct application of greedy algorithm A to the extended parity-check matrix $\hat{H}$ of a given LDPC code. The gap we are interested in is then simply the resulting row gap.

Corollary 1 (Performance of Greedy Algorithm AH): Let (λ, ρ) be a given degree distribution pair with $r(\lambda, \rho) \geq 0$ and choose $\alpha < \alpha^*(\lambda, \rho)$. Pick a code at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$ and let $\hat{H}$ be the associated extended parity-check matrix. Apply greedy algorithm A to $\hat{H}$. Then (asymptotically in n) the gap is concentrated around the

value $(1 - r(\lambda, \rho) - \alpha)n$. Letting $\alpha \uparrow \alpha^*$, we see that the minimum gap achievable with greedy algorithm A is equal to $(1 - r(\lambda, \rho) - \alpha^*)n$.

Example 5 [Gap for the (3,6)-Regular Code and Greedy Algorithm AH]: From Example 2, we know that $r(x^2, x^5) = \frac{1}{2}$ and that $\alpha^*(x^2, x^5) \sim 0.429439$. It follows that the minimum expected gap size for greedy algorithm AH is equal to $g \sim 0.070561n$. $\square$

Note that greedy algorithm A establishes a link between the error-correcting capability on a BEC using a message-passing decoder and the encoding complexity. In simplified terms: Good codes have low encoding complexity!

2) *Greedy Algorithm AHT:* Rather than applying greedy algorithm A directly to the extended parity-check matrix $\hat{H}$ of an LDPC code we can apply it to the transpose of the extended parity-check matrix. In this case, the gap we are interested in is equal to the resulting column gap.

Corollary 2 (Performance of Greedy Algorithm AHT): Let (λ, ρ) be a given degree distribution pair with $r(\lambda, \rho) \geq 0$ and choose $\alpha < \alpha^*(\rho, \lambda)$. Pick a code at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$ and let $\hat{H}$ be the associated extended parity-check matrix. Apply greedy algorithm A to $\hat{H}^T$. Recall that this is equivalent to applying greedy algorithm A to a randomly chosen extended adjacency matrix from the ensemble $\mathcal{C}^{\frac{n}{1-r(\rho, \lambda)}}(\rho, \lambda)$. Therefore, (asymptotically in n) the gap is concentrated around the value $\frac{1-\alpha}{1-r(\rho, \lambda)}n$. Letting $\alpha \uparrow \alpha^*(\rho, \lambda)$, we see that the minimum gap achievable with greedy algorithm AHT is equal to $\frac{1-\alpha^*(\rho, \lambda)}{1-r(\rho, \lambda)}n$.

Example 6 [Gap for the (3,6)-Regular Code and Greedy Algorithm AHT]: From Example 4, we know that $r(x^5, x^2) = -1$ and that $\alpha^*(x^5, x^2) = \frac{3^{18}}{2^{17 \cdot 5^5}}$. It follows that the minimum expected gap size for greedy algorithm AHT is equal to

$$g = \frac{1 - \frac{3^{18}}{2^{17 \cdot 5^5}}}{2} n = \frac{22\,179\,511}{2^{18 \cdot 5^5}} n \sim 0.0270746n.$$

Example 7 (Gap for an “Optimized” Code of Maximal Degree 12 and Greedy Algorithm AHT): Let us determine the threshold $\alpha^*(\rho, \lambda)$ for one of the “optimized” codes listed in [12]. We pick the code with

$$\lambda(x) := 0.251x + 0.309x^2 + 0.002x^3 + 0.438x^9$$

and

$$\rho(x) := 0.637x^6 + 0.363x^7.$$

Quite surprisingly we get $\alpha^*(\rho, \lambda) = 1$! This means that for any $\epsilon > 0$ we can start the process by declaring only an $\epsilon/2$ fraction of all columns to be known and, with high probability, the process will continue until at most an $\epsilon/2$ fraction of all columns is left. Therefore, we can achieve a gap of ϵn for any $\epsilon > 0$. We will later prove a stronger result, namely, that in this case the gap is actually at most of order $\sqrt{n}$, but we will need more sophisticated techniques to prove this stronger result. $\square$

The above example shows that at least for some degree distribution pairs (μ, ν) we have $\alpha^*(\mu, \nu) = 1$. When does this happen? This is answered in the following lemma.

Lemma 2: Let (μ, ν) be a degree distribution pair. Then $\alpha^*(\mu, \nu) = 1$ if and only if for all $\alpha \in [0, 1]$

$$x > \alpha\mu(1 - \nu(1 - x)) \quad \forall x \in (0, \alpha]. \quad (14)$$

Furthermore, if (14) holds and $\nu(0) = \mu(0) = 0$, then

$$\mu_2\nu'(1) \leq 1 \leq \mu'(1)\nu_2. \quad (15)$$

Proof: Clearly, if (14) holds then for any $\alpha \in [0, 1]$ we have

$$x > \alpha\mu(1 - \nu(1 - x)) \quad \forall x \in (0, \alpha].$$

By a compactness argument it follows that $\alpha_\ell(\alpha)$ as defined in (3) converges to 0 as ℓ tends to infinity. Hence, $\alpha^*(\mu, \nu) = 1$.

Assume now that $\alpha^*(\mu, \nu) = 1$. This means that for any $\alpha \in [0, 1]$ we have that $\alpha_\ell(\alpha) \xrightarrow{\ell \rightarrow \infty} 0$. We want to show that (14) holds. Let $f(x, \alpha) := \alpha\mu(1 - \nu(1 - x))$ and note that for $x, \alpha \in [0, 1]$, $f(x, \alpha)$ is an increasing function in both its arguments. Note that because $f(x, \alpha)$ is increasing in x it follows that a necessary condition for $\alpha_\ell(\alpha)$ to converge to zero is that $\alpha > f(\alpha, \alpha)$, i.e., that at least in the first iteration the erasure probability decreases. We will use contraposition to prove (14). Hence, assume that there exist a strictly positive x and an α , $x \leq \alpha < 1$, such that $x \leq f(x, \alpha)$. Since $\alpha > f(\alpha, \alpha)$ and since f is continuous this implies that there exists a strictly positive x and an α , $x \leq \alpha < 1$, such that $x = f(x, \alpha)$. Then

$$\alpha_1(\alpha) = f(\alpha, \alpha) \geq f(x, \alpha) = x.$$

It follows by finite induction that

$$\alpha_\ell(\alpha) = f(\alpha_{\ell-1}, \alpha) \geq f(x, \alpha) = x$$

and, therefore, $\alpha_\ell(\alpha)$ does not converge to zero as ℓ tends to infinity, a contradiction.

Finally, for x close to one we have

$$\mu(1 - \nu(x)) = \mu_2\nu'(1)(1 - x) + O((1 - x)^2)$$

whereas for x tending to zero we have

$$\mu(1 - \nu(x)) = 1 - \mu'(1)\nu_2x + O(x^2).$$

This yields the stability conditions stated in (15). $\square$

B. Greedy Algorithm B

For greedy algorithm A, the elements of the initial set of known columns are chosen independently from each other. We will now show that by allowing dependency in the initial choice, the resulting gap can sometimes be reduced. Of course, this dependency makes the analysis more difficult.

In order to describe and analyze greedy algorithm B we need to introduce some more notation. We call a polynomial

$$\omega(x) := \sum_{i \geq 1} \omega_i x^{i-1}$$

with real nonnegative coefficients in the range $[0, 1]$ a *weight distribution*, and we denote the set of all such weight distributions by $\mathcal{W}$. Let $\tau : \mathcal{D} \times \mathcal{W} \rightarrow \mathcal{D}$ be a map which maps a pair

consisting of a degree distribution ν and a weight distribution ω into a new degree distribution $\tau(\nu, \omega)$. This map is defined as

$$\begin{aligned}\tau(\nu, \omega)_1 &:= \nu_1 + \sum_{i \geq 2} \nu_i \omega_i \\ \tau(\nu, \omega)_i &:= \nu_i (1 - \omega_i), \quad i \geq 2.\end{aligned}$$

We are now ready to state greedy algorithm B.

Greedy Algorithm B:

0. [Initialization] We are given a matrix A and a weight distribution $\omega(x)$. For each row in A perform the following: if the row has weight i then select this row with probability ω_i . For each selected row of weight i declare a random subset of size $(i - 1)$ of its i connected columns to be known. All remaining columns which have not been classified as known are classified as erasures. Let $\tilde{A} := A$.
1. [Stop or Extend] If $\tilde{A}$ neither contains a known column nor a row of degree one then output the present matrix. Otherwise, perform a diagonal extension step.
2. [Declare Variables as Known] Any column in $\tilde{A}$ which is connected to a degree one row is declared to be known. Goto 1

Clearly, greedy algorithm B differs from greedy algorithm A only in the choice of the initial set of columns.

Lemma 3 (Analysis of Greedy Algorithm B): Let (μ, ν) be a given degree distribution pair. Let $\omega(x)$ be a weight distribution such that $\alpha^*(\mu, \tau(\nu, \omega)) = 1$. Define $q := \sum_{i \geq 1} \nu_i \omega_i \frac{i-1}{i}$. Pick a graph at random from the ensemble $\mathcal{C}^l(\mu, \nu)$ and let $\hat{A}$ be its extended adjacency matrix. Apply greedy algorithm B to the extended adjacency matrix $\hat{A}$. Then (asymptotically in l) the row gap is concentrated around the value

$$\left(\frac{\sum_{i \geq 1} [1 - (1 - q)^i] \frac{\mu_i}{i}}{\int \mu} - r(\mu, \nu) \right) l$$

and the column gap is concentrated around the value

$$\frac{\sum_{i \geq 1} [1 - (1 - q)^i] \frac{\mu_i}{i}}{\int \mu} l.$$

Proof: The elements of the initial set of known columns are clearly dependent (since groups of those columns are connected to the same row) and therefore we cannot apply our previous methods directly. But as we will show now there is a one-to-one correspondence between applying greedy algorithm B to the ensemble $\mathcal{C}^l(\mu, \nu)$ with a weight distribution $\omega(x)$ and applying greedy algorithm A to the transformed ensemble $\mathcal{C}^l(\mu, \tau(\nu, \omega))$.

Assume we are given the ensemble $\mathcal{C}^l(\mu, \nu)$ and a weight distribution $\omega(x)$. Assume further that we are given a fixed set of selected right nodes (rows) $r_1, \dots, r_k$, and that the fraction of selected right nodes of degree i is equal to ω_i . Given a graph from $\mathcal{C}^l(\mu, \nu)$ transform it in the following way: replace each selected right node of degree i by i right nodes of degree 1. One can check that this transformation leaves the left degree distribution μ unchanged and that it transforms the right degree dis-

tribution ν to $\tau(\nu, \omega)$. Therefore, the new graph is an element of the ensemble $\mathcal{C}^l(\mu, \tau(\nu, \omega))$. Further, one can check that this map is reversible and, therefore, one-to-one. A closer look reveals now that applying greedy algorithm B to an extended adjacency matrix picked randomly from the ensemble $\mathcal{C}^l(\mu, \nu)$ is equivalent to applying greedy algorithm A with $\alpha = 0$ to the transformed extended adjacency matrix, i.e., the resulting residual graphs (which could be empty) will be the same. Now, since $\tau(\nu, \omega)_1 > 0$ it follows that the greedy algorithm B will get started and since by assumption $\alpha^*(\mu, \tau(\nu, \omega)) = 1$ we know from the analysis of greedy algorithm A that with high probability the diagonalization process will continue until the diagonal has been extended to (essentially) its full length. In this case, the resulting column gap is equal to the size of the set which was initially classified as known. To determine the size of this set we first determine the probability that a randomly chosen edge is one of those edges which connect a selected right node to one of its $(i - 1)$ declared known neighbors. A quick calculation shows that this probability is equal to $q = \sum_{i \geq 1} \nu_i \omega_i \frac{i-1}{i}$. Therefore, the probability that a given left node of degree i is connected to at least one of these edges is equal to $1 - (1 - q)^i$. From this the stated row and column gaps follow easily. $\square$

1) Greedy Algorithm BH: Following our previous notation, by greedy algorithm BH we mean the direct application of greedy algorithm B to the extended parity-check matrix $\hat{H}$ of a given LDPC code. The gap we are interested in is then simply the resulting row gap.

Corollary 3 (Performance of Greedy Algorithm BH): Let (λ, ρ) be a given degree distribution pair with $r(\lambda, \rho) \geq 0$. Let $\omega(x)$ be a weight distribution such that $\alpha^*(\lambda, \tau(\rho, \omega)) = 1$. Define

$$q := \sum_{i \geq 1} \rho_i \omega_i \frac{i-1}{i}.$$

Pick a code at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$ and let $\hat{H}$ be its extended parity-check matrix. Apply greedy algorithm B to the extended parity-check matrix $\hat{H}$. Then (asymptotically in n) the gap is concentrated around the value

$$\left(\frac{\sum_{i \geq 1} [1 - (1 - q)^i] \frac{\lambda_i}{i}}{\int \lambda} - r(\lambda, \rho) \right) n.$$

Let

$$q^* := \inf_{\omega(x)} \left\{ \sum_{i \geq 1} \rho_i \omega_i \frac{i-1}{i} : \alpha^*(\lambda, \tau(\rho, \omega)) = 1 \right\}.$$

Then we see that the minimum gap achievable with greedy algorithm BH is equal to

$$\left(\frac{\sum_{i \geq 1} [1 - (1 - q^*)^i] \frac{\lambda_i}{i}}{\int \lambda} - r(\lambda, \rho) \right) n.$$

Example 8 [Gap for the (3, 6)-Regular Code and Greedy Algorithm BH]: We have $\lambda(x) = x^2$ and $\rho(x) = x^5$ and since $\rho(x)$ has only one nonzero term it follows that we can param-

eterize $\omega(x)$ as $\omega(x) = \omega_6 x^5$. Therefore, we have $\tau(\rho, \omega) = \omega_6 + (1 - \omega_6)x^5$ and since $q = \omega_6 \frac{5}{6}$, it follows that we need to find the smallest value of ω_6 , call it ω_6^* , such that $\alpha^*(x^2, \omega_6 + (1 - \omega_6)x^5) = 1$. From Lemma (14) we see that a necessary and sufficient condition is given by

$$1 - x > (1 - \omega_6 - (1 - \omega_6)x^5)^2 = (1 - \omega_6)^2(1 - x^5)^2 \quad \forall x \in (0, 1).$$

Equivalently, we get

$$(1 - \omega_6)^2 < \frac{1 - x}{(1 - x^5)^2} \quad \forall x \in (0, 1).$$

Differentiating shows that the right-hand side takes on its minimum at the unique positive root of the polynomial $-1 - x - x^2 - x^3 + 9x^4$. If we call this root x^* , with $x^* \sim 0.739429$, then we conclude that

$$\omega_6^* = 1 - \sqrt{\frac{1 - x_0^*}{(1 - (x_0^*)^5)^2}} \sim 0.344683.$$

We then get $q^* = \omega_6^* \frac{5}{6} \sim 0.287236$ and, therefore, the gap is equal to

$$1 - (1 - q^*)^3 - \frac{1}{2} \simeq 0.137893n.$$

Note that in this case the gap is larger than the corresponding gap for greedy algorithm AH. $\square$

2) *Greedy Algorithm BHT*: Again as for greedy algorithm A, rather than applying greedy algorithm B directly to the extended parity-check matrix $\hat{H}$ of an LDPC code we can apply it to the transpose of the extended parity-check matrix. In this case, the gap we are interested in is equal to the resulting column gap.

Corollary 4 (Performance of Greedy Algorithm BHT): Let (λ, ρ) be a given degree distribution pair with $r(\lambda, \rho) \geq 0$. Let $\omega(x)$ be a weight distribution such that $\alpha^*(\rho, \tau(\lambda, \omega)) = 1$. Define $q := \sum_{i \geq 1} \lambda_i \omega_i \frac{i-1}{i}$. Pick a code at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$ and let $\hat{H}$ be its extended parity-check matrix. Apply greedy algorithm B to $\hat{H}^T$. Recall that this is equivalent to applying greedy algorithm B to a randomly chosen extended adjacency matrix from the ensemble $\mathcal{C}^{\frac{n}{1-r(\rho, \lambda)}}(\rho, \lambda)$. Therefore, (asymptotically in n) the gap is concentrated around the value

$$\frac{\sum_{i \geq 1} [1 - (1 - q)^i] \frac{\rho_i}{i}}{(1 - r(\rho, \lambda)) \int \rho} n.$$

Let

$$q^* := \inf_{\omega(x)} \left\{ \sum_{i \geq 1} \lambda_i \omega_i \frac{i-1}{i} : \alpha^*(\rho, \tau(\lambda, \omega)) = 1 \right\}.$$

Then we see that the minimum gap achievable with greedy algorithm BHT is equal to

$$\frac{\sum_{i \geq 1} [1 - (1 - q^*)^i] \frac{\rho_i}{i}}{(1 - r(\rho, \lambda)) \int \rho} n.$$

Example 9 [Gap for the (3, 6)-Regular Code and Greedy Algorithm BHT]: We have $\lambda(x) = x^2$ and $\rho(x) = x^5$ and since $\lambda(x)$ has only one nonzero term we can parameterize $\omega(x)$ as $\omega(x) = \omega_3 x^2$. Therefore, we have $\tau(\lambda, \omega) = \omega_3 + (1 - \omega_3)x^2$ and since $q = \omega_3 \frac{2}{3}$ it follows that we need to find the smallest value of ω_3 , call it ω_3^* , such that $\alpha^*(x^5, \omega_3 + (1 - \omega_3)x^2) = 0$. From Lemma 2 (14) we see that a necessary and sufficient condition is given by

$$1 - x > (1 - \omega_3 - (1 - \omega_3)x^2)^5 = (1 - \omega_3)^5(1 - x^2)^5 \quad \forall x \in [0, 1]$$

which simplifies to

$$1 > (1 - \omega_3)^5(1 - x^2)^4(1 + x).$$

By differentiating $(1 - x^2)^4(1 + x)$ we find that it takes its minimum at $x = 1/9$. Thus, the critical value of ω_3^* is given by

$$\omega_3^* = 1 - \frac{81}{80} \left(\frac{8}{9} \right)^{\frac{1}{5}} = 0.01107.$$

We then get $q^* = \omega_3^* \frac{2}{3} \sim 0.00738$. This corresponds to a gap of

$$\frac{(1 - (1 - q^*)^6)}{2} n \simeq 0.02174n.$$

This is significantly better than the corresponding gap for greedy algorithm AHT. $\square$

C. Greedy Algorithm C

Let (μ, ν) be the given degree distribution pair. Recall that for greedy algorithm B we chose the weight distribution $\omega(x)$ in such a way that $\alpha^*(\mu, \tau(\nu, \omega)) = 1$. Hence, with high probability, the greedy algorithm will extend the diagonal to (essentially) its full length.

Alternatively, we can try to achieve an approximate triangulation in several smaller steps. More precisely, assume that we pick the weight distribution $\omega(x)$ in such a way that $\alpha^*(\mu, \tau(\nu, \omega)) < 1$. Then with high probability the greedy algorithm will not complete the triangulation process. Note that, conditioned on the size and on the degree distribution pair of the resulting residual graph, the edges of this residual graph are still random, i.e., if the residual graph has length $\tilde{n}$ and a degree distribution pair $(\tilde{\mu}, \tilde{\nu})$ then we can think of it as an element of $\mathcal{C}^{\tilde{n}}(\tilde{\mu}, \tilde{\nu})$. This is probably easiest seen by checking that if the destination of two edges which are contained in the residual graph are interchanged in the original graph and if the greedy algorithm B is applied to this new graph then the new residual graph will be equal to the old residual graph except for this interchange. Therefore, if we achieve a triangulation by applying several small steps, then we can still use the previous tools to analyze the expected gap.

There are obviously many degrees of freedom in the choice of step sizes and the choice of weight distribution. In our present discussion, we will focus on the limiting case of infinitesimal small step sizes and a constant weight distribution. Therefore, assume that we are given a fixed weight distribution $\omega(x)$ and let $\epsilon, \epsilon > 0$, be a small scaling parameter for the weights such

that $\alpha^*(\mu, \tau(\nu, \epsilon\omega)) < 1$. Assume that we apply greedy algorithm B to a randomly chosen element of the ensemble $\mathcal{C}^l(\mu, \nu)$ where $\mu(0) = \nu(0) = 0$. We claim that the expected degree distribution pair of the residual graph, call it $(\tilde{\mu}, \tilde{\nu})$, is given by

$$\begin{aligned}\tilde{\mu}_i &\leftarrow \mu_i + \mu_i \frac{\left(\sum_{j \geq 2} \nu_j \omega_j\right) \left(\sum_{j \geq 2} \mu_j j - i\right)}{1 - \mu'(1)\nu_2} \epsilon + O(\epsilon^2) \\ \tilde{\nu}_i &\leftarrow \nu_i + \nu_i \left(\sum_{j \geq 2} \omega_j \nu_j - \omega_i\right) \epsilon + (i\nu_{i+1} - (i-1-\nu_2)\nu_i) \\ &\quad \cdot \frac{\mu'(1) \sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1)\nu_2} \epsilon + O(\epsilon^2).\end{aligned}$$

To see this, first recall from the analysis of greedy algorithm B that the degree distribution pair of the equivalent transformed graph is equal to $(\mu, \tau(\nu, \epsilon\omega))$. Since by assumption $\alpha^*(\mu, \tau(\nu, \epsilon\omega)) < 1$, the recursion given in (3) (with $\alpha = 1$) will have a fixed point, i.e., there exists a real number x , $x < 1$, such that

$$\mu(1 - \tau(\nu, \epsilon\omega)(1 - x)) = x.$$

To determine this fixed point note that if we expand the above in $(1 - x)$ around $(1 - x) = 0$ we obtain

$$\begin{aligned}\mu(1 - \tau(\nu, \epsilon\omega)(1 - x)) \\ = 1 - \mu'(1) \left[\epsilon \sum_{i \geq 2} \nu_i \omega_i + \nu_2(1 - \epsilon\omega_2)(1 - x) \right] \\ + O((1 - x)^2).\end{aligned}$$

Therefore, letting Δx denote $1 - x$, the fixed-point equation is

$$\Delta x = \mu'(1) \left[\epsilon \sum_{i \geq 2} \nu_i \omega_i + \nu_2 \Delta x \right] + O((\Delta x)^2) + O(\epsilon \Delta x).$$

It follows that

$$\Delta x = \frac{\mu'(1) \sum_{i \geq 2} \nu_i \omega_i}{1 - \mu'(1)\nu_2} \epsilon + O(\epsilon^2).$$

In the language of message-passing algorithms, $1 - \Delta x$ is the expected fraction of erasure messages passed from left to right at the time the algorithm stops. The fraction of erasure messages which are passed at that time from right to left is then

$$1 - \tau(\nu, \epsilon\omega)(\Delta x) = 1 - \frac{\sum_{i \geq 2} \nu_i \omega_i}{1 - \mu'(1)\nu_2} \epsilon + O(\epsilon^2). \quad (16)$$

We start by determining the residual degree distribution of left nodes. Note that a left node will not appear in the residual graph

iff at least one of its incoming messages is not an erasure—otherwise, it stays and retains its degree. Using (16) we see that a node of degree i has a probability of

$$\frac{\sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1)\nu_2} i \epsilon + O(\epsilon^2)$$

of being expurgated. Since in the original graph the number of left degree i nodes is proportional to $\frac{\mu_i}{i}$ it follows that in the residual graph the number of left degree i nodes is proportional to

$$\frac{\mu_i}{i} \left[1 - \frac{\sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1)\nu_2} i \epsilon + O(\epsilon^2) \right].$$

From an edge perspective the degree i fraction of the residual graph is, therefore, proportional to

$$\mu_i \left[1 - \frac{\sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1)\nu_2} i \epsilon + O(\epsilon^2) \right].$$

After normalization we find that the left degree distribution of the residual graph, call it $\tilde{\mu}$, is given by

$$\tilde{\mu}_i \leftarrow \mu_i \left[1 + \frac{\left(\sum_{j \geq 2} \nu_j \omega_j\right) \left(\sum_{j \geq 2} \mu_j j - i\right)}{1 - \mu'(1)\nu_2} \epsilon + O(\epsilon^2) \right].$$

We next determine the right degree distribution of the residual graph. Recall that the equivalent transformed graph has a right degree distribution of $\tau(\nu, \epsilon\omega)$. We are only interested in nodes of degree at least two. Hence we have

$$\tau(\nu, \epsilon\omega)_i \propto \nu_i(1 - \epsilon\omega_i), \quad i \geq 2.$$

From a node perspective these fractions are proportional to

$$\frac{\nu_i}{i}(1 - \epsilon\omega_i), \quad i \geq 2.$$

Define the *erasure degree* of a right node to be equal to the number of incoming edges which carry erasure messages. To first order in ϵ , a node of erasure degree i can stem either from a node of regular degree i all of whose incoming messages are erasures or it can stem from a node of regular degree $(i + 1)$ which has one nonerasure message. Hence, at the fixed point the fraction of right nodes with an erasure degree of i is proportional to

$$\begin{aligned}\frac{\nu_i}{i}(1 - \epsilon\omega_i)(1 - \Delta x)^i \\ + \frac{\nu_{i+1}}{i+1}(1 - \epsilon\omega_{i+1})(i+1)\Delta x(1 - \Delta x)^i + O(\epsilon^2).\end{aligned}$$

Converting back to an edge perspective we see that these fractions are proportional to

$$\begin{aligned}\nu_i(1 - \epsilon\omega_i)(1 - \Delta x)^i \\ + \nu_{i+1}i(1 - \epsilon\omega_{i+1})\Delta x(1 - \Delta x)^i + O(\epsilon^2) \\ = \nu_i(1 - (\epsilon\omega_i + i\Delta x)) + \nu_{i+1}i\Delta x + O(\epsilon^2).\end{aligned}$$

Ex. 1022.12

Summing the above over $i \geq 2$ we obtain

$$1 - \epsilon \sum_{i \geq 2} \nu_i \omega_i - (1 + \nu_2) \Delta x.$$

Noting that $\Delta x = O(\epsilon)$ and normalizing we see that the residual right degree distribution, call it $\tilde{\nu}$, is given by

$$\begin{aligned} \tilde{\nu}_i &\leftarrow \nu_i \left(1 - \epsilon \left(\omega_i - \sum_{j \geq 2} \nu_j \omega_j \right) - (i - 1 - \nu_2) \Delta x \right) \\ &\quad + \nu_{i+1} i \Delta x + O(\epsilon^2) \\ &= \nu_i + \nu_i \left(\sum_{j \geq 2} \omega_j \nu_j - \omega_i \right) \epsilon \\ &\quad + (i \nu_{i+1} - (i - 1 - \nu_2) \nu_i) \frac{\mu'(1) \sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1) \nu_2} \epsilon + O(\epsilon^2). \end{aligned}$$

We are ultimately interested in the resulting row and column gaps. Since one can easily be determined from the other we will only write down an expression for the row gap. If we take the expression for the row gap, call it g_r from greedy algorithm B, and keep only the terms which are linear in ϵ then we see that the row gap increased according to

$$g_r \leftarrow g_r + \frac{\sum_{j \geq 2} \nu_j \omega_j \frac{j-1}{j}}{\sum_{j \geq 2} \frac{\mu_j}{j}} l \epsilon + O(\epsilon^2).$$

The length l of the code itself evolves as

$$l \leftarrow l \left[1 - \frac{\sum_{j \geq 2} \nu_j \omega_j}{(1 - \mu'(1) \nu_2) \sum_{j \geq 2} \frac{\mu_j}{j}} \epsilon \right] + O(\epsilon^2).$$

Collecting all results we see that as a function of the independent variable ϵ all quantities evolve according to the system of differential equations

$$\begin{aligned} \frac{\partial \mu_i}{\partial \epsilon} &= \mu_i \frac{\left(\sum_{i \geq 2} \nu_i \omega_i \right) \left(\sum_{j \geq 2} \mu_j j - i \right)}{1 - \mu'(1) \nu_2} \\ \frac{\partial \nu_i}{\partial \epsilon} &= \nu_i \left(\sum_{j \geq 2} \omega_j \nu_j - \omega_i \right) \\ &\quad + (i \nu_{i+1} - (i - 1 - \nu_2) \nu_i) \frac{\mu'(1) \sum_{j \geq 2} \nu_j \omega_j}{1 - \mu'(1) \nu_2} \\ \frac{\partial g_r}{\partial \epsilon} &= \frac{\sum_{j \geq 2} \nu_j \omega_j \frac{j-1}{j}}{\sum_{j \geq 2} \frac{\mu_j}{j}} l \\ \frac{\partial l}{\partial \epsilon} &= - \frac{\sum_{j \geq 2} \nu_j \omega_j}{(1 - \mu'(1) \nu_2) \sum_{j \geq 2} \frac{\mu_j}{j}} \end{aligned}$$

with the value of the initial quantities equal to μ , ν , 0, and l , respectively.

As before, we can apply greedy algorithm C directly to the extended parity-check matrix chosen randomly from an ensemble

$\mathcal{C}^n(\lambda, \rho)$ in which case we are interested in the resulting column gap or we can apply it to the transpose of the extended parity-check matrix in which case we are interested in the column gap. We call these algorithms CH and CHT, respectively.

Example 10 [Gap for (3,6)-Regular Code and Greedy Algorithm CHT]: We choose $\omega_2 = 1$ and let $\omega_3 = \eta$ where η is some very small quantity. Solving the system of differential equations reveals that the resulting gap is equal to $0.017n$. We see that this is the smallest expected gap for all the presented algorithms. $\square$

D. A Practical Greedy Algorithm

In practice, one implements a serial version of greedy algorithm CHT. At each stage, if the residual graph has a degree-one variable node then diagonal extension is applied. If no such degree-one variable node exists then one selects a variable node of lowest possible degree, i say, from the residual graph, and declares $i - 1$ (assuming no multiple edges) of its neighbors to be known. The residual graph now has at least one degree-one node and diagonal extension is applied.

There are many practical concerns. For example, variable nodes which are used in the diagonal extension step correspond to nonsystematic variables. Typically, degree-two nodes have the highest bit-error rates. Thus, it is preferable to use as many low-degree variables in the diagonalization step as possible, e.g., if the subgraph induced by only the degree-two variables has no loops then all degree-two variables can be made nonsystematic using the above algorithm.

IV. CODES WITH LINEAR ENCODING COMPLEXITY

We saw in the preceding section that degree distributions giving rise to codes that allow transmission close to capacity will have gaps that are smaller than an arbitrarily small linear fraction of the length of the code. To prove that these codes have linear encoding complexity more work is needed, namely, one has to show that the gap g satisfies $g \leq O(\sqrt{n})$ with high probability for large enough n . More precisely, we will prove the following.

Theorem 1 (Codes with Linear Encoding Complexity): Let (λ, ρ) be a degree distribution pair satisfying $\alpha^*(\rho, \lambda) = 1$, with minimum right degree $\underline{d}_c > 2$, and satisfying the strict inequality $\lambda'(0)\rho'(1) > 1$. Let G be chosen at random from the ensemble $\mathcal{C}^n(\lambda, \rho)$. Then G is encodable in linear time with probability at least $1 - bc\sqrt{n}$ for some positive constants b and c , where $c < 1$.

Discussion: We note that all optimized degree distribution pairs listed in [12] fulfill the conditions of Theorem 1. Furthermore, in experiments when applying the practical greedy algorithm to graphs based on these degree distribution pairs, the resulting gap is typically in the range of one to three! This is true even for very large lengths like one million. By correctly choosing the first degree-two variable, the gap can nearly always be lowered to one. The primary reason for these very small gaps is the large number of degree-two variable nodes in these degree distributions. The number of degree-two variable nodes is suff

ficiently large so that, with very high probability, the subgraph induced by these nodes has a large (linear size) connected component. Once a single check node belonging to this component is declared known then the remainder of the component will diagonalize in the next diagonal extension step. The diagonalization process then typically completes without further increasing the gap. $\square$

Proof: In order to show that under the stated conditions elements of the ensembles $\mathcal{C}^n(\lambda, \rho)$ are linear time encodable (with high probability) it suffices to show that their corresponding $\hat{H}^T$ can be brought into approximate lower triangular form with a gap of no more than $O(\sqrt{n})$ (with high probability). Note that we are working on the *transpose* of the parity-check matrix. Although one can prove that such an approximate triangulation is achieved by the *practical greedy algorithm* it will be more convenient to consider a slightly different greedy algorithm.⁴ The algorithm we consider has three phases which we will have to investigate separately: *startup*, *main triangulation*, and *cleanup*. In the startup phase, we will declare at most $O(\sqrt{n})$ of the check nodes to be known. Each time we declare one check node to be known we apply the diagonal extension step repeatedly until either there are no degree one variable nodes left or until (we hope) the number of degree-one variable nodes has grown to a linear-sized fraction. Assuming the latter, we then enter the *main triangulation* process. With exponential probability, the process will continue until we are left with at most a small linear fraction of nodes. Now we enter the *cleanup* phase. Here, we will show that with high probability at most $O(\sqrt{n})$ check nodes will be left when the algorithm terminates. So overall, with high probability the gap will be no more than $O(\sqrt{n})$, which will prove the claim. We will now discuss these three phases in detail.

Recall that our aim is to bring a given $\hat{H}^T$, where $\hat{H}$ is a random element from $\mathcal{C}^n(\lambda, \rho)$, into approximate lower triangular form with gap at most $O(\sqrt{n})$ by applying a greedy algorithm.

Startup: Let v be a randomly chosen degree-two variable node and let c_1 and c_2 be its connected check nodes. Declare c_1 to be *known*. Now perform the diagonal extension step. After this step, the columns which correspond to c_1 and c_2 will form the first two columns of the matrix (assuming v does not have a double edge) and the row corresponding to v will form the first row of the matrix. Consider the residual matrix (with the first two columns and the first row deleted) and the corresponding residual graph. If this residual matrix contains a degree-one row then we can apply another diagonal extension step and so on. It will simplify our description if we perform the diagonal extension step to *one degree-one variable node at a time*, instead of to all degree-one variable nodes in parallel. More precisely, we start out with one degree-two variable node which we convert into a degree-one variable node by declaring one of its neighbors to be known. Then, at any stage of the procedure, choose one of the degree-one variable

nodes (assuming that at least one such node exists) and perform the diagonal extension step only on this variable.

Let Y_t denote the number of degree-one variable nodes after the t th such step, where we have $Y_0 = 1$. If by X_t , we denote the number of *additional* degree-one variable nodes which are generated in the t th step then we get

$$Y_t = Y_{t-1} + X_t - 1 \quad (17)$$

where the -1 term stems from the fact that one degree-one variable node is used up during the diagonal extension step. Equation (17) is an instance of a *branching process*. Note that the process continues until $Y_t = 0$, i.e., until there are no more degree-one variable nodes available for the diagonal extension step. We would like the process to continue until Y_t has reached “linear size,” i.e., until Y_t is a small fixed fraction of the number of variable nodes.

Assume that we have performed at most ϵn steps. Let $(\lambda^{\text{res}}, \rho^{\text{res}})$ denote the *residual* degree distribution pair. If ϵ is small, it is intuitively clear that $(\lambda^{\text{res}}, \rho^{\text{res}})$ is “close” to (λ, ρ) . Indeed, in Lemma 4 in Appendix A it is shown that, given a degree distribution pair (λ, ρ) such that $\lambda_2 \rho'(1) > 1$, then there exists an $\epsilon > 0$ and a $\gamma > 1$ such that $\lambda_2^{\text{res}} \rho^{\text{res}'}(1) \geq \gamma > 1$ regardless which check nodes have been removed, as long as their total number is no more than ϵn .

So, assume that we have performed at most ϵn steps. What is the expected value of X_t ? Consider an edge e emanating from a degree-one variable node. With probability ρ_j^{res} it is connected to a degree- j check node, call this node c . This check node has $(j-1)$ other edges, each of which has probability λ_2^{res} of being connected to a degree-two node. Therefore, if c has degree j then the expected number of *new* degree-one nodes that will be generated is equal to $(j-1)\lambda_2^{\text{res}}$. Averaging over all degrees we get that X_t has expected value $\lambda_2^{\text{res}} \rho^{\text{res}'}(1) \geq \gamma > 1$. In other words, we have $\mathbb{E}[X_t | X_1, \dots, X_{t-1}] \geq \gamma$ for $1 \leq t \leq \epsilon n$. Furthermore, X_t is upper-bounded by the maximum right degree $\bar{d}_c$.

Let us define $T := \min_t \{Y_t = 0\}$ to be the stopping time of Y_t . We will say that the branching process *stops prematurely* if $T \leq \sqrt{n}$ and we will say that it is *successful* if $T \geq \frac{\epsilon}{2} n$ and $Y_{\frac{\epsilon}{2} n} \geq \eta \frac{\epsilon}{2} n$, where η can be chosen freely in the range $0 < \eta < 1 - \gamma$. Assume now that we employ the following strategy. Start a process, by choosing a degree-two variable node and declaring one of its neighbors to be known. If this process stops prematurely then start another process if the number of prematurely stopped processes so far is less than $\frac{\epsilon}{2} \sqrt{n}$ or declare a failure otherwise. If the current process has not stopped prematurely then declare a success if $T \geq \frac{\epsilon}{2} n$ and $Y_{\frac{\epsilon}{2} n} \geq \eta \frac{\epsilon}{2} n$ and stop the process at that time, and declare a failure otherwise. Note that the total number of steps taken for this strategy is at most $\sqrt{n} \frac{\epsilon}{2} \sqrt{n} + \frac{\epsilon}{2} n = \epsilon n$. Although the branching process which we consider always stops at a finite time and although we will only be interested in the process for at most ϵn steps it is convenient to think of an infinite process $Y_0, Y_1, \dots$ with the property that

$$1 < \gamma \leq \mathbb{E}[X_t | X_1, \dots, X_{t-1}] < \bar{d}_c, \quad l \geq 1.$$

This will allow us to write statements like

$$\Pr\{T < \sqrt{n}\} \leq \Pr\{T < \infty\}.$$

Ex. 1022.14

⁴In Appendix C, we define the notion of a “stopping set.” Stopping sets determine the termination points of diagonal extension steps regardless of the implementation of the diagonal extension. Thus, the particular three-phase formulation used here is only for convenience of presentation.

The probability of failure can be easily bounded in the following way. From iv) Lemma 5 in Appendix B, the probability that a process has stopping time less than $\sqrt{n}$, i.e., $\Pr\{T < \sqrt{n}\}$, can be upper-bounded by

$$\Pr\{T < \sqrt{n}\} \leq \Pr\{T < \infty\} \leq p(\gamma, \bar{d}_c) < 1.$$

Therefore, the probability that $\frac{\epsilon}{2} \sqrt{n}$ processes have stopping time less than $\sqrt{n}$ has an upper bound of the form $c_1 \sqrt{n}$, $c_1 < 1$. Using i), ii), and iii) of Lemma 5, the probability that a process failed assuming that it did not stop prematurely can be upper-bounded as follows:

$$\begin{aligned} & \Pr\left\{\left\{T \leq \frac{\epsilon}{2} n\right\} \cup \left\{Y_{\frac{\epsilon}{2} n} \leq \eta_{\frac{\epsilon}{2} n}\right\} \mid T \geq \sqrt{n}\right\} \\ & \leq \Pr\left\{T \leq \frac{\epsilon}{2} n \mid T \geq \sqrt{n}\right\} + \Pr\left\{Y_{\frac{\epsilon}{2} n} \leq \eta_{\frac{\epsilon}{2} n} \mid T \geq \sqrt{n}\right\} \\ & \leq \Pr\{T < \infty \mid T \geq \sqrt{n}\} + \frac{\Pr\{Y_{\frac{\epsilon}{2} n} \leq \eta_{\frac{\epsilon}{2} n}\}}{\Pr\{T \geq \sqrt{n}\}} \\ & \leq \Pr\{T < \infty \mid T \geq \sqrt{n}\} + \frac{\Pr\{Y_{\frac{\epsilon}{2} n} \leq \eta_{\frac{\epsilon}{2} n}\}}{\Pr\{T = \infty\}} \\ & \leq b_1 \xi^{\sqrt{n}} + b_2 c_2^n \\ & \leq b_3 \xi^{\sqrt{n}} \end{aligned}$$

for some constants b_1, b_2, b_3 , and $c_2 < 1$ for some $\gamma > 1$, and a constant $\pi < 1$ defined in Appendix B. Combining these two results, we conclude that the probability of failure is upper-bounded by $b_4 c_3^{\sqrt{n}}$ for some constant b_4 and $c_3 < 1$.

Main Upper Triangulation: With high probability we will have succeeded in the startup phase. Consider now the output of this startup phase assuming it was successful. From Lemma 4 in Appendix A we know that the residual degree distribution pair fulfills $\underline{d}_c^{\text{res}} > 2$ and $\alpha^*(\rho^{\text{res}}, \lambda^{\text{res}}) = 1$. Furthermore, it is easy to see that conditioned, on $(\lambda^{\text{res}}, \rho^{\text{res}})$ and the length n^{res} , the resulting residual graph can be thought of as an element of $\mathcal{C}^{n^{\text{res}}}(\lambda^{\text{res}}, \rho^{\text{res}})$. This is most easily seen as follows: Given the original graph, the residual graph is the result of removing a certain set of edges, where the choice of these edges is the result of certain random experiments. Consider another element of $\mathcal{C}^n(\lambda, \rho)$ which agrees with the original graph in those edges but is arbitrary otherwise. Assume now that we run the startup phase on this new graph with the same random choices. It is easy to see that the sequence of degree distribution pairs will be the same and that at each step the probability for the given choice of node which gets chosen is identical. So the resulting residual graphs will have identical degree distribution pairs, identical length, and the same probability of being generated. Further, each element of $\mathcal{C}^{n^{\text{res}}}(\lambda^{\text{res}}, \rho^{\text{res}})$ is reachable and, by the above discussion, they have equal probability of being generated, which proves the claim. Since $\lambda_1^{\text{res}} > 0$ it follows that we can now simply use greedy algorithm AHT to continue the lower triangulation process. From the analysis of greedy algorithm AHT we know that, with exponential probability, the process will not stop until at most a small linear fraction σ of check nodes is left, where this fraction σ can be made as small as desired.

Cleanup Operation: So far we have increased the gap to at most $\frac{\epsilon}{2} \sqrt{n}$ and, with high probability, we have accomplished

a partial lower triangulation with at most a small fraction σ of check nodes left. In Lemma 6 in Appendix C it is now shown that the probability at $1 - \frac{1}{2} \sqrt{n}$ actually fewer than $\sqrt{n}$ check nodes will be left.

Combining all these statements we see that the probability that the gap exceeds $2\sqrt{n}$ is at most $bc^{\sqrt{n}}$, where $c < 1$. $\square$

APPENDIX A

RESIDUAL DEGREE DISTRIBUTION PAIRS

Lemma 4: Let (λ, ρ) be a degree distribution pair satisfying $\underline{d}_c > 2$, the strict inequality $\lambda_2 \rho'(1) > 1$ and $\alpha^*(\rho, \lambda) = 1$. Let B_ϵ be the set of all residual degree distribution pairs obtainable from (λ, ρ) by removing at most an ϵ fraction of check nodes from a graph with degree distribution pair (λ, ρ) . Then, for ϵ sufficiently small, there exists a $\gamma > 1$ such that any $(\lambda^{\text{res}}, \rho^{\text{res}}) \in B_\epsilon$ will satisfy $\underline{d}_c^{\text{res}} > 2$ and the strict inequality $\lambda_2^{\text{res}} \rho^{\text{res}'}(1) > \gamma$. If, moreover, for some $\alpha > 0$ we have $\lambda_1^{\text{res}} \geq \alpha\epsilon$ then $\alpha^*(\rho^{\text{res}}, \lambda^{\text{res}}) = 1$.

Proof: The conclusion $\underline{d}_c^{\text{res}} > 2$ is immediate since $\underline{d}_c > 2$ and since we either remove a check node completely or leave its degree unchanged. By continuity and since, for some $\gamma > 1$, $\lambda_2 \rho'(1) > \gamma$, it is also clear that $\lambda_2^{\text{res}} \rho^{\text{res}'}(1) > \gamma$ if ϵ is sufficiently small.

It remains to show that $\alpha^*(\rho^{\text{res}}, \lambda^{\text{res}}) = 1$ if $\lambda_1^{\text{res}} \geq \alpha\epsilon$ for some $\eta > 0$. Let $\Delta := \max\{\bar{d}_v, \bar{d}_c\}$. Let δ be a positive number such that $\lambda_2 \rho'(1) > 1 + \delta$. It follows by continuity that for ϵ small enough we have

$$\lambda_2^{\text{res}} \rho^{\text{res}'}(1 - \lambda_1^{\text{res}}) > 1 + \delta/2.$$

Define $\nu(x) := \rho^{\text{res}}(1 - \lambda^{\text{res}}(x))$ and note that $\sum |\nu_i| \leq 2^\Delta$. Since $\rho^{\text{res}'}(1) \geq 2$ we have $\rho^{\text{res}}(1 - \lambda_1^{\text{res}}) \leq 1 - \eta\epsilon$ for ϵ sufficiently small. Hence, for ϵ sufficiently small, we have

$$\begin{aligned} & \rho^{\text{res}}(1 - \lambda^{\text{res}}(x)) \\ & \leq \rho^{\text{res}}(1 - \lambda_1^{\text{res}}) - \rho^{\text{res}'}(1 - \lambda_1^{\text{res}}) \lambda_2^{\text{res}} x + 2^\Delta x^2 \\ & \leq 1 - \eta\epsilon - x(1 + \delta/2 - 2^\Delta x) \\ & \leq 1 - x - \eta\epsilon \quad \forall x \in [0, \delta/2^{\Delta+1}). \end{aligned}$$

In a similar manner

$$\begin{aligned} & \rho^{\text{res}}(1 - \lambda^{\text{res}}(x)) \leq \rho_2^{\text{res}} \lambda^{\text{res}'}(1)(1 - x) + 2^\Delta (1 - x)^2 \\ & \leq 2^\Delta (1 - x)^2 \\ & < 1 - x \quad \forall x \in (1 - 2^{-\Delta}, 1). \end{aligned}$$

In the compact range $[\delta/2^{\Delta+1}, 1 - 2^{-(\Delta+1)}]$, $\rho^{\text{res}}(1 - \lambda^{\text{res}}(x))$ is a continuous function in the perturbation and since the degree distribution pair (λ, ρ) fulfills the strict inequality (14) in this range, it follows that there exists an ϵ_2 such that if $\epsilon \leq \epsilon_2$ then $\rho^{\text{res}}(1 - \lambda(x)) < 1 - x - \epsilon$ in this range. Let us further assume $\epsilon < 2^{-(\Delta+2)}$. Then it follows that on the interval $[1 - 2^{-(\Delta+1)}, 1 - 2\epsilon]$ we have

$$2^\Delta (1 - x)^2 - (1 - x) + \epsilon < 0$$

and hence $\rho^{\text{res}}(1 - \lambda(x)) < 1 - x - \epsilon$. This shows that $\alpha^*(\rho^{\text{res}}, \lambda^{\text{res}}) = 1$.

APPENDIX B BRANCHING PROCESSES

Let $X_1, X_2, \dots$ be a sequence of independent and identically distributed (i.i.d.) random variables. Define the sequence $Y_0, Y_1, \dots$ by the recursion $Y_0 := 1$ and $Y_t := Y_{t-1} + X_t - 1$ and let T be the least time t such that $Y_t = 0$. If no such time t exists then we define $T = \infty$. The process $Y_0, Y_1, \dots$ is usually referred to as a *branching process* and originated in the context of the study of population growths [18]. It is well known that if $\mathbb{E}[X_t] < 1$ then $\Pr\{T = \infty\} = 0$ but that if $\mathbb{E}[X_t] > 1$ then $\Pr\{T = \infty\} > 0$. We will now show that under suitable conditions the same conclusions hold even if we allow (limited) dependency within the process $X_1, X_2, \dots$

Lemma 5 (Branching Processes): Let $X_1, X_2, \dots$, be a sequence of random variables taking values in $\{0, \dots, d\}$ such that

$$1 < \gamma \leq \mathbb{E}[X_t | X_1, \dots, X_{t-1}] \leq d$$

for all $t \geq 1$. Define the *branching process* $Y_0, Y_1, \dots$, by $Y_0 := 1$ and $Y_t := Y_{t-1} + X_t - 1$, $t \geq 1$. Let the *stopping time* T be defined by $T := \min_{t \geq 1} \{t : Y_t = 0\}$, where $T := \infty$ if no such t exists.

i) For any $\alpha \in (0, 1)$

$$\Pr\{Y_t \leq (\alpha\gamma - 1)t\} \leq e^{-t^{\frac{\gamma}{d}}(1-\alpha(1-\ln \alpha))}.$$

ii) Define $\xi := \xi(\gamma, d) := e^{-\frac{1}{d}(\gamma-1-\ln \gamma)} < 1$. Then

$$\Pr\{T < \infty | T \geq t\} \leq \frac{\xi^t}{1-\xi}.$$

iii) For any $t \geq 0$

$$\Pr\{T > t\} > \left(\frac{\gamma}{d}\right)^t.$$

iv) Define

$$p := p(\gamma, d) := \min_{t \geq 1} 1 - \frac{1-\xi-\xi^t}{1-\xi} \left(\frac{\gamma}{d}\right)^t < 1.$$

Then

$$\Pr\{T < \infty\} \leq p.$$

Proof: We start with a proof of the tail inequality. For any $s > 0$ we have

$$\begin{aligned} \Pr\{Y_t \leq (\alpha\gamma - 1)t\} &\leq \Pr\left\{\sum_{i=1}^t X_i \leq \alpha\gamma t\right\} \\ &= \Pr\left\{-\sum_{i=1}^t X_i \geq -\alpha\gamma t\right\} \\ &= \Pr\left\{\exp\left(-s \sum_{i=1}^t X_i\right) \geq e^{-s\alpha\gamma t}\right\} \\ &\leq e^{s\alpha\gamma t} \mathbb{E}\left[\exp\left(-s \sum_{i=1}^t X_i\right)\right] \end{aligned}$$

where the last step follows from the well-known Markov inequality. We proceed by bounding $\mathbb{E}[\exp(-s \sum_{i=1}^t X_i)]$.

Recall the following basic fact from the theory of linear programming, [19]. A *primal* linear program in *normal form*, $\min\{cu | u \geq 0; Au \geq b\}$, has the associated *dual* linear program $\max\{vb | v \geq 0; vA \leq c\}$. Further, if u^* and v^* are feasible solutions for the primal and dual linear program such that their *values* are equal, i.e., $cu^* = v^*b$, then u^* and v^* are *optimal* solutions.

Now note that

$$\begin{aligned} \max \{ &\mathbb{E}[e^{-sX_i} | X_1, \dots, X_{i-1}] | \mathbb{E}[X_i | X_1, \dots, X_{i-1}] \\ &\geq \gamma; X_i \in \{0, \dots, d\} \} \\ &\leq \max_{u_k} \left\{ \sum_{k=0}^d e^{-sk} u_k | u_k \geq 0; \sum_{k=0}^d k u_k \geq \gamma; \right. \\ &\quad \left. \sum_{k=0}^d u_k = 1 \right\} \\ &= -\min_{u_k} \left\{ -\sum_{k=0}^d e^{-sk} u_k | u_k \geq 0; \sum_{k=1}^d k u_k \geq \gamma; \right. \\ &\quad \left. -\sum_{k=0}^d u_k \geq -1 \right\}. \end{aligned}$$

The last step warrants some remarks. First, we rewrote the maximization as a minimization to bring the linear program into normal form. Second, a simple scaling argument shows that one can replace the equality condition $\sum_{k=0}^d u_k = 1$ with the inequality $\sum_{k=0}^d u_k \leq 1$ without changing the value of the linear program. The linear program in the last line is our primal problem. It is easy to check that this primal problem has the feasible solution

$$u_0 = 1 - \frac{\gamma}{d}; u_k = 0, 1 \leq k \leq d-1; u_d = \frac{\gamma}{d}$$

with value

$$-\left(1 - \frac{\gamma}{d}(1 - e^{-sd})\right). \quad (18)$$

To see that this solution is optimal consider the associated dual program

$$\max\{v_1\gamma - v_2 | v_1 \geq 0; v_2 \geq 0; kv_1 - v_2 \leq -e^{-sk}, 0 \leq k \leq d\}.$$

The solution $(v_1, v_2) = (\frac{1}{d}(1 - e^{-sd}), 1)$ gives rise to the same value as in (18). Hence, to prove optimality it suffices to prove that this solution is feasible. For this we need to show that

$$\frac{k}{d}(1 - e^{-sd}) - 1 \leq -e^{-sk}, \quad 0 \leq k \leq d.$$

This is trivially true for $k = 0$ and for $k \geq 1$ this is equivalent to

$$\frac{1}{sd}(1 - e^{-sd}) \leq \frac{1}{sk}(1 - e^{-sk}), \quad 1 \leq k \leq d.$$

The claim now follows since $\frac{1-e^{-x}}{x}$ is a decreasing function in x for $x \geq 0$.

Ex. 1022.16

We get

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-s \sum_{i=1}^t X_i \right) \right] \\ &= \mathbb{E} \left[\exp \left(-s \sum_{i=1}^{t-1} X_i \right) \mathbb{E} [e^{-sX_t} | X_1, \dots, X_{t-1}] \right] \\ &\leq \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right) \mathbb{E} \left[\exp \left(-s \sum_{i=1}^{t-1} X_i \right) \right] \\ &\leq \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right)^t. \end{aligned}$$

It follows that

$$\begin{aligned} \Pr \left\{ \sum_{i=1}^t X_i \leq \alpha \gamma t \right\} &\leq e^{s\alpha \gamma t} \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right)^t \\ &= \exp \left(s\alpha \gamma t + t \ln \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right) \right) \\ &\leq \exp \left(-t \frac{\gamma}{d}(1 - e^{-sd} - sd\alpha) \right). \end{aligned}$$

The desired inequality now follows by choosing $s = \frac{1}{d} \ln \frac{1}{\alpha}$.

Next we show that

$$\Pr\{T < \infty | T \geq t\} \leq \frac{\pi^t}{1 - \pi}.$$

The proof will be very similar to the preceding one and we will be short.

$$\begin{aligned} & \Pr\{T < \infty | T \geq t\} \\ &= \Pr \left\{ \bigcup_{i \geq 1} \{Y_i = 0\} \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &\leq \sum_{i=t}^{\infty} \Pr \left\{ Y_i = 0 \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &= \sum_{i=t}^{\infty} \Pr \left\{ \sum_{k=1}^i X_k = (i-1) \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &\leq \sum_{i=t}^{\infty} \Pr \left\{ \sum_{k=1}^i X_k \leq i \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &= \sum_{i=t}^{\infty} \Pr \left\{ -\sum_{k=1}^i X_k \geq -i \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &= \sum_{i=t}^{\infty} \Pr \left\{ \exp \left(-s \sum_{k=1}^i X_k \right) \geq e^{-si} \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right\} \\ &\leq \sum_{i=t}^{\infty} e^{si} \mathbb{E} \left[\exp \left(-s \sum_{k=1}^i X_k \right) \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right]. \end{aligned}$$

Now for $i \geq t$ we have

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-s \sum_{k=1}^i X_k \right) \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right] \\ &\leq \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right)^{i-t+1} \end{aligned}$$

$$\begin{aligned} & \cdot \mathbb{E} \left[\exp \left(-s \sum_{k=1}^{t-1} X_k \right) \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right] \\ &\leq \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right)^{i-t+1} \mathbb{E} \left[\exp \left(-s \sum_{k=1}^{t-2} X_k \right) \right. \\ &\quad \cdot \mathbb{E} \left[e^{-sX_{t-1}} \middle| X_1, \dots, \right. \\ &\quad \left. \left. X_{t-2}, X_{t-1} \geq (t-1) - \sum_{l=1}^{t-2} X_l \right] \middle| \bigcap_{j=1}^{t-3} \{Y_j \geq 1\} \right]. \end{aligned}$$

But

$$\begin{aligned} & \max_{u_k} \left\{ \mathbb{E} \left[e^{-sX_i} \middle| X_1, \dots, X_{i-1}, X_i \geq i - \sum_{l=1}^{i-1} X_l \right] \right. \\ &\quad \cdot \mathbb{E}[X_i | X_1, \dots, X_{i-1}] \geq \gamma; X_i \in \{0, \dots, d\} \left. \right\} \\ &\leq \max_{c, u_k} \left\{ \frac{\sum_{k=c}^d e^{-sk} u_k}{\sum_{k=c}^d u_k} \middle| u_k \geq 0; \sum_{k=0}^d k u_k \geq \gamma; \sum_{k=0}^d u_k = 1 \right\} \\ &\leq \max_{c, u_k} \left\{ \sum_{k=c}^d e^{-sk} u_k \middle| u_k \geq 0; \sum_{k=c}^d k u_k \geq \gamma; \sum_{k=c}^d u_k = 1 \right\} \\ &\leq \max_{u_k} \left\{ \sum_{k=0}^d e^{-sk} u_k \middle| u_k \geq 0; \sum_{k=0}^d k u_k \geq \gamma; \sum_{k=0}^d u_k = 1 \right\} \\ &= -\min_{u_k} \left\{ -\sum_{k=0}^d e^{-sk} u_k \middle| u_k \geq 0; \sum_{k=0}^d k u_k \geq \gamma; \right. \\ &\quad \left. -\sum_{k=0}^d u_k \geq -1 \right\} \\ &= 1 - \frac{\gamma}{d}(1 - e^{-sd}) \end{aligned}$$

so that

$$\begin{aligned} & \mathbb{E} \left[\exp \left(-s \sum_{k=1}^i X_k \right) \middle| \bigcap_{j=1}^{t-1} \{Y_j \geq 1\} \right] \\ &\leq \left(1 - \frac{\gamma}{d}(1 - e^{-sd}) \right)^i = \pi^i. \end{aligned}$$

It follows that

$$\begin{aligned} \Pr\{T < \infty | T \geq t\} &\leq \sum_{i=t}^{\infty} \xi^i \\ &= \frac{\xi^t}{1 - \xi}. \end{aligned}$$

To prove that for any $t \geq 0$, $\Pr\{T > t\} > 0$ note that from $\mathbb{E}[X_i | X_1, \dots, X_{i-1}] \geq \gamma$ and $X_i \in \{0, \dots, d\}$ we conclude that $\Pr\{X_i \geq 1 | X_1, \dots, X_{i-1}\} \geq \frac{\gamma}{d}$. Therefore

$$\begin{aligned} \Pr\{T > t\} &= \Pr\{Y_1 \geq 1, \dots, Y_t \geq 1\} \\ &\geq \Pr\{X_1 \geq 1, \dots, X_t \geq 1\} \\ &= \prod_{i=1}^t \Pr\{X_i \geq 1 | X_1 \geq 1, \dots, X_{i-1} \geq 1\} \\ &\geq \left(\frac{\gamma}{d} \right)^t. \end{aligned}$$

Ex. 1022.17

It remains to prove that $\Pr\{T < \infty\} \leq p$. Note that for any $t \geq 1$ we have

$$\begin{aligned} \Pr\{T < \infty\} &= 1 - \Pr\{T = \infty\} \\ &= 1 - \Pr\{T = \infty | T > t\} \Pr\{T > t\} \\ &\quad - \Pr\{T = \infty | T \leq t\} \Pr\{T \leq t\} \\ &= 1 - \Pr\{T = \infty | T > t\} \Pr\{T > t\} \\ &\leq 1 - \frac{1 - \xi - \xi^t}{1 - \xi} \left(\frac{\gamma}{d}\right)^t. \end{aligned} \quad \square$$

APPENDIX C

NO SMALL STOPPING SETS

Given a bipartite graph G with constraint set C we define a *stopping set* S to be any subset of C with the property that there does not exist a variable node in G with exactly one edge into S . The union of any two stopping sets is a stopping set. Hence, there exists a unique maximal stopping set. Thus, if the graph G is operated on by the diagonal extension stage of an approximate upper triangulation algorithm which looks for degree-one variable nodes, then it always terminates with the residual graph determined by the maximal stopping set. That is, after a diagonal extension stage, the residual graph consists of the constraint nodes in S , the maximal stopping set, the edges emanating from S and the variable nodes incident to these edges. Thus, one can show that the diagonal extension step will not terminate prematurely if one can show that there are no stopping sets.

Lemma 6 (No Small Stopping Sets): Let $\mathcal{C}^n(\lambda, \rho)$ be an ensemble of LDPC codes with $\underline{d}_c > 2$. Then there exist a positive number α and a natural number $n(\alpha)$ such that for all $n \geq n(\alpha)$, a randomly chosen element of $\mathcal{C}^n(\lambda, \rho)$ has probability at most $2^{-\sqrt{n}}$ of containing a stopping set of size in the range $[\sqrt{n}, \alpha n]$.

Proof: Recall the following simple estimates which we will use frequently in the sequel:

$$\frac{1}{n+1} 2^{nh(\frac{k}{n})} \leq \binom{n}{k} \leq 2^{nh(\frac{k}{n})}$$

as well as the fact that $h(x) = x \log \frac{1}{x} + O(x)$, for $0 \leq x \leq 1$.

Recall that the total number of edges is equal to nd_v^{avg} . Consider e edges. More precisely, fix e check node sockets and for each such check node socket choose a variable node socket at random. Clearly, this can be done in

$$\binom{nd_v^{\text{avg}}}{e} e! \geq \frac{1}{n+1} 2^{nd_v^{\text{avg}} h(\frac{e}{nd_v^{\text{avg}}})} e! = 2^{nh(\frac{e}{n}) + O(e)} e!$$

ways. We will say that the e edges are *doubly connected* if each variable is either not connected at all or connected at least twice (with respect to these e edges). We claim that there are at most

$$\binom{n}{e/2} \left(\frac{e}{2\bar{d}_v}\right) e! \leq 2^{n(h(\frac{e}{2n}) + \frac{e\bar{d}_v}{2n} h(\frac{2}{\bar{d}_v}))} e! = 2^{n\frac{1}{2}h(\frac{e}{n}) + O(e)} e!$$

doubly connected constellations. To see this claim note that a doubly connected constellation with e edges involves at most $e/2$ variable nodes. Therefore, we get an upper bound if we

count as follows: first choose $e/2$ out of the n variable nodes. These chosen variable nodes have at most $e/2\bar{d}_v$ sockets. Choose e of these sockets and connect to them the edges in any order. From this argument we see that the probability that the chosen set of edges is doubly connected is upper-bounded by

$$2^{-(n\frac{1}{2}h(\frac{e}{n}) + O(e))}. \quad (19)$$

Note that for sufficiently small e , (19) is *decreasing* in e .

Now consider a set of c check nodes. There are

$$\binom{n(1-r)}{c} \leq 2^{n(1-r)h(\frac{c}{n(1-r)})} = 2^{nh(\frac{c}{n}) + O(c)}$$

such sets. Each such set has at least $c\underline{d}_c$ edges and therefore, assuming that c is sufficiently small, we see from (19) that the probability that one such set is a stopping set is at most

$$2^{-(n\frac{d_c}{2}h(\frac{c}{n}) + O(c))}.$$

By the union bound it follows that the probability that a randomly chosen element of $\mathcal{C}^n(\lambda, \rho)$ does have a stopping set of size c is upper-bounded by

$$2^{-(\frac{\underline{d}_c-2}{2}nh(\frac{c}{n}) + O(c))} = 2^{-f(n, c, \underline{d}_c)}$$

where we defined

$$f(n, c, \underline{d}_c) := \frac{\underline{d}_c-2}{2} nh\left(\frac{c}{n}\right) + O(c).$$

It remains to show that there exist constants $\alpha > 0$ and $n(\alpha) > 0$ such that for all $c \in [\sqrt{n}, \alpha n]$ and for all $n \geq n(\alpha)$ we have $f(n, c, \underline{d}_c) \geq \sqrt{n}$. Recall that $\underline{d}_c > 2$. Therefore, if $\alpha > 0$ and $\beta \in (\frac{1}{2}, 1)$ then

$$\begin{aligned} f(n, \alpha n^\beta, \underline{d}_c) &= \frac{\underline{d}_c-2}{2} \alpha n^\beta \left(\log_2 \left(\frac{n^{1-\beta}}{\alpha} \right) + O(1) \right) \\ &\geq \frac{\underline{d}_c-2}{2} \alpha n^\beta \\ &\geq \sqrt{n} \end{aligned}$$

where the second step is true if we choose α small enough and the third step is true for sufficiently large n . $\square$

REFERENCES

- [1] R. G. Gallager, *Low-Density Parity-Check Codes*. Cambridge, MA: MIT Press, 1963. Available at <http://justice.mit.edu/people/gallager.html>.
- [2] V. Zyablov and M. Pinsker, "Estimation of the error-correction complexity of Gallager low-density codes," *Probl. Pered. Inform.*, vol. 11, pp. 23–26, Jan. 1975.
- [3] G. A. Margulis, "Explicit construction of graphs without short cycles and low density codes," *Combinatorica*, vol. 2, no. 1, pp. 71–78, 1982.
- [4] R. Tanner, "A recursive approach to low complexity codes," *IEEE Trans. Inform. Theory*, vol. IT-27, pp. 533–547, Sept. 1981.
- [5] D. J. C. MacKay and R. M. Neal, "Near Shannon limit performance of low density parity check codes," *Electron. Lett.*, vol. 32, pp. 1645–1646, Aug. 1996.
- [6] N. Wiberg, "Codes and decoding on general graphs," Dissertation no. 440, Dept. Elect. Eng. Linköping Univ., Linköping, Sweden, 1996.

- [7] N. Sourlas, "Spin-glass models as error-correcting codes," *Nature*, no. 339, pp. 693–695, 1989.
- [8] I. Kanter and D. Saad, "Error-correcting codes that nearly saturate Shannon's bound," *Phys. Rev. Lett.*, vol. 83, pp. 2660–2663, 1999.
- [9] M. Luby, M. Mitzenmacher, A. Shokrollahi, D. Spielman, and V. Stemann, "Practical loss-resilient codes," in *Proc. 29th Annual ACM Symp. Theory of Computing*, 1997, pp. 150–159.
- [10] M. Luby, M. Mitzenmacher, A. Shokrollahi, and D. Spielman, "Analysis of low density codes and improved designs using irregular graphs," in *Proc. 30th Annu. ACM Symp. Theory of Computing*, 1998, pp. 249–258.
- [11] T. Richardson and R. Urbanke, "The capacity of low-density parity-check codes under message-passing decoding," *IEEE Trans. Inform. Theory*, vol. 47, pp. 599–618, Feb. 2001.
- [12] T. Richardson, A. Shokrollahi, and R. Urbanke, "Design of capacity-approaching low-density parity check codes," *IEEE Trans. Inform. Theory*, vol. 47, pp. 619–637, Feb. 2001.
- [13] M. Sipser and D. Spielman, "Expander codes," *IEEE Trans. Inform. Theory*, vol. 42, pp. 1710–1722, Nov. 1996.
- [14] D. J. C. MacKay, S. T. Wilson, and M. C. Davey, "Comparison of constructions of irregular Gallager codes," in *Proc. 36th Allerton Conf. Communication, Control, and Computing*, Sept. 1998.
- [15] D. Spielman, "Linear-time encodeable and decodable error-correcting codes," *IEEE Trans. Inform. Theory*, vol. 42, pp. 1723–1731, Nov. 1996.
- [16] M. Luby, M. Mitzenmacher, and A. Shokrollahi, "Analysis of random processes via and-or tree evaluation," in *Proc. 9th Annu. ACM-SIAM Symp. Discrete Algorithms*, 1998, pp. 364–373.
- [17] L. Bazzi, T. Richardson, and R. Urbanke, "Exact thresholds and optimal codes for the binary symmetric channel and Gallager's decoding algorithm A," *IEEE Trans. Inform. Theory*, to be published.
- [18] N. Alon, J. Spencer, and P. Erdős, *The Probabilistic Method*. New York: Wiley, 1992.
- [19] A. Schrijver, *Theory of Linear and Integer Programming*. New York: Wiley, 1986.