

Probabilistic Analysis of Linear Programming Decoding

Constantinos Daskalakis, Alexandros G. Dimakis, *Student Member, IEEE*, Richard M. Karp, and Martin J. Wainwright, *Member, IEEE*

Abstract—We initiate the probabilistic analysis of linear programming (LP) decoding of low-density parity-check (LDPC) codes. Specifically, we show that for a random LDPC code ensemble, the linear programming decoder of Feldman *et al.* succeeds in correcting a constant fraction of errors with high probability. The fraction of correctable errors guaranteed by our analysis surpasses previous nonasymptotic results for LDPC codes, and in particular, exceeds the best previous finite-length result on LP decoding by a factor greater than ten. This improvement stems in part from our analysis of probabilistic bit-flipping channels, as opposed to adversarial channels. At the core of our analysis is a novel combinatorial characterization of LP decoding success, based on the notion of a flow on the Tanner graph of the code. An interesting by-product of our analysis is to establish the existence of “probabilistic expansion” in random bipartite graphs, in which one requires only that almost every (as opposed to every) set of a certain size expands, for sets much larger than in the classical worst case setting.

Index Terms—Binary-symmetric channel (BSC), channel coding, error-control coding, expanders, factor graphs, linear programming decoding, low-density parity-check (LDPC) codes, randomized algorithms, sum-product algorithm.

I. INTRODUCTION

LOW-density parity-check (LDPC) codes are a class of sparse binary linear codes, first introduced by Gallager [15], and subsequently studied extensively by various researchers [21]–[23]. See the book by Richardson and Urbanke [24] for a comprehensive treatment of the subject. When decoded with efficient iterative algorithms (e.g., the sum-product algorithm [20]), suitably designed classes of LDPC codes yield

Manuscript received January 23, 2007. The work of C. Daskalakis was supported by the National Science Foundation under Grant CCF-0515259. The work of A. G. Dimakis was supported by the National Science Foundation under Grants DMS-0528488, CCF-0635372, and a University of California MICRO Grant with Marvell Corporation. The work of R. M. Karp was supported by the National Science Foundation under Grant CCF-0515259. The work of M. J. Wainwright was supported by an Alfred P. Sloan Foundation Fellowship, the National Science Foundation under Grants DMS-0528488, CCF-0635372, and a University of California MICRO Grant with Marvell Corporation. The material in this paper was presented in part at the SIAM Symposium on Discrete Algorithms (SODA), New Orleans, LA, January 2007. This work was performed, in part, while C. Daskalakis was visiting Microsoft Research, Redmont, WA, and A. G. Dimakis was visiting EPFL, Lausanne, Switzerland.

The authors are with the Department of Electrical Engineering and Computer Science, University of Berkeley, Berkeley, CA 94704 USA (e-mail: costos@eecs.berkeley.edu; adim@eecs.berkeley.edu; karp@cs.berkeley.edu; wainwright@eecs.berkeley.edu).

Communicated by T. J. Richardson, Associate Editor for Coding Theory.

Color versions of Figures 1 and 2 in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIT.2008.926452

error-correcting performance extremely close to the Shannon capacity of noisy channels for very large codes [6]. Most extant methods for analyzing the performance of iterative decoding algorithms for LDPC codes—notably the method of density evolution [21], [23]—are asymptotic in nature, based on exploiting the high girth of very large random graphs. Therefore, the thresholds computed using density evolution are only estimates of the true algorithm behavior, since they assume a cycle-free message history. In fact, the predictions of such methods are well known to be inaccurate for specific codes of intermediate block length (e.g., codes with a few hundreds or thousands of bits). For this reason, our current understanding of practical decoders for smaller codes, which are required for applications with delay constraints, is relatively limited.

The focus of this paper is the probabilistic analysis of linear programming (LP) decoding, a technique first introduced by Feldman *et al.* [10], [14] as an alternative to iterative algorithms for decoding LDPC codes. The underlying idea is a standard one in combinatorial optimization—namely, to solve a particular *linear programming* (LP) relaxation of the integer program corresponding to maximum-likelihood (ML) (optimal) decoding. Although the practical performance of LP decoding is comparable to message-passing decoding, a significant advantage is its relative amenability to nonasymptotic analysis. Moreover, there turn out to be a number of important theoretical connections between the LP decoding and standard forms of iterative decoding [19], [31]. These connections allow theoretical insight from the LP decoding perspective to be transferred to iterative decoding algorithms.

A. Previous Work

The technique of LP decoding was introduced for turbo-like codes [10], extended to LDPC codes [11], [14], and further studied by various researchers (e.g., [28], [5], [12], [8], [13], [16]). Significant recent interest has focused on postprocessing algorithms that use the ML-certificate property of LP decoding to achieve near ML performance (see [8], [4]) and also [9], [26]).

For concatenated expander codes, Feldman and Stein [13] showed that LP decoding can achieve capacity; see also [1], [18]. For the standard LDPC codes used in practice, the best positive result from previous work [12] is the existence of a constant $\beta > 0$, depending on the rate of the code, such that LP decoding can correct *any* bit-flipping pattern consisting of at most βn bit flips. (In short, we say that LP decoding can correct a β -fraction of errors.) As a concrete example for suitable classes of rate- $1/2$ LDPC codes, Feldman *et al.* [12] established that $\beta = 0.000177$ is achievable. However, this analysis [12] was worst case in nature, essentially assuming an adversarial

channel model. Such analysis yields overly conservative predictions for the probabilistic channel models that are of more practical interest. Consequently, an important direction—and the goal of this paper—is to develop methods for finite-length and *average-case analysis* of the LP decoding method.

B. Our Contributions

This paper initiates the average-case analysis of LP decoding for LDPC codes. In particular, we analyze the following question: what is the probability, given that a random subset of αn bits is flipped by the channel, that LP decoding succeeds in recovering correctly the transmitted codeword? As a concrete example, we prove that for bit-degree-regular LDPC codes of rate $1/2$ and a random error pattern with αn bit flips, LP decoding will recover the correct codeword, with probability converging to one,¹ for all α up to at least 0.002. This guarantee is roughly ten times higher than the best guarantee from prior work [12], derived in the setting of an adversarial channel. Our proof is based on analyzing the dual of the decoding linear program and obtaining a simple graph-theoretic condition for certifying a zero-valued solution of the dual LP, which (by strong duality) ensures that the LP decoder recovers the transmitted codeword. We show that this dual witness has an intuitive interpretation in terms of the existence of hyperflow from the flipped to the unflipped bits. Although this paper focuses exclusively on the binary-symmetric channel (BSC), the poison hyperflow is an exact characterization of LP decoding for any memoryless binary input symmetric output (MBIOS) channel. We then show that such a hyperflow witness exists with high probability under random errors in the bit-degree-regular LDPC ensemble. The argument itself entails a fairly delicate sequence of union bounds and concentration inequalities, exploiting expansion and matchings on random bipartite graphs.

C. Probabilistic Expanders

An interesting by-product of our analysis is the proof of the existence of *probabilistic expanders*—that is, bipartite graphs in which *almost all sets* of vertices of size up to αn and their *subsets* have large expansion. One key point is that it is not sufficient to require a random subset of vertices to expand with high probability (w.h.p.), because we use the expansion combined with Hall's theorem to guarantee large matchings. What we need instead is that a random subset of vertices and all its subsets will expand w.h.p. which by Hall's theorem will guarantee that a random subset will have a matching. In effect, by relaxing the expansion requirement from every set to almost all sets of a given size, we show that one can obtain much larger expansion factors, and corresponding stronger guarantees on error correction. Our analysis relies on the fact that a random bipartite graph, conditioned on all the small sets having some expansion, will also have this probabilistic expansion for much larger constants α . This innovation allows us to go beyond the worst case fraction of errors guaranteed by traditional expansion arguments [12], [25].

The remainder of the paper is organized as follows. We begin in Section II with background on error-control coding and

LDPC codes, as well as the method of LP decoding. Section III describes our main result and Section IV provides the proof in a series of lemmas, with more technical details deferred to the Appendix. We conclude in Section V with a discussion.

II. BACKGROUND AND PROBLEM FORMULATION

We begin with some background on LDPC codes. We then describe the LP decoding method, and formulate the problem to be studied in this paper.

A. LDPC Codes

The purpose of an error-correcting code is to introduce redundancy into a data sequence so as to achieve error-free communication over a noisy channel. Given a binary vector of length k (representing information to be conveyed), the encoder maps it onto a codeword, corresponding to a binary vector of length $n > k$. The code rate is given by $R = k/n$, corresponding to the ratio of information bits to transmitted bits. In a binary linear code, the set of all possible codewords corresponds to a subspace of $\{0, 1\}^n$, with a total of 2^k elements (one for each possible information sequence). The codeword is then transmitted over a noisy channel. In this paper, we focus on the *binary-symmetric channel* (BSC), in which each bit is flipped independently with probability α . Given the received sequence from the channel, the goal of the decoder is to correctly reconstruct the transmitted codeword (and hence the underlying information sequence).

Any binary linear code can be described as the null space of a parity-check matrix $H \in \{0, 1\}^{(n-k) \times n}$; more concretely, the code \mathcal{C} is given by the set of all binary strings $x \in \{0, 1\}^n$ such that $Hx = 0$ in modulo-two arithmetic. A convenient graphical representation of such a binary linear code is in terms of its factor graph [20]. The factor graph associated with a code \mathcal{C} is a bipartite graph $G = (V, C)$, with $n = |V|$ variable nodes corresponding to the codeword bits (columns of the matrix H), and $m = n - k = |C|$ nodes corresponding to the parity checks (rows of the matrix H). Edges in the factor graph connect each variable node to the parity checks that constrain it, so that the parity-check matrix H specifies the adjacency matrix of the graph. An LDPC code is a binary linear code that can be expressed with a sparse factor graph (i.e., one with $\Theta(1)$ edges per row).

Given a received sequence $\hat{y} \in \{0, 1\}^n$ from the BSC, the ML decoding problem is to determine the closest codeword (in Hamming distance). It is well known that ML decoding for general binary linear codes is NP-hard [2], which motivates the study of suboptimal but practical algorithms for decoding.

B. LP Decoding

We now describe how the problem of optimal decoding can be reformulated as a linear program over the *codeword polytope*, i.e., the convex hull of all codewords of the code \mathcal{C} . For every bit \hat{y}_i of the received sequence \hat{y} , define its log-likelihood as $\gamma_i = \log \left(\frac{\mathbb{P}[\hat{y}_i | y_i = 0]}{\mathbb{P}[\hat{y}_i | y_i = 1]} \right)$, where y_i represents the corresponding bit of the transmitted codeword y . Using the memoryless property of the channel, it can be seen that the ML codeword is

$$\hat{y}_{\text{ML}} = \arg \min_{y \in \mathcal{C}} \sum_{i=1}^n \gamma_i y_i. \quad (1)$$

¹Note that our analysis yields a bound on the probability of failure for every finite block length n .

Without changing the outcome of the maximization, we can replace the code \mathcal{C} by its convex hull $\text{conv}(\mathcal{C})$, and thus express ML decoding as the linear program

$$\hat{y}_{\text{ML}} = \arg \min_{y \in \text{conv}(\mathcal{C})} \sum_{i=1}^n \gamma_i y_i. \quad (2)$$

Although we have converted the decoding problem from an integer program to a linear program, it remains intractable because for general factor graphs with cycles, the codeword polytope does not have a concise description.

A natural approach, and one that is standard in operations research and polyhedral combinatorics, is to relax the linear program by taking only a polynomial set of constraints that approximate the codeword polytope $\text{conv}(\mathcal{C})$. The first-order LP decoding method [14] makes use of a relaxation that results from looking at each parity check, or equivalently at each row of H , in an independent manner. For each check $a \in C$ in the code, denote by \mathcal{C}_a the set of binary sequences that satisfy it—that is, \mathcal{C}_a corresponds to the local parity-check subcode defined by check a and its bit neighbors. Observe that the full code \mathcal{C} is simply the intersection of all the local codes, and the codeword polytope has the *exact* representation $\text{conv}(\mathcal{C}) = \text{conv}(\bigcap_{a=1}^m \mathcal{C}_a)$. The first-order LP decoder simply ignores interactions between the various local codes, and performs the optimization over the relaxed polytope given by $\mathcal{P} := \bigcap_{a=1}^m \text{conv}(\mathcal{C}_a)$. Note that \mathcal{P} is a convex set that contains the codeword polytope $\text{conv}(\mathcal{C})$, but also includes additional vertices with fractional coordinates (called *pseudocodewords* in the coding literature). It can be shown [31] that if the LDPC graph had no cycles and hence were tree-structured, this relaxation would be exact; consequently, this relaxation can be thought of as tree-based.

In sharp contrast to the codeword polytope for a general factor graph with cycles, the relaxed polytope \mathcal{P} for LDPC codes is always defined by a linear number of constraints. Consequently, LP decoding based on solving the relaxed linear program

$$\hat{y}_{\text{LP}} = \arg \min_{y \in \mathcal{P}} \sum_{i=1}^n \gamma_i y_i \quad (3)$$

can be solved exactly in polynomial time using standard LP solvers (e.g., interior point or simplex methods), or even faster with iterative methods tailored to the problem structure [11], [29]–[31].

For completeness, we now provide an explicit inequality description of the relaxed polytope \mathcal{P} . For every check a connected to neighboring variables in the set $N(a)$ and for all subsets $S \subseteq N(a)$, $|S|$ odd, we introduce the following constraints:

$$\sum_{i \in N(a) \setminus S} y_i + \sum_{i \in S} (1 - y_i) \geq 1. \quad (4)$$

Each such inequality corresponds to constraining the ℓ_1 distance of the polytope from the sequences not satisfying check a —the forbidden sequences—to be at least one. It can be shown that these *forbidding inequalities* do not exclude valid codewords from the relaxed polytope. We also need to add a set of $2n$ *box inequalities*—namely, $0 \leq y_i \leq 1$ —in order to ensure that we remain inside the unit hypercube. The set of forbidding inequalities along with the $[0, 1]$ -box inequalities define the relaxed polytope.

Note that, given a check a of degree d_c , there are 2^{d_c-1} local forbidden sequences, i.e., sequences of bits in the check neighborhood $N(a)$ that do not satisfy the check a . Consequently, for a constant check degree code, the total number of local forbidden sequences is $2^{d_c-1}m$, so that the number of forbidding inequalities scales linearly in the block length n . Fortunately, in the case of LDPC codes, the degree d_c is usually either a fixed constant (for regular constructions) or small with high probability (for irregular constructions) so that the number of local forbidden sequences remains small. Overall, in the cases of practical interest, the relaxed polytope can be characterized by a linear number of inequalities in the way that we have described. (We refer the interested reader to [14], [32] for alternative descriptions more suitable for the case of large d_c .)

III. MAIN RESULT AND PROOF OUTLINE

In this section, we state our main result characterizing the performance of LP decoding for a random ensemble of LDPC codes, and provide an outline of the main steps. Section IV completes the technical details of the proof.

A. Random Code Ensemble

We consider the random ensemble of codes constructed according to the following procedure. Given a code rate $R \in (0, 1)$, form a bipartite factor graph $G = (V, C)$ with a set of $n = |V|$ variable nodes, and $m = |C| = \lfloor (1 - R)n \rfloor$ check nodes as follows: i) Fix a variable degree $d_v \in \mathbb{N}$; and ii) For each variable $j \in V$, choose a random subset $N(j)$ of size d_v from C , and connect variable j to each check in $N(j)$. For obvious reasons, we refer to the resulting ensemble as the *bit-degree-regular* random ensemble, and use $\mathcal{C}(d_v)$ to denote a randomly chosen LDPC code from this ensemble.

The analysis of this paper focuses primarily on the BSC, in which each bit of the transmitted codeword is flipped independently with some probability α . By concentration of measure for the binomial distribution, it is equivalent (at least asymptotically) to assume that a constant fraction αn of bits are flipped by the channel. Let \mathbb{P} denote the joint distribution, over both the space of bit-degree-regular random codes, and the space of αn bit flips. With the goal of obtaining upper bounds on the LP error probability $\mathbb{P}[\text{LP fails}]$, our analysis is based on the *expansion* of the factor graph of the code. Specifically, the factor graph of a code with block length n is a (μ, p) -*expander* if all sets S of *variable nodes*, of size $|S| \leq \mu n$, are connected to at least $p|S|$ checks.²

B. Statement of Main Result

Our main result is a novel bound on the probability of error for LP decoding, applicable for finite block length n and the bit-degree-regular LDPC ensemble. The main idea is to show that, under certain expansion properties of the code, LP decoding will succeed in recovering the correct codeword with high probability. We note that a random graph will have the required expansion properties with high probability.

²Throughout this paper, we work with codes with simple parity-check constraints (LDPC codes) which are different from the generalized expander codes [25], [13] that can have large linear subcodes as constraints.

In particular, we show that for the joint distribution over random expander bit-degree-regular codes and $\lceil \alpha_c n \rceil$ (or less) bit flips by the channel, there exists a constant α_c , depending on the expansion properties of the ensemble, such that LP decoding succeeds with high probability. More formally,

Theorem 1: For every bit-degree-regular LDPC code ensemble with parameters R, d_v, n , we specify quadruples (α_c, c, μ, p) for which the LP decoder succeeds with high probability over the space of (μ, p) -expander bit-regular random codes and at most $\lceil \alpha_c n \rceil$ bit flips. The probability of failure decreases exponentially in c —namely

$$\mathbb{P}[\text{LP success} | \mathcal{C}(d_v) \text{ is an } (\mu, p) \text{ expander}] \geq 1 - e^{-cn}. \quad (5)$$

We note that any factor graph sampled from the bit-regular ensemble will be an expander with high probability. In general, the fraction of correctable errors α_c guaranteed by Theorem 1 is a function of the code ensemble, specified by the code rate R , the bit degree d_v , its expansion parameters μ and p , and the error exponent c . For any code rate, the maximum fraction of correctable errors α_c achieved by our analysis is provably larger than that of the best previously known result [12] for LP decoding, which guaranteed correction of a fraction $\frac{3p-2}{2p-1} \mu$ of errors. As a particular illustration of the stated Theorem 1, we have the following guarantee for rate $R = \frac{1}{2}$ codes.

Corollary 1: For code rate $R = \frac{1}{2}$, bit degree $d_v = 8$, and error fraction $\alpha \in (0, 0.002)$, the LP decoder succeeds with probability $1 - o(1)$ over the space of bit-degree-regular random codes of degree d_v and $\lceil \alpha n \rceil$ bit flips.

More generally, for any code rate R , our analysis in Section IV (see discussion following Lemma 8) specifies conditions for the bit flipping probability α_c and the expansion parameters μ and p so that the condition (5) is satisfied with a suitable choice of error exponent c .

C. Outline of Main Steps

We now describe the main steps involved in the Proof of Theorem 1.

1) *Hyperflow Witness:* As in previous work [12], we prove that the LP decoder succeeds by constructing a *dual witness*: a dual feasible vector with zero dual cost, which guarantees that the transmitted codeword is optimal for the primal linear program. Using the symmetry of the relaxed polytope, it can be shown [14] that the failure or success of LP decoding depends only on the subset of bits flipped by the channel and *not* on the transmitted codeword. Consequently, we may assume without loss of generality that the all-zero codeword was transmitted. Moreover, note that, for the BSC with flip probability ϵ , the log-likelihood of each received bit is either $\log(\frac{1-\epsilon}{\epsilon})$ or $-\log(\frac{1-\epsilon}{\epsilon})$. Since the optimum of the primal is not affected by rescaling, we may assume without loss of generality that all γ_i are either 1 or -1 . Then, every flipped bit i will be assigned $\gamma_i = -1$, whereas every unflipped bit $\gamma_i = 1$. Under these assumptions, Feldman *et al.* [12] demonstrated that a dual witness can be graphically interpreted as a set of weights on the edges of the factor graph of the code.




Fig. 1. Example of a valid hyperflow: bits x_1 and x_4 have been flipped in a BSC, and hence each are contaminated with one unit of poison. Each of the unflipped bits can absorb up to one unit of poison, whereas the checks act as hyperedges and replicate any incoming flow in all directions other than the incoming one. The valid hyperflow shown in this figure certifies that LP decoding can correct these two flipped bits.

Lemma 1 (Dual Witness [12]): Suppose that all bits in the set F are flipped by the channel, whereas all bits in the complementary set $F^c := V \setminus F$ are left unchanged. Set $\gamma_i = -1$, for all $i \in F$, and $\gamma_i = 1$, for all $i \in F^c$. Then LP decoding succeeds for this error pattern if there exist weights τ_{ia} for all checks $a \in C$ and distinct adjacent bits $i \in N(a)$ such that the following conditions hold:

$$\tau_{ia} + \tau_{ja} \geq 0, \quad \forall \text{ checks } a \in C, \text{ and} \quad (6a)$$

$$\forall \text{ adjacent bits } i, j \in N(a).$$

$$\sum_{a \in N(i)} \tau_{ia} < \gamma_i \quad \forall i \in V. \quad (6b)$$

We next introduce a sufficient condition for the success of LP decoding, one which is equivalent but arguably more intuitive than the dual witness definition.

Definition 1: A *hyperflow* for γ is a set of edge weights τ_{ij} that satisfy condition (6b) and moreover, have the following additional property: for every check $j \in C$, there exists a $\mathcal{P}_j \geq 0$ such that for exactly one variable $i \in N(j)$, $\tau_{ij} = -\mathcal{P}_j$, and for all the other $i' \in N(j) \setminus \{i\}$, $\tau_{i'j} = \mathcal{P}_j$.

The flow interpretation is that each check corresponds to a hyperedge connecting its adjacent variables; the function of any check is to replicate the flow incoming from one variable toward all its other adjacent variables. With this setup, condition (6b) corresponds to the requirement that all the variables i with $\gamma_i < 0$ need to get rid of at least $-\gamma_i$ units of “poison,” whereas each variable i with $\gamma_i > 0$ can absorb at most γ_i units of “poison.” Fig. 1 illustrates a valid hyperflow for a simple code.

We claim that the existence of a valid hyperflow is equivalent to the dual witness:

Proposition 1: There exists a weight assignment τ_{ij} satisfying the conditions of Lemma 1 if and only if there exists a hyperflow τ'_{ij} for γ .

See Appendix A for a proof of this claim.

2) *Hyperflow From (p, q) Matching:* Let $N(F)$ denote the subset of checks that are adjacent to the set F of flipped bits. One way to construct a hyperflow for γ is to match each bit i in the set F of flipped bits, with some number of checks, say $p \leq d_v$ checks, to which it has the exclusive privilege to push flow, suppose in a uniform fashion. This follows because in a matching, each check is used at most once. Let us refer to the checks that are actually used in such a matching as *dirty*, and to all the checks in $N(F)$ as *potentially dirty*. The challenge is that there might be unflipped variables that are adjacent to a large number of dirty checks, and hence fail to satisfy condition (6b); i.e., they receive more flow than they can absorb. Thus, the goal is to construct the matching of the flipped bits in a careful way so that no unflipped bit has too many dirty neighbors. The δ -matching witness, used by Feldman *et al.* [12], avoids this difficulty by matching *all* of the bits adjacent to potentially dirty checks with $\delta = p$ checks each. Our approach circumvents this difficulty using a more refined combinatorial object that we call a (p, q) -matching. For each bit $j \in F^c$, let $Z_j := |N(j) \cap N(F)|$ be the number of its edges adjacent to checks in $N(F)$.

Definition 2: Given nonnegative integers p and q , a (p, q) -matching is defined by the following conditions:

- a) each bit $i \in F$ must be matched with p distinct checks, and
- b) each bit $j \in F^c$ must be matched with

$$X_j := \max\{q - d_v + Z_j, 0\} \tag{7}$$

distinct checks from the set $N(F)$.

In all theoretical analysis in this paper, it is technically convenient to consider only pairs (p, q) such that

$$p \geq q, \quad 2p + q > 2d_v, \quad \text{and} \quad d_v \geq p + 2.$$

(The lone exception is Fig. 2, which is shown only for illustrative purposes.)

We refer to the number of checks with which each variable node needs to be matched as its *request number*. In this language, all flipped bits have p requests while each unflipped bit j has a variable number of requests X_j which depends on how many of its edges land on checks which have flipped neighbors. The following lemma justifies why requests are selected in this way and illustrates the key property of the (p, q) -matching:

Lemma 2: A (p, q) -matching guarantees that all the flipped bits are matched with p checks, and all the nonflipped bits have q or more non-dirty check neighbors.

This fact follows by observing that any unflipped bit j with Z_j edges in $N(F)$ has $d_v - Z_j$ clean neighboring checks, and requests $q - (d_v - Z_j)$ extra checks from the potentially dirty ones.




Fig. 2. Illustration of a (p, q) -generalized matching with $p = 2, q = 3$, and $d_v = 4$. The first three bits are flipped, and form the poisoned set S ; each flipped bit must be matched with $p = 2$ checks from its neighborhood (edges drawn with arrows). The bit node labeled 2 lies in S^c : it connects to $Z_2 = |N(2) \cap N(F)| = 3$ checks within the set $N(F)$, and so must be matched with $X_2 = q - (d_v - 3) = 2$ checks from $N(F)$ (two incoming arrows). By construction, bit 2 then has $2 + 1 = q$ checks that are not contaminated. Similarly, bit 1 connects $Z_1 = 2$ checks from $N(F)$, and so must be matched with $X_1 = q - (d_v - 2) = 1$ check from $N(F)$ (incoming arrow). It has a total of $1 + 2 = q$ checks that are not contaminated.

Fig. 2 illustrates a generalized matching for a degree $d_v = 4$ factor graph, and $(p, q) = (2, 3)$. Note that the bit node labeled 2 has $Z_2 = 3$ neighbors in the potentially dirty set $N(F)$, and so it makes $X_2 = 3 - (4 - 3) = 2$ requests for matching. This ensures that it is connected to $2 + 1 = q$ checks that are not dirty. A similar argument applies to bit 1, with $Z_1 = 2$ and $X_1 = 1$.

We next claim that a (p, q) -matching is a certificate of LP decoding success.

Lemma 3: For any integers p and q such that $2p + q > 2d_v$, a (p, q) -matching can be used to generate a set of weights τ_{ia} which constitute a hyperflow for γ and, hence, satisfy the dual witness conditions (6).

Proof: Each flipped bit is matched with p checks: suppose it sends χ units of poison to each of these checks. In the worst case, the remaining $d_v - p$ edges are connected to checks to which other flipped bits are sending poison. Therefore, each flipped bit (in the worst case) can purge itself of $p\chi - (d_v - p)\chi$ units of its own poison, so that we require that $p\chi - (d_v - p)\chi > 1$.

By Lemma 2, each unflipped bits has at least q checks that do not send any poison. In the worst case, then, an unflipped bit can receive $(d_v - q)\chi$ units of poison, which we require to be less than 1. Overall, a valid routing parameter χ will exist if $\frac{1}{2p-d_v} < \frac{1}{d_v-q}$, or equivalently, if $p + q > 2d_v$ as claimed. \square

In fact, it can verified that our combinatorial witness for LP decoding success is easier to satisfy than the condition used by Feldman *et al.* [12]. Our use of this improved witness, along with our focus on the probabilistic setting, are the two ingredients that allow us to establish a much larger fraction α_c of correctable errors.

3) *From Expansion to Matching Via Hall's Theorem:* The remainder (and bulk) of the analysis involves establishing that, with high probability over the selection of random expander bit-degree-regular codes and random subsets of $\lceil \alpha n \rceil$ flipped bits, a (p, q) -matching exists, for suitable values $q \leq p$ to be specified later. It is well known [3], [25] that random regular bipartite graphs will have good expansion, with high probability:

Lemma 4 (Good Expansion): For any fixed code rate $R \in (0, 1)$, degree d_v , and $p \leq d_v - 2$, there exist constants $\mu, c > 0$ so that a code $\mathcal{C}(d_v)$ from the bit-degree-regular ensemble

of degree d_v is a $(\mu n, p)$ expander with probability at least $1 - O(1/n)$.

Therefore, conditioned that the event that the random graph is an expander, the next step is to analyze the existence of a (p, q) -matching. We use Hall's theorem [27], which in our context, states that a matching exists if and only if *every subset of the variable nodes* have (jointly) enough neighbors in $N(F)$ to cover the sum of their requests.

Given our random graph and channel models, an equivalent description of the neighborhood choices for each variable $j \in F^c$ is as follows. Each node $j \in F^c$ picks a random number $Z_j \in \{0, 1, \dots, d_v\}$ according to the binomial distribution $\text{Bin}(d_v, \frac{|N(F)|}{m})$, and picks a subset of $N(F)$ of size Z_j . This subset corresponds to the intersection of its check neighborhood $N(j)$ with the check neighborhood $N(F)$ of the flipped bits. The remaining $d_v - Z_j$ edges from bit j connect to checks outside $N(F)$. With this setup, we now define the a "bad event" \mathcal{E} , defined by the existence of a pair $(S_1, S_2) \in 2^F \times 2^{F^c}$ of sets that *contracts*, meaning that it has more requests than neighbors, so that

$$\begin{aligned} & |N(S_1) \cup [N(S_2) \cap N(F)]| \\ & < p|S_1| + \sum_{j \in S_2} \max\{0, q - (d_v - Z_j)\}. \end{aligned} \quad (8)$$

Notice that only the neighbors in $N(F)$ are counted, since a (p, q) -matching involves only checks in $N(F)$. By Lemma 3, the event \mathcal{E} must occur whenever LP decoding fails so that we have the inequality $\mathbb{P}[\text{LP decoding fails}] \leq \mathbb{P}[\mathcal{E}]$. Defining the event

$$\mathcal{G}(d_v, \mu, p) := \{\mathcal{C}(d_v) \text{ is a } (\mu n, p) \text{ expander}\} \quad (9)$$

we make use of the following conditional form of this inequality:

$$\mathbb{P}[\text{LP decoding fails} | \mathcal{G}] \leq \mathbb{P}[\mathcal{E} | \mathcal{G}].$$

It is useful to partition the space $2^F \times 2^{F^c}$ into three subsets controlled by the parameters $\epsilon_2, \mu > 0$. Parameter $\epsilon_2 > 0$ is a small constant to be specified later in the proof and μ is the expansion coefficient. The three subsets of interest are given by

$$A_1 := \{(S_1, S_2) | (S_1, S_2) \in A, |S_1| + |S_2| < \mu n\} \quad (10a)$$

$$A_2 := \{(S_1, S_2) | (S_1, S_2) \in A - A_1, |S_1| \geq \epsilon_2 n\} \quad (10b)$$

$$A_3 := A \setminus (A_1 \cup A_2). \quad (10c)$$

This partition, as illustrated in Fig. 3, decomposes \mathcal{E} into subevents

$$\mathcal{E}(A_i) := \{\exists (S_1, S_2) \in A_i | \text{equation (8) holds}\}$$

for $i = 1, 2, 3$. Then, via a series of union bounds, we have the following upper bound on the probability of failure:

$$\mathbb{P}[\text{LP fails} | \mathcal{G}] \leq \mathbb{P}[\mathcal{E} | \mathcal{G}] \leq \sum_{i=1}^3 \mathbb{P}[\mathcal{E}(A_i) | \mathcal{G}].$$

However, all subsets of variable nodes of size at most μn in a $(\mu n, p)$ expander have a p -matching and, because $q \leq p$, it follows that

$$\mathbb{P}[\mathcal{E}(A_1) | \mathcal{G}] = 0. \quad (11)$$

Consequently, we only have to deal with the remaining two terms of the summation for $i = 2$ and $i = 3$. Before proceeding, an important side remark here is that (11) by itself implies that

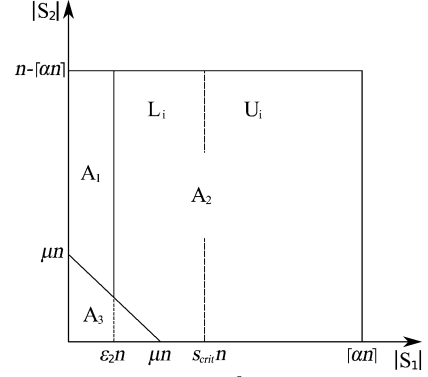


Fig. 3. Partitioning the space $2^F \times 2^{F^c}$.

the LP decoder can correct a constant fraction of errors; indeed, it is precisely this observation that was exploited by Feldman *et al.* [12]. However, our ultimate goal in this paper is to establish higher fractions of correctable errors, so we need to continue our analysis further.

For $i = 2, 3$, we have

$$\begin{aligned} \mathbb{P}[\mathcal{E}(A_i) | \mathcal{G}] &= \frac{\mathbb{P}[\mathcal{E}(A_i) \wedge \mathcal{G}]}{\mathbb{P}[\mathcal{G}]} \\ &\leq \frac{\mathbb{P}[\mathcal{E}(A_i)]}{\mathbb{P}[\mathcal{G}]} \\ &\leq 2\mathbb{P}[\mathcal{E}(A_i)] \end{aligned} \quad (12)$$

where the last inequality follows from Lemma 4. Overall, putting everything together, we conclude that

$$\mathbb{P}[\text{LP fails} | \mathcal{G}] \leq 2 \sum_{i=2}^3 \mathbb{P}[\mathcal{E}(A_i)]. \quad (13)$$

The remainder of the proof consists of careful analysis of these two error terms. It turns out to be convenient to use an alternative probabilistic model in the analysis. In particular, observe that there is an inconvenient asymmetry in the definition of our generalized matching: the bits of set F^c need to be matched with checks from the neighborhood of the flipped bits F , and not from the whole set of checks from which they select their neighbors. This correlation between $N(F)$ and the number of requests from set F^c creates severe complications in the analysis. Indeed, any attempt to use Hall's condition through union bounds seems to require independence among different edges; moreover, crude upper bounds on the number of requests from set F^c seem inadequate to decorrelate the requests of F^c from the size of $N(F)$. For this reason, we use an alternative probabilistic model, as described in Section IV-A.

IV. PROOF OF THEOREM 1

We now turn to the remaining (somewhat more technical) steps involved in the Proof of Theorem 1.

A. Simplifying the Probability Model

In order to decouple the distribution of the requests of F^c from the size of $N(F)$, observe that the number of requests X_j from each bit j in F^c grows linearly with the number of edges that this bit has in $N(F)$. Notice that the checks are selected with replacement and the degree of a variable can be strictly smaller than d_v , although this will not be an issue asymptotically. This observation combined with a coupling argument shows that, if

$x, x' \in \{0, \dots, d_v\}^{|F^c|}$ are two vectors of requests from the bits in F^c , where $x \leq x'$ elementwise, then the probability that a (p, q) -matching exists is larger conditioned on x than on x' .

This observation suggests the following alternative experiment.

- A node $j \in F^c$ first picks a random number $Z_j \in \{0, 1, \dots, d_v\}$ according to the modified binomial distribution $\text{Bin}(d_v, \frac{d_v \lceil \alpha n \rceil}{m})$.
- Node j then chooses Z_j checks from $N(F)$ with replacement.

This procedure is repeated independently for each $j \in F^c$. Since $|N(F)| \leq d_v \lceil \alpha n \rceil$, the bits of set F^c will tend to have more edges in $N(F)$ and, therefore, more requests in this new experiment than in the original one (as suggested by the natural coupling between the two processes). Moreover, since checks are now chosen with replacement, for each bit $j \in F^c$, the size of the intersection $N(j) \cap N(F)$ is less than or equal to Z_j , since the same check might be chosen more than once. Intuitively, the existence of matchings is less likely in the new experiment than in the original one; this claim follows rigorously by combining these observations with the coupling argument used in the previous paragraph. The benefit of switching from the original experiment to this new experiment is in allowing us to *decouple* the process of deciding the number of requests made by each bit in F^c from the cardinality of the random variable $N(F)$.

Let us use \mathbb{Q} to denote the probability distribution over random graphs in this new model. Setting $F^c(q) = \{j \in F^c \mid q > d_v - Z_j\}$, we can define the alternative “bad event” \mathcal{B} , meaning that there exist $S_1 \subseteq F$, and $S_2 \subseteq F^c(q)$ such that

$$|N(S_1) \cup [N(S_2) \cap N(F)]| \leq p|S_1| + \sum_{j \in S_2} [q - (d_v - Z_j)]. \quad (14)$$

In addition, we define the corresponding subevents $\mathcal{B}(A_i)$ for $i = 1, 2, 3$. As argued above, it must hold that

$$\mathbb{P}[\mathcal{E}(A_i)] \leq \mathbb{Q}[\mathcal{B}(A_i)], \quad \text{for all } i = 1, 2, 3$$

and, therefore, as inequality (13) suggests, in order to upper-bound the probability of LP decoding failure, it suffices to obtain upper bounds on the probabilities $\mathbb{Q}[\mathcal{B}(A_i)]$ for $i = 2, 3$.

For future use, we define for fixed subsets $S_1 \subseteq F$ and $S_2 \subseteq F^c(q)$, the event $\mathcal{B}(S_1, S_2)$ that (14) holds for S_1 and S_2 . We now proceed, in a series of steps, to obtain suitable upper bounds on the probabilities $\mathbb{Q}[\mathcal{B}(A_i)]$ and, hence, on the probability of LP decoding failure.

B. Conditioning on Requests From F^c

For each $i \in \{1, \dots, q\}$, we define the random variable

$$Y_i := |\{j \in F^c \mid Z_j = d_v - (q - i)\}| \quad (15)$$

corresponding to the number of bits in F^c with $d_v - (q - i)$ edges that lie inside the “contaminated” neighborhood $N(F)$

and, hence, with i requests each. Note that each Y_i is binomial with parameters $\lfloor (1 - \alpha)n \rfloor$ and

$$b_i := \binom{d_v}{d_v - q + i} \left(\frac{\lceil \alpha n \rceil d_v}{m} \right)^{d_v - q + i} \left(1 - \frac{\lceil \alpha n \rceil d_v}{m} \right)^{q - i}. \quad (16)$$

Since $\mathbb{E}[Y_i] = b_i \lfloor (1 - \alpha)n \rfloor$, applying Hoeffding’s inequality [17] yields the sharp concentration

$$\mathbb{Q}[|Y_i - b_i \lfloor (1 - \alpha)n \rfloor| \geq \epsilon_1 n] \leq 2 \exp(-2\epsilon_1^2 n)$$

for any $\epsilon_1 > 0$. Hence, if we define the event

$$\mathcal{T}(\epsilon_1) := \bigcap_{i=1}^q \{|Y_i - b_i \lfloor (1 - \alpha)n \rfloor| \leq \epsilon_1 n\} \quad (17)$$

then a simple union bound yields that

$$\mathbb{Q}[\mathcal{T}(\epsilon_1)] \geq 1 - 2q \exp(-2\epsilon_1^2 n)$$

so that it suffices to bound the conditional probabilities $\mathbb{Q}[\mathcal{B}(A_i) \mid \mathcal{T}(\epsilon_1)]$, $i = 2, 3$. Note that conditioned on the event $\mathcal{T}(\epsilon_1)$, we are guaranteed that

$$\frac{Y_i}{n} \leq b_i(1 - \alpha) + \epsilon_1 =: y_i^{-\text{up}}. \quad (18)$$

We now turn to bounding the probability of the bad event \mathcal{B} . Since, by symmetry, the probability of the event $\mathcal{B}(S_1, S_2)$ is the same for different sets S_1 of the same size, a union bound gives $\mathbb{Q}[\mathcal{B}(A_2) \mid \mathcal{T}(\epsilon_1)] \leq \sum_{s_1 = \lceil \epsilon_2 n \rceil}^{\lceil \alpha n \rceil} D(s_1)$, where

$$D(s_1) := \binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}[\exists S_2 \subseteq F^c(q), \text{ with } (S_1, S_2) \in A_2 \text{ s.t. } \mathcal{B}(S_1, S_2) \mid \mathcal{T}(\epsilon_1)]$$

with S_1 is any fixed set of size s_1 .

Before bounding these terms, we first partition the values of s_1 into two sets $\{\lceil \epsilon_2 n \rceil, \dots, \lceil \bar{s}_{\text{crit}} n \rceil\}$ and $\{\lceil \bar{s}_{\text{crit}} n \rceil + 1, \dots, \lceil \alpha n \rceil\}$ for some value of \bar{s}_{crit} to be specified formally in Lemma 5. To give some intuition, in the conditional space $\mathcal{T}(\epsilon_1)$, the total number of matching-requests from the bits of set F^c is at most

$$V := n \sum_{i=1}^q i y_i^{-\text{up}}. \quad (19)$$

Therefore, if $\mathbb{Q}[\mathcal{B}(A_2) \mid \mathcal{T}(\epsilon_1)]$ is indeed small, we would expect that, if the set S_1 is large enough (say $|S_1| \approx |F|$), then with high probability, the size of its image $N(S_1)$ should be large enough not only to cover its own requests but also V additional requests—viz. $|N(S_1)| \geq p|S_1| + V$. If this condition holds, then there cannot exist any set S_2 such that the event $\mathcal{B}(S_1, S_2)$ occurs. We formalize this intuition in the following result, proved in Appendix C.

Lemma 5 (Upper Regime): Define $\bar{v} := \sum_{i=1}^q i \bar{y}_i^{\text{up}}$, and the function

$$f(s) := \alpha H\left(\frac{s}{\alpha}\right) + (1-R)H\left(\frac{ps + \bar{v}}{(1-R)}\right) + d_v s \log_2\left(\frac{ps + \bar{v}}{(1-R)}\right)$$

where $H(\cdot)$ is the binary entropy function, and set

$$\bar{s}_{\text{crit}} := \min\{\alpha, \inf\{s \in [0, \alpha] | f(s') < 0, \forall s' \in [s, \alpha]\}\}. \quad (20)$$

Then, for all $s_1 \in \{\lceil \bar{s}_{\text{crit}} n \rceil + 1, \dots, \lceil \alpha n \rceil\}$, the quantity $D(s_1)$ decays exponentially fast in n .

It remains to bound $D(s_1)$ for $s_1 \in L_I$, where

$$L_I := \{\lceil \epsilon_2 n \rceil, \dots, \lceil \bar{s}_{\text{crit}} n \rceil\}.$$

For a randomly chosen set S_1 , define the event

$$\mathcal{F}(s_1, \gamma_1) := \{|S_1| = s_1, |N(S_1)| = \gamma_1\}. \quad (21)$$

By conditioning, we have the decomposition

$$D(s_1) = \sum_{\gamma_1=1}^{d_v s_1} E(\gamma_1, s_1)$$

where

$$E(\gamma_1, s_1) := \binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}'[\mathcal{F}(s_1, \gamma_1)] \times \mathbb{Q}'[\exists S_2 \text{ with } (S_1, S_2) \in A_2 \text{ s.t. } \mathcal{B}(S_1, S_2) | \mathcal{F}(s_1, \gamma_1)].$$

Here \mathbb{Q}' denotes the conditional probability distribution of \mathbb{Q} conditioned on the event $\mathcal{T}(\epsilon_1)$.

The following lemma allows us to restrict our attention to linearly sized check neighborhoods $N(S_1)$ in analyzing the individual terms $E(\gamma_1, s_1)$ of the summation; the proof is provided in Appendix D.

Lemma 6 ($\bar{\gamma}_1$ small): Define the critical point $\bar{\gamma}_{\text{crit}}(\bar{s}_1)$

$$\sup\left\{\bar{\gamma}_1 \in (0, d_v \bar{s}_1) \mid 2 + d_v \bar{s}_1 \log_2\left(\frac{\bar{\gamma}_1}{(1-R)}\right) < 0\right\}. \quad (22)$$

Then, for set sizes $s_1 \geq \lceil \epsilon_2 n \rceil$ and neighborhood sizes $\gamma_1 \leq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$, the quantity $E(\gamma_1, s_1)$ decays exponentially fast in n .

Note that the supremum (22) is always finite. This lemma essentially says that if s_1 has linear size, its neighborhood γ_1 must also have linear size.

To summarize our progress thus far, we first argued that in order to bound the probability $\mathbb{Q}[\mathcal{B}(A_2) | \mathcal{T}(\epsilon_1)]$, it suffices to bound the quantities $D(s_1)$, for $s_1 \in \{\lceil \epsilon_2 n \rceil, \dots, \lceil \alpha n \rceil\}$. Next, we partitioned the range of s_1 into two sets: the lower set $L_I = \{\lceil \epsilon_2 n \rceil, \dots, \lceil \bar{s}_{\text{crit}} n \rceil\}$, and the upper set $U_I := \{\lceil \bar{s}_{\text{crit}} n \rceil + 1, \dots, \lceil \alpha n \rceil\}$. The upper set has the property that for all sets $S_1 \subseteq F$ of size $|S_1| \in U_I$, with high probability, the neighborhood $N(S_1)$ is big enough to accommodate not only the

matching requests from set S_1 , but also all possible matching-requests from any set $S_2 \subseteq F^c$. Having established this property of large S_1 sets, it remains to focus on small S_1 . In this regime, the neighborhood $N(S_1)$ on its own is no longer sufficient to cover the joint set of requests from S_1 and from any possible set $S_2 \subseteq F^c$. Consequently, one has to consider for every choice $(S_1, S_2) \in A_2$, whether the joint neighborhood $N(S_1) \cup (N(S_2) \cap N(F))$ is large enough to cover the matching requests from S_1 and S_2 .

At this point, one might imagine that a rough concentration argument applied to the sizes of $N(S_1)$ and $N(S_2) \cap N(F) \setminus N(S_1)$ would suffice to complete the proof. Unfortunately, any concentration result must be sufficiently strong to dominate the factor $\binom{\lceil \alpha n \rceil}{s_1}$ that leads the expression $D(s_1)$. Consequently, we study the exact distribution of the size of $N(S_1)$, and bound the quantities $E(\gamma_1, s_1)$ for $s_1 \in L_I$ and $\gamma_1 \in \{1, \dots, d_v s_1\}$. Of course, since s_1 is linear in size, the bulk of the probability mass is concentrated on linear values for γ_1 . Therefore, by Lemma 6, we need only bound $E(\gamma_1, s_1)$ for $s_1 \in L_I$ and $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$. We complete these steps in the following subsection.

C. Completing the Bound

Let us fix sizes $s_1 \in L_I$ and $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$. For a set S_1 of size s_1 with neighborhood $N(S_1)$ of size γ_1 , define its *residual neighborhood* to be the set $N(F) \setminus N(S_1)$ and use $\gamma_2 := |N(F) \setminus N(S_1)|$ to denote its size. Moreover, define the vector of requests³ $y \in \otimes_{i=1}^q \{0, \dots, \lceil \bar{y}_i^{\text{up}} n \rceil\}$, let us denote by $\beta(s_1, \gamma_1, y)$ the number of checks missing from the neighborhood of S_1 to cover the total number of requests from S_1 and a set S_2 with configuration of requests y . Also, let $\nu(y)$ be the number of edges from S_2 to $N(F)$. More precisely, the quantities $\beta(s_1, \gamma_1, y)$ and $\nu(y)$ are given by the formulas

$$\beta(s_1, \gamma_1, y) := ps_1 - \gamma_1 + \sum_{i=1}^q i y_i \quad (23a)$$

$$\nu(y) := \sum_{i=1}^q (d_v - q + i) y_i. \quad (23b)$$

Note that for $s_1 \in L_I$ and $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$, the quantity β also grows linearly in n ; as usual, we use $\bar{\beta}$ to denote the rescaled quantity β/n . Also, recall the definition of \bar{y}_i^{up} from (18).

Letting $\bar{y} := (\bar{y}_1, \dots, \bar{y}_q)$ be a vector of request fractions in $[0, 1]^q$, we define

$$G(\bar{s}_1, \bar{\gamma}_1, \bar{\gamma}_2, y) := \sum_{k=1}^4 \min\{0, G_k(\bar{s}_1, \bar{\gamma}_1, \bar{\gamma}_2, \bar{y})\}$$

where

$$G_1 := \alpha H\left(\frac{\bar{s}_1}{\alpha}\right) \sum_{i=1}^q \bar{y}_i^{\text{up}} H\left(\frac{\bar{y}_i}{\bar{y}_i^{\text{up}}}\right)$$

$$G_2 := (1-R)H\left(\frac{\bar{\gamma}_1}{(1-R)}\right) + d_v \bar{s}_1 \log_2\left(\frac{\bar{\gamma}_1}{(1-R)}\right)$$

³Recall that we have conditioned on the event $\mathcal{T}(\epsilon_1)$, so that the number of bits in F^c with i matching requests is concentrated, for every $i \in \{1, \dots, q\}$.

$$\begin{aligned}
G_3 &:= ((1-R) - \bar{\gamma}_1) H \left(\frac{\bar{\gamma}_2}{((1-R) - \bar{\gamma}_1)} \right) \\
&\quad + d_v(\alpha - \bar{s}_1) \log_2 \left(\frac{\bar{\gamma}_1 + \bar{\gamma}_2}{(1-R)} \right) \\
G_4 &:= \bar{\gamma}_2 H \left(\frac{\min\{\bar{\gamma}_2, \bar{\beta}(\bar{s}_1, \bar{\gamma}_1, \bar{y})\}}{\bar{\gamma}_2} \right) \\
&\quad + \nu(\bar{y}) \log_2 \left(\frac{\bar{\gamma}_1 + \min\{\bar{\gamma}_2, \bar{\beta}(\bar{s}_1, \bar{\gamma}_1, \bar{y})\}}{\bar{\gamma}_1 + \bar{\gamma}_2} \right).
\end{aligned}$$

With these definitions, we have the following result.

Lemma 7 (Exponential Upper Bound): Suppose that the following inequalities hold

$$\bar{s}_{\text{crit}} < \frac{\alpha}{2} \quad (25a)$$

$$\alpha d_v < \frac{(1-R) - d_v \bar{s}_{\text{crit}}}{2} \quad \text{and} \quad (25b)$$

$$\alpha H \left(\frac{\bar{s}_{\text{crit}}}{\alpha} \right) d_v(\alpha - \bar{s}_{\text{crit}}) \log_2 \left(\frac{d_v \bar{s}_{\text{crit}}}{(1-R)} \right) < 0. \quad (25c)$$

Then for some $c > 0$, we have the exponential upper bound

$$\mathbb{Q}[\mathcal{B}(A_2) | \mathcal{T}(\epsilon_1)] \leq 2^{nF(\alpha)} + \exp(-cn)$$

where the function in the exponent is given by

$$F(\alpha) := \sup_{\bar{s}_1, \bar{\gamma}_1, \bar{\gamma}_2, \{\bar{y}_i\}} G(\bar{s}_1, \bar{\gamma}_1, \bar{\gamma}_2, \bar{y}_1, \dots, \bar{y}_q) \quad (26)$$

with the maximization over

$$\begin{aligned}
\bar{s}_1 &\in [0, \bar{s}_{\text{crit}}] \\
\bar{\gamma}_1 &\in [0, d_v \bar{s}_1] \\
\bar{\gamma}_2 &\in [0, d_v(\alpha - \bar{s}_1)] \quad \text{and} \\
\bar{y}_i &\in [\bar{y}_i^{\text{up}}/2, \bar{y}_i^{\text{up}}].
\end{aligned}$$

See Appendix E for a proof of this lemma.

It remains to upper-bound the probability of the bad event $\mathcal{B}(A_3)$ which is equivalent to the existence of a pair of contracting sets (S_1, S_2) , where the size of set $S_1 \subseteq F$ is at most $\epsilon_2 n$ and the size of set $S_2 \subseteq F^c$ is at least $(\mu - \epsilon_2)n$. Note that we have not yet specified the constant ϵ_2 . The following lemma establishes that there exists a value of ϵ_2 so that $\mathbb{Q}[\mathcal{B}(A_3) | \mathcal{T}(\epsilon_1)]$ is bounded by an exponentially decreasing function in n provided that the function $F(\alpha)$ from (26) is negative. The proof of this final lemma is provided in Appendix F.

Lemma 8: If $F(\alpha) < 0$, then there exists ϵ_2 so that the probability $\mathbb{Q}[\mathcal{B}(A_3) | \mathcal{T}(\epsilon_1)]$ is decreasing exponentially in n .

We may now complete the Proof of Theorem 1. For a given rate R , fix the bit degree d_v and the matching parameters (p, q) such that

$$2p + q > 2d_v, \quad p \geq, \quad \text{and} \quad d_v - p \geq 2 \quad (27)$$

and recall the definition (20) of \bar{s}_{crit} . Suppose that the three inequalities (25) hold, and that the function F defined in (26) satisfies

$$F(\alpha) < 0. \quad (28)$$

Then the probability

$$\mathbb{P}[\text{LP decoding fails} | \mathcal{C}(d_v) \text{ is a } (\mu n, p) \text{ expander}]$$

decays exponentially in n , where \mathbb{P} is the uniform distribution over the set of bit-degree-regular codes of degree d_v and selections of $\lceil \alpha n \rceil$ bit flips.

In particular, these explicit conditions allow us to investigate fractions of correctable errors on specific code ensembles. As a concrete example, for code rate $R = 1/2$, if we choose variable degrees $d_v = 8$ and generalized matching parameters $(p, q) = (6, 5)$, one can numerically verify that the conditions (27), (28), and (25) are satisfied for all $\alpha \leq \alpha_{\text{crit}} = 0.002$. Therefore, for that rate, we establish a fraction of correctable errors which is more than ten times higher than the previously known worst case results, as claimed.

V. CONCLUSION

The main contribution of this paper is to perform probabilistic analysis of linear programming (LP) decoding of low-density parity-check (LDPC) codes in the finite-length regime. Specifically, we showed that for a random LDPC code ensemble, the linear programming decoder of Feldman *et al.* succeeds (with high probability) in correcting a constant fraction of errors that surpasses all prior nonasymptotic results for LDPC codes. For a rate-1/2 code, it exceeds the best previous finite-length result on LP decoding by a factor greater than ten. Despite these substantial improvements, it should be noted that our analysis still yields very conservative results, roughly a factor of 50 lower than the typical empirical performance of these codes, as well as the associated asymptotic thresholds.

Perhaps more important than specific numerical improvements over past results are the technical innovations that underlie our analysis: a direct treatment of the probabilistic nature of bit-flipping channels (as opposed to adversarial analysis in previous work), and a novel combinatorial characterization of LP decoding, based on the notion of a poison hyperflow witness. This hyperflow perspective illustrates that the factor graph defining a good code should have good flow properties, in the sense that no matter which subset of bits are flipped, the poison associated with errors can be diffused and routed to the unflipped bits. For more general MBIOS channels, the amount of poison corresponds exactly to the negative log-likelihood that the channel is assigning to each bit, and the same characterization of LP decoding holds.

This intuition suggests that the property of supporting sufficient hyperflow could provide a useful design principle in the finite-length setting, for example, small sets of variables which contract (are jointly adjacent to few checks) will cause pseudocodewords of small pseudoweight.

There are a number of ways in which specific technical aspects of the current analysis can likely be sharpened, which await further work. In addition, it remains to further explore the consequences of our analysis technique for other channels and code ensembles, beyond the particular LDPC ensemble and binary-symmetric channel considered here.

APPENDIX

A. Proof of Proposition 1

One direction of the claim is immediate: given the weights τ'_{ij} of any hyperflow, they must satisfy condition (6b) by definition, and moreover it is easy to see that condition (6a) will be automatically satisfied. In the other direction, we transform the edge weights τ_{ij} to new weights τ'_{ij} that satisfy the hyperflow constraints. For each check j separately, we replace the weights τ_{ij} on the adjacent edges with new weights that satisfy the hyperflow constraints and, at the same time, do not violate any of the constraints in condition (6b). Consider a check j and order the weights on the adjacent edges in increasing order. Assuming that the check has degree $d(j)$, consider the following cases.

Case I: $0 \leq \tau_{1j} \leq \tau_{2j} \leq \dots \leq \tau_{d(j)j}$. In this case, set $\tau'_{ij} = 0$ for all i . The new weights are clearly hyperflow weights and, moreover, it is not hard to verify that none of the conditions (6b) are violated by the transformation.

Case II: $\tau_{1j} \leq 0 \leq \tau_{2j} \leq \dots \leq \tau_{d(j)j}$. Set $\mathcal{P}_j = -\tau_{1j}$, $\tau'_{1j} = -\mathcal{P}_j$, and $\tau'_{i'j} = \mathcal{P}_j, \forall i' \in N(j) \setminus \{1\}$. This is a hyperflow weight assignment by construction. Observe, also, that none of the conditions (6b) for the variables in $N(j)$ are violated by this transformation: indeed, for each variable $k \in N(j)$, we have that $\tau_{kj} \geq -\tau_{1j}$ since the weights τ_{ij} satisfy (6a); therefore, setting $\tau'_{kj} = -\tau_{1j}$ only makes the sum of the edges adjacent to variable k smaller, and the sum was already satisfying condition (6b) before the transformation.

To conclude the claim, notice that at most one edge adjacent to every check j can have negative weight in assignment τ_{ij} ; otherwise, condition (6b) would be violated for that check. \square

B. Elementary Bounds on Binomial Coefficients

For each $\beta \in (0, 1)$, define the binomial entropy

$$H(\beta) := -\beta \log_2 \beta - (1 - \beta) \log_2 (1 - \beta)$$

with $H(0) = H(1) = 0$ by continuity. We make use of standard asymptotics of binomial coefficients: for all integers k in the interval $[0, n]$, we have

$$\frac{1}{n} \log \binom{n}{k} = H\left(\frac{k}{n}\right) \pm o(1) \quad (29)$$

as n tends to infinity (e.g., see Cover and Thomas [7]).

C. Proof of Lemma 5

Note that conditioned on the event $\mathcal{T}(\epsilon_1)$, we are guaranteed that $\sum_{i=1}^q iY_i \leq \bar{v}n$. Letting S_1^* be the fixed subset $\{1, \dots, s_1\}$, define the event $\mathcal{E}(s_1) := \{|N(S_1^*)| \leq ps_1 + \bar{v}n\}$ and the quantity

$$P(s_1) := \binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}[\mathcal{E}(s_1)].$$

Using the nature of the bit-regular random ensemble, we have

$$P(s_1) \leq \binom{\lceil \alpha n \rceil}{s_1} \binom{\lfloor (1-R)n \rfloor}{\lfloor ps_1 + \bar{v}n \rfloor} \left(\frac{\lfloor ps_1 + \bar{v}n \rfloor}{\lfloor (1-R)n \rfloor} \right)^{d_v s_1}.$$

Setting $\bar{s}_1 = \frac{s_1}{n}$ and using standard bounds on binomial coefficients (29), the quantity $\frac{1}{n} \log P(s_1)$ is upper-bounded by

$$\left[\alpha H\left(\frac{\bar{s}_1}{\alpha}\right) + (1-R)H\left(\frac{p\bar{s}_1 + \bar{v}}{1-R}\right) + d_v \bar{s}_1 \log_2 \left(\frac{p\bar{s}_1 + \bar{v}}{1-R} \right) + o(1) \right].$$

Defining the function f and value \bar{s}_{crit} as in the lemma statement, we are guaranteed that $P(s_1)$ decays exponentially in n for all $s_1 \in \{\lceil \bar{s}_{\text{crit}} n \rceil + 1, \dots, \lceil \alpha n \rceil\}$. To complete the proof of the claim, we claim that $D(s_1)$ can be upper-bounded by $P(s_1)$. Indeed, for $s_1 \in \{\lceil \bar{s}_{\text{crit}} n \rceil + 1, \dots, \lceil \alpha n \rceil\}$, we can either condition on $\mathcal{E}(s_1)$ or its complement to obtain that $D(s_1)$ is upper-bounded by

$$\binom{\lceil \alpha n \rceil}{s_1} \left\{ \mathbb{Q}[\exists(S_1^*, S_2) \in A_2 \text{ s.t. } \mathcal{B}(S_1^*, S_2) | \mathcal{E}^c(s_1)] + \mathbb{Q}[\mathcal{E}(s_1)] \right\}$$

which is equal to $P(s_1)$ because, as argued in Section IV-B, conditioned on the event $\mathcal{E}^c(s_1)$, there can be no S_2 such that the event $\mathcal{B}(S_1^*, S_2)$ holds.

D. Proof of Lemma 6

We have the bound

$$E(\gamma_1, s_1) \leq \binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}[|N(S_1)| = \gamma_1 | |S_1| = s_1]$$

where we have used the fact that the event $\{|N(S_1)| = \gamma_1\}$ is independent of $\mathcal{T}(\epsilon_1)$ under the probability distribution \mathbb{Q} . An exact computation yields that

$$\frac{1}{n} \log \{ \mathbb{Q}[|N(S_1)| = \gamma_1 | |S_1| = s_1] \}$$

is upper-bounded by

$$\frac{1}{n} \log \left\{ \binom{\lfloor (1-R)n \rfloor}{\gamma_1} \left(\frac{\gamma_1}{\lfloor (1-R)n \rfloor} \right)^{d_v s_1} \right\}$$

which is in turn upper-bounded by

$$\ell(s_1) := \left\{ (1-R)H\left(\frac{\gamma_1}{\lfloor (1-R)n \rfloor}\right) + d_v \frac{s_1}{n} \log_2 \left(\frac{\gamma_1}{\lfloor (1-R)n \rfloor} \right) \right\} + o(1)$$

where we have used standard bounds on binomial coefficients (29). Overall, we have

$$\begin{aligned} \frac{1}{n} \log E(\gamma_1, s_1) &\leq \alpha H\left(\frac{s_1}{\lceil \alpha n \rceil}\right) + \ell(s_1) \\ &\leq \left\{ 2 + d_v \frac{s_1}{n} \log_2 \left(\frac{\gamma_1}{\lfloor (1-R)n \rfloor} \right) \right\} \end{aligned}$$

since $\alpha, R \in (0, 1)$, and each entropy term remains bounded within $[0, 1]$.

Finally, setting $\bar{s}_1 = s_1/n$ and $\bar{\gamma}_1 = \gamma_1/n$, consider the function

$$g(\bar{\gamma}_1) := 2 + d_v \bar{s}_1 \log_2 \left(\frac{\bar{\gamma}_1}{(1-R)} \right).$$

We have $\lim_{\bar{\gamma}_1 \rightarrow 0^+} g(\bar{\gamma}_1) = -\infty$, implying that $E(\gamma_1, s_1)$ decays exponentially fast in n for all $s_1 \geq \lceil \epsilon_2 n \rceil$ and neighborhood sizes $\gamma_1 \leq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$, where $\bar{\gamma}_{\text{crit}}(\cdot)$ is defined as in the statement of the lemma.

E. Proof of Lemma 7

We begin by proving the following lemma, which provides an upper bound on the quantity $E(\gamma_1, s_1)$.

Lemma 9 (Lower Regime): If the three conditions (25) hold, then, for all $s_1 \in \{\lceil \epsilon_2 n \rceil, \dots, \lceil \bar{s}_{\text{crit}} n \rceil\}$ and $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$, there exists some $\gamma_2^* = \gamma_2^*(\bar{s}_{\text{crit}}, \epsilon_2) > 0$ such that

$$\frac{1}{n} \log E(s_1, \gamma_1) \leq \sum_{k=1}^3 \min\{0, T_k(s_1, \gamma_1)\} + T_4(s_1, \gamma_1) + o(1)$$

where

$$\begin{aligned} T_1 &= \frac{1}{n} \log \binom{(1-R)n}{\gamma_1} + \frac{1}{n} \log \left(\frac{\gamma_1}{(1-R)n} \right)^{d_v s_1} \\ T_2 &= \frac{1}{n} \log \binom{(1-R)n - \gamma_1}{\gamma_2} \\ &\quad + \frac{1}{n} \log \left(\frac{\gamma_2 + \gamma_1}{(1-R)n} \right)^{(\alpha n - s_1) d_v} \\ T_3 &= \frac{1}{n} \log \left(\min\{\beta(s_1, \gamma_1, y), \gamma_2\} \right) \\ &\quad + \frac{1}{n} \log \left(\frac{\gamma_1 + \min\{\gamma_2, \beta(s_1, \gamma_1, y)\}}{\gamma_1 + \gamma_2} \right)^{\nu(y)} \\ T_4 &= \max_{\gamma_2 \in \Gamma} \max_{y_i \in \mathcal{Y}_i} \left[\sum_{i=1}^q \frac{1}{n} \log \binom{\lceil \bar{y}_i^{\text{up}} n \rceil}{y_i} \right] \end{aligned}$$

where

$$\begin{aligned} \Gamma &:= \{\lceil \gamma_2^* n \rceil, \lceil \gamma_2^* n \rceil + 1, \dots, d_v(\alpha n - s_1)\} \\ \mathcal{Y}_i &:= \left\{ \left\lfloor \frac{\bar{y}_i^{\text{up}} n}{2} \right\rfloor, \dots, \lceil \bar{y}_i^{\text{up}} n \rceil \right\}, \quad \text{for } i = 1, \dots, q. \end{aligned}$$

Proof: We begin with the decomposition

$$E(s_1, \gamma_1) = \binom{\lceil \alpha n \rceil}{s_1} \sum_{\gamma_2=1}^{d_v \lceil \alpha n \rceil - s_1} U_1(\gamma_1, \gamma_2) U_2(\gamma_1, \gamma_2) \quad (30)$$

where

$$\begin{aligned} U_1(\gamma_1, \gamma_2) &:= BBQ'[\exists S_2 \text{ with } (S_1, S_2) \in A_2 \\ &\quad \text{s.t. } \mathcal{B}(S_1, S_2) \mid |N(S_1)| = \gamma_1, \\ &\quad |N(F) \setminus N(S_1)| = \gamma_2, |S_1| = s_1], \\ U_2(\gamma_1, \gamma_2) &:= BBQ'[\mid N(S_1)| = \gamma_1, \\ &\quad |N(F) \setminus N(S_1)| = \gamma_2 \mid |S_1| = s_1] \end{aligned}$$

and recall that \mathbb{Q}' is the distribution \mathbb{Q} conditioned on the event $\mathcal{T}(\epsilon_1)$. We now require a lemma that allows us to restrict appropriately the range of summation over to values of γ_2 that scale linearly in n .

Lemma 10: (γ_2 small): The conditions of Lemma 9 imply that there exists some value $\gamma_2^* = \gamma_2^*(\bar{s}_{\text{crit}}) > 0$ for which the quantity

$$G(s_1, \gamma_1) := \binom{\lceil \alpha n \rceil}{s_1} \sum_{\gamma_2=1}^{\lceil \gamma_2^* n \rceil} U_1(\gamma_1, \gamma_2) U_2(\gamma_1, \gamma_2) \quad (32)$$

decays exponentially in n for any s_1, γ_1 that satisfy $\lceil \epsilon_2 n \rceil \leq s_1 \leq \lceil \bar{s}_{\text{crit}} n \rceil$ and $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$.

Proof: The proof is similar in spirit to the Proof of Lemma 6. Taking a term in the summation (30), we can bound it as follows:

$$\begin{aligned} B(s_1, \gamma_1, \gamma_2) &:= \binom{\lceil \alpha n \rceil}{s_1} U_1(\gamma_1, \gamma_2) U_2(\gamma_1, \gamma_2) \\ &\leq \binom{\lceil \alpha n \rceil}{s_1} U_2(\gamma_1, \gamma_2) \end{aligned}$$

which is upper-bounded by

$$\binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}'[\mid N(F) \setminus N(S_1)| = \gamma_2 \mid \mid N(S_1)| = \gamma_1, \mid S_1| = s_1].$$

Note that \mathbb{Q}' term is upper-bounded by

$$\binom{\lceil (1-R)n \rceil - \gamma_1}{\gamma_2} \left(\frac{\gamma_2 + \gamma_1}{\lceil (1-R)n \rceil} \right)^{(\lceil \alpha n \rceil - s_1) d_v} \quad (33)$$

so that $\frac{1}{n} \log_2 B(s_1, \gamma_1, \gamma_2) \leq \sum_{i=1}^3 C_i(s_1, \gamma_1, \gamma_2) + o(1)$, where

$$\begin{aligned} C_1 &= \alpha H \left(\frac{s_1/n}{\alpha} \right) \\ C_2 &= H \left(\frac{\gamma_2/n}{(1-R) - \gamma_1/n} \right) \\ C_3 &= d_v \left(\alpha - \frac{s_1}{n} \right) \log_2 \left(\frac{\gamma_2/n + \gamma_1/n}{(1-R)} \right). \end{aligned}$$

Using conditions (25), since $\bar{s}_{\text{crit}} < \frac{\alpha}{2}$, the term C_1 is increasing in s_1/n . Moreover, since $\alpha d_v < \frac{(1-R) - d_v \bar{s}_{\text{crit}}}{2}$, second entropy term is the term C_2 increasing in γ_1 . Finally, the term C_3 is increasing in s_1/n and in γ_1/n .

Consequently, $\frac{1}{n} \log_2 B(s_1, \gamma_1, \gamma_2)$ is upper-bounded by the function

$$\begin{aligned} b(\gamma) &:= \alpha H \left(\frac{\bar{s}_{\text{crit}}}{\alpha} \right) + H \left(\frac{\gamma}{(1-R) - d_v \bar{s}_{\text{crit}}} \right) \\ &\quad + d_v \left(\alpha - \bar{s}_{\text{crit}} \right) \log_2 \left(\frac{\gamma + d_v \bar{s}_{\text{crit}}}{(1-R)} \right). \end{aligned}$$

Note that $\lim_{\gamma \rightarrow 0} b(\gamma) < 0$ follows from the third condition in the series (25). The remainder of the proof is entirely analogous to that of Lemma 6. \square

By Lemma 10, it suffices to provide upper bounds for the terms $B(s_1, \gamma_1, \gamma_2)$ for $s_1 \in \{\lceil \epsilon_2 n \rceil, \dots, \lceil \bar{s}_{\text{crit}} n \rceil\}$, $\gamma_1 \geq \bar{\gamma}_{\text{crit}}(\epsilon_2)n$, and $\gamma_2 \geq \gamma_2^* n$. Recall the bound (33) on

$\mathbb{Q}'[|N(F) \setminus N(S_1)| = \gamma_2 \mid |N(S_1)| = \gamma_1, |S_1| = s_1]$. Similarly, recall from Appendix D that $\mathbb{Q}'[|N(S_1)| = \gamma_1 \mid |S_1| = s_1]$ is upper-bounded by

$$\binom{\lfloor (1-R)n \rfloor}{\gamma_1} \left(\frac{\gamma_1}{\lfloor (1-R)n \rfloor} \right)^{d_v s_1}.$$

Recalling the notation S_1^* for the fixed set $\{1, \dots, s_1\}$, he only missing piece is an upper bound on

$$\mathbb{Q}'[\exists(S_1^*, S_2) \in A_2 \text{ s.t. } \mathcal{B}(S_1^*, S_2) \mid N(S_1^*) = \gamma_1, |N(F) \setminus N(S_1^*)| = \gamma_2].$$

Recall that \mathbb{Q}' is the conditional probability given the event $\{\mathcal{T}(\epsilon_1)\}$. In this space, every set $S_2 \in F^c(q)$ corresponds to a request vector $y \in \prod_{i=1}^q \{0, \dots, \lceil \bar{y}_i^{\text{up}} n \rceil\}$. Moreover, for a set $S_2 \in F^c(q)$ and its corresponding request vector y , the event $\mathcal{B}(S_1^*, S_2)$ is equivalent to the following condition being satisfied:

$$\mathcal{B}(S_1^*, S_2) \Leftrightarrow |(N(S_2) \cap N(F)) - N(S_1^*)| \leq \beta(s_1, \gamma_1, y).$$

Therefore, a union bound over all the possible choices of sets S_2 (shown at the bottom of the page) gives the upper bound for the probability of interest, where $\lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$ is the probability, under the distribution \mathbb{Q}' of the event

$$\{|(N(S_2) \cap N(F)) \setminus N(S_1^*)| \leq \beta(s_1, \gamma_1, r) \mid |N(S_1)| = \gamma_1, |N(F) \setminus N(S_1^*)| = \gamma_2\}$$

where S_2 corresponds to request vector y .

In order to complete the proof, we need a final observation.

Lemma 11: For all $i = 1, \dots, q$, if $\{y_j\}_{j \neq i}, \gamma_1, \gamma_2$ are fixed, then the function $\Lambda(y_1, y_2, \dots, y_q, \gamma_1, \gamma_2)$ is increasing in the scalar variable $y_i \in \{1, 2, \dots, \lfloor \frac{\bar{y}_1^{\text{up}} n}{2} \rfloor\}$.

Proof: Clearly $\binom{\lceil \bar{y}_i^{\text{up}} n \rceil}{y_i}$ is increasing for

$$y_i \in \left\{ 1, \dots, \left\lfloor \frac{\bar{y}_1^{\text{up}} n}{2} \right\rfloor \right\}.$$

Therefore, it is enough to establish that the probability $\lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$ is increasing for

$$y_i \in \left\{ 1, \dots, \left\lfloor \frac{\bar{y}_1^{\text{up}} n}{2} \right\rfloor \right\}.$$

This fact follows from the same coupling argument used in Section IV-A: for a variable $j \in F^c$, the number of requests X_j and the size of the intersection $|N(j) \cap N(F)|$ are positively correlated. Therefore, increasing the number of edges can only increase the probability $\lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$ of the bad event $\mathcal{B}(S_1, S_2)$. \square

Using Lemma 11, we can now conclude the Proof of Lemma 9. Denote by $\mathcal{Y}_i := \{\lfloor \frac{\bar{y}_i^{\text{up}} n}{2} \rfloor, \dots, \lceil \bar{y}_i^{\text{up}} n \rceil\}$, we have that $\frac{1}{n} \log \prod_{i=1}^q \sum_{y_i=0}^{\lceil \bar{y}_i^{\text{up}} n \rceil} \lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$ is upper-bounded by

$$\max_{y_i \in \mathcal{Y}_i} \left\{ \sum_{i=1}^q \frac{1}{n} \log \binom{\lceil \bar{y}_i^{\text{up}} n \rceil}{y_i} + \frac{1}{n} \log \lambda(y_1, \dots, y_q, \gamma_1, \gamma_2) \right\}.$$

By union bound, the quantity $\frac{1}{n} \log \lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$ is upper-bounded by

$$\frac{1}{n} \log \left\{ \binom{\gamma_2}{\min\{\beta(s_1, \gamma_1, r), \gamma_2\}} \times \left(\frac{\gamma_1 + \min\{\beta(s_1, \gamma_1, r), \gamma_2\}}{\gamma_1 + \gamma_2} \right)^{\nu(r)} \right\}.$$

Putting everything together yields the claim of Lemma 9. \square

Based on the preceding analysis, we can now complete our Proof of Lemma 7. Indeed, using Lemmas 5, 6, and 9 we can upper-bound $\frac{1}{n} \log \mathbb{Q}[\mathcal{B}(A_2) \mid \mathcal{T}(\epsilon_1)]$. In this upper bound, all the relevant quantities (i.e., $s_1, \gamma_1, \gamma_2, y_1, y_2, \dots, y_q$) scale linearly with n . Therefore, standard bounds on binomial coefficients (29) lead to the claimed form of F .

F. Proof of Lemma 8

The last thing we need to do is bound the probability of the bad event $\mathbb{Q}'[\mathcal{B}(A_3)]$ (S_1 small, S_2 large). As usual, we do a union bound over S_1 sets of various sizes contracting. Define

$$D'(s_1) = \binom{\lceil \alpha n \rceil}{s_1} \mathbb{Q}'[\exists S_2 \subseteq F^c(q) \text{ with } (S_1, S_2) \in A_3 \text{ s.t. } \mathcal{B}(S_1, S_2) \mid S_1 \text{ is some fixed set of size } s_1]$$

and therefore

$$\mathbb{Q}'[\mathcal{B}(A_3)] \leq \sum_{s_1=1}^{\lfloor \epsilon_2 n \rfloor} D'(s_1).$$

Intuitively, it should be clear that this is the easiest regime, because handling requests from S_1 is always harder compared to requests from S_2 (because variables in S_2 have fewer requests). We will make ϵ_2 small enough so that the requests from S_1 are completely covered from the neighborhood of S_2 (which is always larger than a linear fraction). The function we obtain is strictly dominated by $F(\alpha)$ for sufficiently small ϵ_2 , as one would expect, since $F(\alpha)$ is satisfying the requests in a harder regime. We make a formal argument using continuity to establish this fact.

$$\sum_{y_1=0}^{\lceil \bar{y}_1^{\text{up}} n \rceil} \dots \sum_{y_q=0}^{\lceil \bar{y}_q^{\text{up}} n \rceil} \underbrace{\binom{\lceil \bar{y}_1^{\text{up}} n \rceil}{y_1} \dots \binom{\lceil \bar{y}_q^{\text{up}} n \rceil}{y_q}}_{\Lambda(y_1, y_2, \dots, y_q, \gamma_1, \gamma_2)} \lambda(y_1, \dots, y_q, \gamma_1, \gamma_2)$$

For ϵ_2 sufficiently small, we have that, for all $s_1 \in \{1, \dots, \lfloor \epsilon_2 n \rfloor\}$

$$\binom{\lfloor \alpha n \rfloor}{s_1} \leq \binom{\lfloor \alpha n \rfloor}{\lfloor \epsilon_2 n \rfloor} \leq n \left(\alpha H \left(\frac{\epsilon_2}{\alpha} \right) + o(1) \right).$$

The remainder of the analysis exploits the fact that for ϵ_2 sufficiently small and any set S_2 of size at least ρn , if y is the vector of requests from S_2 , then, with high probability

$$|N(S_2) \cap (N(F) \setminus N(S_1))| \geq \underbrace{\sum_{i=1}^q i y_i}_{\beta'(\epsilon_2, y)} + p \epsilon_2 n.$$

In words, the neighborhood of set S_2 inside $N(F) \setminus N(S_1)$ is sufficiently large not only to cover the requests from set S_2 but also from S_1 . We are going to bound the probability of failure, by only allowing S_2 to cover all the requests

$$\begin{aligned} & \mathbb{Q}'[\exists S_2 \subseteq F^c(q) \text{ with } (S_1, S_2) \in A_3 \\ & \quad \text{s.t. } \mathcal{B}(S_1, S_2) |_{S_1} \text{ some fixed set of size } s_1] \leq \\ & \mathbb{Q}'[\exists S_2 \subseteq F^c(q) \text{ with } (S_1, S_2) \in A_3 \\ & \quad \text{s.t. } |N(S_2) \cap (N(F) - N(S_1))| \leq \\ & \quad \beta'(\epsilon_2, y), |S_1| \text{ some fixed set of size } s_1]. \end{aligned}$$

By similar analysis as in the Proof of Lemma 9, we obtain $\frac{1}{n} \log D'(S_1) \leq F'(\alpha, \epsilon_2) + o(1)$, where

$$F'(\alpha, \epsilon_2) := \sup_{\bar{\gamma}_2 \in [0, d_v \alpha]} \sup_{\bar{y}_i \in [\bar{y}_i^{\text{LP}}/2, \bar{y}_i^{\text{UP}}]} G'(\bar{\gamma}_2, \bar{y}_1, \bar{y}_2, \dots, \bar{y}_q, \epsilon_2)$$

and the intermediate function

$$\begin{aligned} G' &= G'(\bar{\gamma}_2, \bar{y}_1, \dots, \bar{y}_q, \epsilon_2) \\ &= \sum_{i=1}^2 \min\{0, G'_i(\bar{\gamma}_2, \bar{y})\} + G'_3(\bar{\gamma}_2, \bar{y}) \end{aligned}$$

has terms

$$\begin{aligned} G'_1 &= ((1-R) - d_v \epsilon_2) H \left(\frac{\bar{\gamma}_2}{((1-R) - d_v \epsilon_2)} \right) \\ & \quad + d_v (\alpha - \epsilon_2) \log_2 \left(\frac{d_v \epsilon_2 + \bar{\gamma}_2}{(1-R)} \right) \\ G'_2 &= \bar{\gamma}_2 H \left(\frac{\min\{\bar{\gamma}_2, \bar{\beta}'(\epsilon_2, \bar{y})\}}{\bar{\gamma}_2} \right) \\ & \quad + \nu(\bar{y}) \log_2 \left(\frac{d_v \epsilon_2 + \min\{\bar{\gamma}_2, \bar{\beta}'(\epsilon_2, \bar{y})\}}{d_v \epsilon_2 + \bar{\gamma}_2} \right) \\ G'_3 &= \alpha H \left(\frac{\epsilon_2}{\alpha} \right) + \sum_{i=1}^q \bar{y}_i^{\text{LP}} H \left(\frac{\bar{y}_i}{\bar{y}_i^{\text{LP}}} \right). \end{aligned}$$

Note that

$$\begin{aligned} \lim_{\epsilon_2 \rightarrow 0} G'(\bar{\gamma}_2, \bar{y}_1, \dots, \bar{y}_q, \epsilon_2) \\ = \lim_{\bar{s}_1 \rightarrow 0, \bar{\gamma}_1 \rightarrow 0} G(\bar{s}_1, \bar{\gamma}_1, \bar{\gamma}_2, \bar{y}_1, \dots, \bar{y}_q, \epsilon_2) \end{aligned}$$

where the limit is taken by setting $\bar{\gamma}_1 = \Theta(\bar{s}_1)$ (which will be true by concentration). Therefore, we have

$$\lim_{\epsilon_2 \rightarrow 0} F'(\alpha, \epsilon_2) \leq F(\alpha).$$

Consequently, if $F(\alpha) < 0$, it then follows that $\lim_{\epsilon_2 \rightarrow 0} F'(\alpha, \epsilon_2) < 0$. By continuity, there exists some value $\epsilon_2 > 0$ such that $F'(\alpha, \epsilon_2) < 0$; for this value of ϵ_2 , the probability $\mathbb{Q}[\mathcal{B}(A_3) | \mathcal{T}(\epsilon_1)]$ decreases exponentially in n .

ACKNOWLEDGMENT

The authors would like to thank the anonymous referees for their numerous and detailed suggestions that helped to improve this paper.

REFERENCES

- [1] A. Barg and G. Zémor, "Error exponents of expander codes under linear-complexity decoding," *SIAM J. Discr. Math.*, vol. 17, no. 3, pp. 426–445, 2004.
- [2] E. Berlekamp, R. McEliece, and H. van Tilborg, "On the inherent intractability of certain coding problems," *IEEE Trans. Inf. Theory*, vol. IT-24, no. 3, pp. 384–386, May 1978.
- [3] D. Burshtein and G. Miller, "Expander graph arguments for message-passing algorithms," *IEEE Trans. Inf. Theory*, vol. 48, no. 2, pp. 782–790, Feb. 2002.
- [4] M. Chertkov, "Reducing the error floor," in *Proc. IEEE Information Theory Workshop ITW*, 2007.
- [5] M. Chertkov and M. G. Stepanov, "An Efficient Pseudo-Codeword Search Algorithm for Linear Programming Decoding of LDPC Codes," Los Alamos National Laboratory, Tech. Rep., Sep. 2006 [Online]. Available: arxivss.IT/0601113
- [6] S. Y. Chung, T. Richardson, and R. Urbanke, "Analysis of sum-product decoding of low-density parity check codes using a Gaussian approximation," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 657–670, Feb. 2001.
- [7] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, 1991.
- [8] A. G. Dimakis and M. J. Wainwright, "Guessing facets: Improved LP decoding and polytope structure," in *Proc. IEEE Int. Symp. Information Theory*, Seattle, WA, Jul. 2006, pp. 1369–1373.
- [9] S. Draper, J. Yedidia, and Y. Wang, "ML decoding via mixed-integer adaptive linear programming," in *Proc. IEEE Int. Symp. Information Theory*, Nice, France, Jun. 2007, pp. 1656–1660.
- [10] J. Feldman and D. Karger, "Decoding turbo-like codes in polynomial time with provably good error-correcting performance via linear programming," in *Proc. Conf. Foundations of Computer Science (FOCS)*, Jul. 2002.
- [11] J. Feldman, D. R. Karger, and M. J. Wainwright, "Linear programming-based decoding of turbo-like codes and its relation to iterative approaches," in *Proc. 40th Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Oct. 2002.
- [12] J. Feldman, T. Malkin, R. A. Servedio, C. Stein, and M. J. Wainwright, "LP decoding corrects a constant fraction of errors," *IEEE Trans. Inf. Theory*, vol. 53, no. 1, pp. 82–89, Jan. 2007.
- [13] J. Feldman and C. Stein, "LP decoding achieves capacity," in *Proc. Symp. Discrete Algorithms (SODA'05)*, Jan. 2005.
- [14] J. Feldman, M. J. Wainwright, and D. R. Karger, "Using linear programming to decode binary linear codes," *IEEE Trans. Inf. Theory*, vol. 51, no. 3, pp. 954–972, Mar. 2005.
- [15] R. G. Gallager, *Low-Density Parity Check Codes*. Cambridge, MA: MIT Press, 1963.
- [16] N. Halabi and G. Even, "Improved bounds on the word error probability of RA(2) codes with linear-programming-based decoding," *IEEE Trans. Inf. Theory*, vol. 51, no. 1, pp. 265–280, Jan. 2005.
- [17] W. Hoeffding, "Probability inequalities for sums of bounded random variables," *J. Amer. Statist. Assoc.*, vol. 58, pp. 13–30, 1963.
- [18] C. A. Kelley and D. Sridhara, "Eigenvalue bounds on pseudocodeword weight of expander codes," *AMC Trans. Adv. Math. of Commun.*, 2007.
- [19] R. Koetter and P. O. Vontobel, "Graph-covers and iterative decoding of finite length codes," in *Proc. 3rd Int. Symp. Turbo Codes*, Sep. 2003.
- [20] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 498–519, Feb. 2001.
- [21] M. Luby, M. Mitzenmacher, M. A. Shokrollahi, and D. Spielman, "Improved low-density parity check codes using irregular graphs," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 585–598, Feb. 2001.

- [22] D. MacKay, "Good error correcting codes based on very sparse matrices," *IEEE Trans. Inf. Theory*, vol. 45, no. 2, pp. 399–431, Mar. 1999.
- [23] T. Richardson and R. Urbanke, "The capacity of low-density parity check codes under message-passing decoding," *IEEE Trans. Inf. Theory*, vol. 47, no. 2, pp. 599–618, Feb. 2001.
- [24] T. Richardson and R. Urbanke, *Modern Coding Theory*. Cambridge, U.K.: Cambridge Univ. Press, 2007.
- [25] M. Sipser and D. Spielman, "Expander codes," *IEEE Trans. Inf. Theory*, vol. 42, no. 6, pp. 1710–1722, Nov. 1996.
- [26] M. H. Taghavi and P. H. Siegel, "Adaptive linear programming decoding," in *Proc. IEEE Int. Symp. Information Theory*, Seattle, WA, Jul. 2006, pp. 1374–1378.
- [27] J. H. van Lint and R. M. Wilson, *A Course in Combinatorics*. Cambridge, U.K.: Cambridge Univ. Press, 1992.
- [28] P. O. Vontobel and R. Koetter, "Bounds on the threshold of linear programming decoding," in *Proc. IEEE Information Theory Workshop*, Mar. 2006.
- [29] P. O. Vontobel and R. Koetter, "Toward low-complexity linear-programming decoding," in *Proc. 4th Int. Symp. Turbo Codes and Related Topics*, Munich, Germany, Apr. 2006.
- [30] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, "MAP estimation via agreement on (hyper) trees: Message-passing and linear programming approaches," in *Proc. Allerton Conf. Communication, Control and Computing*, Monticello, IL, Oct. 2002.
- [31] M. J. Wainwright, T. S. Jaakkola, and A. S. Willsky, "Exact MAP estimates via agreement on (hyper) trees: Linear programming and message-passing," *IEEE Trans. Inf. Theory*, vol. 51, no. 11, pp. 3697–3717, Nov. 2005.
- [32] K. Yang, X. Wang, and J. Feldman, "Cascaded formulation of the fundamental polytope of general linear block codes," in *Proc. IEEE Int. Symp. Information Theory*, Nice, France, Jun. 2007, pp. 1361–1365.