Finite-Length LDPC Codes on the q-ary Multi-Bit Channel

Rami Cohen and Yuval Cassuto
Andrew and Erna Viterbi Faculty of Electrical Engineering, Technion - Israel Institute of Technology
rc@campus.technion.ac.il, ycassuto@ee.technion.ac.il

***Abstract*—In this paper, we address the finite-length decoding performance of LDPC codes over the q-ary multi-bit channel (QMBC). The QMBC is defined over the full q-ary symbols, while addressing the differences in reliability between the bits composing the symbols. We show that unlike the binary erasure channel, the QMBC iterative decoder does not necessarily halt at stopping sets. Instead, its performance depends on the edge-label configuration of the LDPC code graph. We characterize good edge-label configurations, and propose an edge-labeling algorithm for improved iterative-decoding performance. We then provide finite-length maximum-likelihood decoding analysis for both the standard non-binary random ensemble and LDPC ensembles. Finally, simulations are presented to demonstrate the advantages of the proposed edge-labeling algorithm.**

I. Introduction

In a memory device, information is often stored in the form of $q = 2^s$ voltage/current level ranges. As an example, Flash memory chips with triple-level cell (TLC) technology use $q = 8$ levels. It is a common practice to split the q levels to s bits, and encode those bits with a binary code. However, by doing so one eliminates the relations between bits of the same symbol, and may hence lose a lot in coding performance. At the other extreme, applying a q-ary code for symmetric errors also loses for ignoring the real error characteristics, e.g. higher likelihood for errors with small magnitudes. Our objective in this paper is to follow a middle path of coding the s bits of the symbol jointly, but considering the different effect symbol errors have on each bit. In particular, we are interested in the most natural model where the s bits form a hierarchy representing the q-ary symbol from the least significant bit to the most significant bit.

We study coding for a channel model we call the *q-ary multi-bit channel (QMBC)*. In the QMBC[1] model, a stored q-ary symbol is represented by s bits, and the channel parameters are the probabilities to lose the j lower bits of the symbol, for $j = 1, \ldots, s$. This results in a channel falling under the class of *partial-erasure channels* [2], and mimicking specifically the graded reliabilities of the s bits. One use of this channel is when the level-measurement process can return partial-precision read values. Another use is as a loyal and theoretically manageable proxy for designing LDPC codes for graded-magnitude errors, similarly to binary erasures being a good proxy for symmetric bit errors.

In this paper, we are interested in design and analysis of *finite-length* LDPC codes for the QMBC. When iterative decoding is applied over the QMBC, in addition to the stopping sets [3], the finite-length performance depends strongly on the edge labels. We theoretically characterize this dependence by analyzing the algebraic structure of the partial-erasure sets within the finite field, and propose an edge-labeling algorithm that considerably mitigates the harmful effect of stopping sets (Section III). In that, our work extends previous label-optimization algorithms (e.g., [4], [5]) to the special structure of the QMBC. The advantage here is that the QMBC has strong solvability conditions that are local to a single check, and thus allow neutralizing stopping sets even without relying on the cycle structure of the graph. Later in Section IV we study the QMBC finite-length maximum-likelihood decoding performance, both for the standard non-binary ensemble and regular LDPC ensembles. Because QMBC erasures are *subsets* of the field GF(q), the main analytical challenge here is in losing the linear structure. Finally, simulation results in Section V show that our edge-labeling algorithm offers significant improvement over uniform labeling, and even more so compared to using a binary LDPC code.

[1]In our previous work [1] we used the different initials QBMC for the same channel, describing its more specific application as q-ary bit-measurement channel.

II. Preliminaries

A. Channel model

The QMBC input alphabet consists of $q = 2^s$ symbols: $\mathcal{X} = \{0, 1, ..., q-1\}$, for some integer s. For each input symbol x and $j = 0, 1, 2, ..., s$, a *partial-erasure* event occurs when only the $s - j$ left bits of x (in binary representation) are known. The output in this case is a *set* of 2^j consecutive symbols that have the same $s-j$ left bits as x. We denote this output set by $\mathcal{M}_x^j$. Note that the correct input symbol always belongs to the output set. In addition, the input symbol is completely known when $j = 0$. The transition probabilities governing the QMBC are:

$$\Pr\left(Y = \mathcal{M}_x^j \middle| X = x\right) = \varepsilon_j, \tag{1}$$

where ε_j for $j = 0, 1, ..., s$ are the partial-erasure probabilities. Note that for $q = 2$ the QMBC is equivalent to the binary erasure channel (BEC).

B. GF(q) LDPC codes and set iterative decoder

For analysis purposes, we map the symbols in $\mathcal{X}$ to GF($q = 2^s$) elements. Consider a basis $\{\omega_1, \omega_2, ..., \omega_s\}$ of GF($q = 2^s$) over GF(2). Denote by $\langle \omega_1, \omega_2, ..., \omega_j \rangle$ the span of the basis elements $\omega_1, \omega_2, ..., \omega_j$ for $j = 1, 2, ..., s$. As an example, $\langle \omega_1, \omega_2 \rangle = \{a \cdot \omega_1 + b \cdot \omega_2 : a, b \in \{0, 1\}\}$. We map the sets $\mathcal{M}_0^j$ for $j = 1, 2, ..., s$ to $\langle \omega_1, \omega_2, ..., \omega_j \rangle$, which are *subgroups* of the additive group of GF(q). More generally, for each $j = 1, 2, ..., s$ and $x \in \mathcal{X}$ we map $\mathcal{M}_x^j$ to one of the 2^{s-j} cosets of $\langle \omega_1, \omega_2, ..., \omega_j \rangle$, where the coset representatives are taken from $\langle \omega_{j+1}, \omega_{j+2}, ..., \omega_s \rangle$.

Example 1. Let α designate a root of the primitive polynomial $x^2 + x + 1$ such that $\{1, \alpha\}$ is a basis of GF(4) over GF(2). The sets $\mathcal{M}_0^0, \mathcal{M}_0^1$ and $\mathcal{M}_0^2$ are mapped to the subgroups $\{0\}, \{0, 1\}$ and $\{0, 1, \alpha, \alpha + 1\}$, respectively. The cosets of $\{0, 1\}$ are $\{0, 1\}$ and $\{\alpha, \alpha + 1\}$. Thus, $\mathcal{M}_1^1$ is mapped to $\{0, 1\}$, while $\mathcal{M}_2^1$ and $\mathcal{M}_3^1$ are mapped to $\{\alpha, \alpha + 1\}$.

The error-correcting codes we consider for dealing with the QMBC are GF(q) LDPC codes [6]. These codes are defined by a sparse parity-check matrix with elements taken from GF(q). This matrix is commonly visualized as a (bipartite) graph with *variable* (left) nodes corresponding to codeword symbols, and *check* (right) nodes corresponding to parity-check equations. The edge labels on the graph are taken from the non-zero elements of GF(q). For clarity and simplicity, we concentrate on regular (d_v, d_c) LDPC codes.

Since the QMBC belongs to the class of partial-erasure channels, we use the iterative decoder suggested for such channels in [2]. In this decoder, sets of symbols are exchanged as messages in the decoding process. As usual, we have *variable-to-check* (VTC) and *check-to-variable* (CTV) messages. A CTV message $\mathrm{CTV}_{\mathtt{c}\to\mathtt{v}}^{(l)}$ contains all the possible symbol values of $\mathtt{v}$ that satisfy the parity-check equation at $\mathtt{c}$ given the VTC messages to $\mathtt{c}$ at iteration $l-1$. To simplify its calculation, we use the *sumset* operation, defined for two sets $\mathcal{A}$ and $\mathcal{B}$ that contain GF(q) elements as

$$\mathcal{A}+\mathcal{B} \triangleq \{a+b : a\in\mathcal{A}, b\in\mathcal{B}\}, \tag{2}$$

where the addition is performed using the GF(q) arithmetic. Denote the nodes adjacent to $\mathtt{c}$ by $\mathcal{N}(\mathtt{c})$. The CTV message from check node $\mathtt{c}$ to variable node $\mathtt{v}$ is then:

$$\mathrm{CTV}_{\mathtt{c}\to\mathtt{v}}^{(l)} = \sum_{\mathtt{v}'\in\{\mathcal{N}(\mathtt{c})\setminus\mathtt{v}\}} \left(\frac{h_{\mathtt{c},\mathtt{v}'}}{h_{\mathtt{c},\mathtt{v}}}\right)\cdot \mathrm{VTC}_{\mathtt{v}'\to\mathtt{c}}^{(l-1)}, \tag{3}$$

where the sum is a sumset operation and the multiplications are performed element-wise. Once all the CTV messages are calculated, a VTC message is the *intersection* of the channel-information set and the incoming CTV message sets:

$$\mathrm{VTC}_{\mathtt{v}\to\mathtt{c}}^{(l)} = \mathrm{VTC}_{\mathtt{v}}^{(0)} \bigcap \left\{ \bigcap_{\mathtt{c}'\in\{\mathcal{N}(\mathtt{v})\setminus\mathtt{c}\}} \mathrm{CTV}_{\mathtt{c}'\to\mathtt{v}}^{(l)} \right\}. \tag{4}$$

A decoding failure occurs if unresolved variable nodes remain after the decoder terminates.

III. Edge-labeling Algorithm for Improved Finite-Length Performance

A stopping set $\mathcal{S}$ is defined as a subset of variable nodes, such that all the neighboring check nodes of $\mathcal{S}$ are connected to $\mathcal{S}$ at least twice. A key result in BEC finite-length analysis is that (fully) erased variable nodes in the maximal stopping set remain erased when the decoder stops [3]. *For the QMBC, however, partially-erased coordinates within a stopping set as defined above do not necessarily halt the iterative decoder, and may eventually be resolved*. The reason is that with partial erasures the iterative decoder can make progress even if two or more neighbours of a check node are partially erased. This is demonstrated in the following example.

Example 2. Consider the Tanner graph in Figure 1, where the variable nodes $\mathtt{v}_1$ and $\mathtt{v}_2$ form a stopping set. Suppose $\mathtt{v}_1$ and $\mathtt{v}_2$ are partially erased, each with the set $\{0,1\}$ ($q=4$ is assumed). The initial CTV messages from the lower check node are $\{0, h_1/h_2\}$ to $\mathtt{v}_1$ and $\{0, h_2/h_1\}$ to $\mathtt{v}_2$. If $h_1 = h_2$, the variable nodes are not resolved, as the intersection operation at variable nodes results in $\{0,1\}$. But in all label combinations satisfying $h_1 \neq h_2$, they are resolved as $\{0\}$.

As shown in Example 2, the edge labels play an important role in the resolvability of the stopping sets. Still, non-resolved partial erasures must belong to a stopping set. Denote by $\mathcal{E}$ the index set of partially-erased variable nodes.

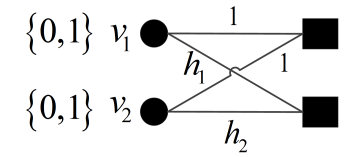

Fig. 1: $\mathtt{v}_1$ and $\mathtt{v}_2$ form a partially-erased stopping set (the channel information sets appear to the left). The resolvability of $\mathtt{v}_1$ and $\mathtt{v}_2$ depends on the value of the edge labels h_1, h_2.

Lemma 1. *The variable nodes that remain unresolved when the iterative QMBC decoder terminates belong to the maximal stopping set contained in $\mathcal{E}$.*

This lemma is based on the intuitive fact that for a variable node to remain partially-erased, each of its adjacent check nodes must observe at least two partial erasures. Considering Example 2 and Lemma 1, an iterative-decoding failure occurs on stopping sets whose edges satisfy a certain edge-labeling configuration. We leverage this observation and propose an edge-labeling algorithm for improved finite-length iterative-decoding performance. Our approach is to increase the probabilty of resolving a stopping set with the QMBC *iterative* decoder, unlike the common assumption in the literature [7] that the stopping sets are decoded with a maximum-likelihood decoder.In the following, we assume the transmission of the all-zero codeword [1]. The partial-erasure sets on the stopping sets are assumed to be of the maximal order $j_{\max}$ allowed by the channel parameters, that is, $\mathcal{M}_0^{j_{\max}}$. Consider a check node connected to κ variable nodes $\mathtt{v}_1, \mathtt{v}_2, ..., \mathtt{v}_\kappa$ partially-erased to the set $\mathcal{M}_0^{j_{\max}}$ via edges labeled $h_1, h_2, ..., h_\kappa$.

Definition 1. *The edge labels $h_1, h_2, ..., h_\kappa$ are said to be κ-resolvable if $\mathtt{v}_1, \mathtt{v}_2, ..., \mathtt{v}_\kappa$ are resolvable (i.e., decoded correctly), independently of other variable nodes.*

The existence of resolvable edge labels is proved in the following theorem. We consider the basis $\{1, \alpha, \alpha^2, ..., \alpha^{s-1}\}$ to GF(q) over GF(2), where α is a primitive element of GF(q). In this case, the elements in $\mathcal{M}_0^{j_{\max}}$ are all the polynomials in α of degree at most $j_{\max}-1$.

Theorem 2. *The edge labels $\left\{\alpha^{ij_{\max}}\right\}_{i=0}^{\lfloor s/j_{\max}\rfloor-1}$ are $\lfloor s/j_{\max}\rfloor$-resolvable.*

Proof. Suppose that $\kappa \leq \lfloor s/j_{\max}\rfloor$ variable nodes are connected to a check node with distinct edge labels taken from $\left\{\alpha^{ij_{\max}}\right\}_{i=0}^{\lfloor s/j_{\max}\rfloor-1}$. Denote the symbol of variable node $\mathtt{v}_i$ by x_i, and its edge label by h_i; we claim that the only solution to the parity-check equation $\sum_{i=1}^{\kappa} h_i \cdot x_i = 0$ is the trivial one. To see this, observe that h_i are represented by monomials in α having degrees separated by at least $j_{\max}$. That implies non-zero polynomials in the sets $h_i \cdot \mathcal{M}_0^{j_{\max}}$ for distinct i have different degrees. As polynomials of different degrees cannot sum to zero, no non-trivial solution exists in this case. From the zero-codeword assumption, having no non-trivial solution implies the resolvability of the κ variable nodes. □

Based on Lemma 1 and Theorem 2, we propose an edge-labeling algorithm for improved finite-length decoding performance. From now on we set $\kappa = \lfloor s/j_{\max}\rfloor$ and use the edge labels of Theorem 2. Note that any κ'-subset of these edge labels is also resolvable. We initialize the edge labels to be uniformly selected non-zero elements from GF(q).

Algorithm 1. *(Edge labeling)*

1) Run the BEC iterative decoder with the channel parameter $\varepsilon = \varepsilon_{j_{\max}}$ for a predefined number of times. After each run, store the set of unresolved variable nodes.
2) Initialize Σ as the subgraph induced by the variable nodes from the sets of Step 1. Rank the variable nodes by their number of occurrences in the sets.
3) Modify the edge labels of check nodes of degree at most κ in Σ to resolvable edge labels. Set the rank of connected variable nodes to 0.
4) Run over the sets found in Step 1 by ascending cardinality. For each check node connected to a set:
   a) Set κ' as the minimum between κ and the number of non-zero ranking variable nodes.
   b) Modify κ' edge labels connected to non-zero highest-ranking variable nodes to κ'-resolvable edge labels.

The steps of Algorithm 1 are explained as follows. First, we focus on stopping sets, as they are a necessary condition for a decoding failure. Finding stopping sets is computationally hard [8], so we run the BEC decoder that fails on stopping sets. Resolvable edge labels are then distributed on edges connected to stopping sets, prioritizing edges connected to variable nodes common among the sets found in Step 1. A major advantage of Algorithm 1 is that the graph topology remains unchanged, such that the degree distributions are not affected. The performance improvement of the algorithm is shown in Section V.

IV. Finite-Length Analysis of Maximum-Likelihood Decoding

In this section, we study the QMBC maximum-likelihood (ML) decoding performance for both the standard non-binary linear ensemble and LDPC ensembles, as a function of the code length. As in the iterative-decoding case, we assume (w.l.o.g.) that the all-zero codeword was transmitted. Denote by $\mathcal{E}_j$ the index set of the variable nodes that are partially-erased to the set $\mathcal{M}_0^j$ ($j = 1, 2, ..., s$), and define $\mathcal{E} \triangleq \bigcup_{j=1}^{s} \mathcal{E}_j$. An ML decoding failure occurs if and only if there exists a non-zero codeword, whose symbols indexed in $\mathcal{E}_j$ belong to $\mathcal{M}_0^j$. This leads us to the following definition.

Definition 2. *A vector is said to be consistent with respect to $\{\mathcal{E}_j\}_{j=1}^s$ if at every coordinate indexed in $\mathcal{E}_j$ it has a symbol that belongs to $\mathcal{M}_0^j$.*

Example 3. Suppose $q = 4$, $n = 3$, $\mathcal{E}_1 = \{2\}$ and $\mathcal{E}_2 = \{3\}$. There are 8 consistent vectors: $(0,0,0)$ (the all-zero codeword is always consistent), $(0,1,0), (0,0,1), (0,1,1), (0,0,2), (0,1,2), (0,0,3)$ and $(0,1,3)$.

A. Standard non-binary random ensemble

Consider the standard non-binary random ensemble (SNBRE) of linear codes. Each code in the SNBRE is defined by a parity-check matrix $\mathbf{H}$ of dimensions $(n-k)\times n$, whose entries are i.i.d. uniform random variables taken from the GF(q) elements. For each $\mathbf{H}$, we denote by $\mathbf{H}_\mathcal{E}$ the submatrix formed by the columns indexed in $\mathcal{E}$. We start with the following definition.

Definition 3. *The columns of $\mathbf{H}_\mathcal{E}$ are said to be partially linearly independent if no non-zero consistent vector exists in the null space of $\mathbf{H}_\mathcal{E}$.*

Definition 3 reduces to the ordinary linear independence definition when the partial erasures are full erasures. But in our case, the columns of $\mathbf{H}_\mathcal{E}$ can be partially linearly independent even if they are not linearly independent under the ordinary definition. This is because the GF(q) coefficients that map $\mathbf{H}_\mathcal{E}$ to the zero vector may be outside the sets in $\mathcal{E}_j$, and thus not consistent. An ML decoding failure occurs if and only if the columns of $\mathbf{H}_\mathcal{E}$ are partially linearly dependent. Let us define the set

$$\mathcal{M}_0^{j,j'} \triangleq \left\{ \frac{h_j}{h_{j'}} : h_j \in \mathcal{M}_0^j, h_{j'} \in \mathcal{M}_0^{j'} \setminus \{0\} \right\}, \quad (5)$$

obtained by an element-wise division of the set $\mathcal{M}_0^j$ by $\mathcal{M}_0^{j'} \setminus \{0\}$ (for certain $j, j' \leq s$). We denote the cardinality of $\mathcal{M}_0^{j,j'}$ by $\chi^{j,j'}$:

$$\chi^{j,j'} \triangleq \left| \mathcal{M}_0^{j,j'} \right|. \quad (6)$$

Example 4. *Consider the finite field GF(4). Then $\chi^{1,1} = |\{0,1\}| = 2$ and $\chi^{j,j'}$ for $j \neq 1$ or $j' \neq 1$ are 4.*

Let ψ denote the probability that the columns of a randomly drawn $\mathbf{H}_\mathcal{E}$ are partially linearly independent. We make the dependence of ψ on $\{\mathcal{E}_j\}_{j=1}^s$ implicit for notational convenience. We define $x^+ \triangleq \max(0, x)$.

Lemma 3. *Given $\{\mathcal{E}_j\}_{j=1}^s$, let $\mathcal{O}$ contain all vectors of length $|\mathcal{E}|$ in which j occurs $|\mathcal{E}_j|$ times. Then*

$$\psi \geq \max_{\boldsymbol{o}\in\mathcal{O}} \prod_{i=1}^{|\mathcal{E}|} \left(1 - \left(\prod_{l=1}^{i-1} \chi^{o_l, o_i} \right) / q^{n-k} \right)^+. \quad (7)$$

Proof. As the matrices in the SNBRE are equiprobable, ψ is a function of $\{|\mathcal{E}_j|\}_{j=1}^s$ rather than of $\{\mathcal{E}_j\}_{j=1}^s$. Let us concentrate on some fixed but arbitrary choice of index sets with cardinalities $\{|\mathcal{E}_j|\}_{j=1}^s$. This choice is represented by a vector $\boldsymbol{o}$ that contains j in indices of codeword symbols partially-erased to $\mathcal{M}_0^j$. Consider a matrix $\mathbf{H}_\mathcal{E}$ with columns $\boldsymbol{e}_i$ and denote by $\mathcal{A}_i$ the partial-erasure sets indexed in o_i ($i = 1, 2, ..., |\mathcal{E}|$). We count in how many ways partially linearly independent columns can be placed in $\mathbf{H}_\mathcal{E}$.

Assume that the first $i'-1$ columns of $\mathbf{H}_\mathcal{E}$ are partially linearly independent. The next column, $\boldsymbol{e}_{i'}$, must satisfy $\{\mathbf{e}_{i'} \cdot \mathcal{A}_{i'}\} \cap \left\{ \sum_{l=1}^{i'-1} \mathbf{e}_l \cdot \mathcal{A}_l \right\} = \emptyset$. Thus, $\boldsymbol{e}_{i'}$ must not be contained in the set $\Gamma = \sum_{l=1}^{i'-1} \mathbf{e}_l \cdot \{\mathcal{A}_l / \{\mathcal{A}_{i'} \setminus 0\}\}$. $|\Gamma|$ is *upper bounded* by $\prod_{l=1}^{i'-1} \chi^{o_l, o_i}$ as the linear combinations of $\boldsymbol{e}_l$ in Γ might not be distinct. We maximize over $\boldsymbol{o} \in \mathcal{O}$ to tighten the bound, and to obtain a probability we normalize by $q^{(n-k)|\mathcal{E}|}$, which is the number of possible $\mathbf{H}_\mathcal{E}$ matrices. □

In fact, we can obtain the exact value of ψ in certain cases. Consider a subset $\mathcal{J}^*$ of $\{1, 2, ..., s\}$ such that each element in $\mathcal{J}^*$ divides s and j' divides j for each $j, j' \in \mathcal{J}^*$, $j' \leq j$. As an example, the possible choices of $\mathcal{J}^*$ for $q = 4$ are $\{1\}$, $\{2\}$ and $\{1, 2\}$. For the following lemma, we assume that the basis elements $\{\omega_1, \omega_2, ..., \omega_s\}$ are chosen such that for j that divides s, the span of $\{\omega_1, \omega_2, ..., \omega_j\}$ forms a *subfield* of GF(q).

Lemma 4. *Assume that $\mathcal{E}_j = \emptyset$ for $j \notin \mathcal{J}^*$. Denote by $\boldsymbol{o}$ the (now specific) vector of length $|\mathcal{E}|$ with s in its first $|\mathcal{E}_s|$*

entries, $s-1$ *in its next* $|\mathcal{E}_{s-1}|$ *entries downto* 1 *in its last* $|\mathcal{E}_1|$ *entries. Then*

$$\psi = \prod_{i=1}^{|\mathcal{E}|} \left(1 - \left(\prod_{l=1}^{i-1} 2^{o_l}\right) / q^{n-k}\right)^{+}. \tag{8}$$

Proof. We first note that if j divides s, the elements in $\mathcal{M}_0^j$ with the arithmetic operations of GF($q=2^s$) form a *subfield* of GF(q). In a similar manner, $\mathcal{M}_0^{j'}$ is a subfield of $\mathcal{M}_0^j$ if j' divides j. This implies that $\chi^{j,j'} = 2^j$ for $j' \leq j$ in $\mathcal{J}^*$. Consider the placement process depicted in the proof of Lemma 3 and assume that we place the i'th column. From the ordering of $\boldsymbol{o}$ we get that $\chi^{o_l,o_{i'}} = 2^{o_l}$ for $l < i'$. In choosing the vector $\boldsymbol{e}_{i'}$ we exclude all combinations of previous vectors $\boldsymbol{e}_l$ with coefficients in $\mathcal{M}_0^{o_l,o_{i'}}$. Assume by contradiction that two of these $\prod_{l=1}^{i'-1} 2^{o_l}$ combinations result in the same vector. But this would imply an $\boldsymbol{e}_{i''}$, $i'' < i'$ that is a combination of vectors $\boldsymbol{e}_l$, $l < i''$, with coefficients in $\mathcal{M}_0^{o_l,o_{i'}}$. Since for any l , $\mathcal{M}_0^{o_l,o_{i'}} = \mathcal{M}_0^{o_l,o_{i''}}$, this is a contradiction because it means that at step i'' we did not exclude all partially dependent vectors, and thus the count is exact with no over-subtraction. □

Let us denote by $P^{\mathrm{ML}}(\mathbf{H})$ the probability of decoding failure of a particular linear code defined by the parity-check matrix $\mathbf{H}$, given the QMBC parameters $\{\varepsilon_j\}_{j=0}^s$. We now calculate the expected value of $P^{\mathrm{ML}}(\mathbf{H})$, over codes (parity-check matrices) drawn from the SNBRE.

Theorem 5.

$$\mathbb{E}_{\mathrm{SNBRE}}\left[P^{\mathrm{ML}}(\mathbf{H})\right] \leq \sum_{\substack{|\mathcal{E}_0|,|\mathcal{E}_1|,\ldots,|\mathcal{E}_s|:\\ \sum_{j=0}^{s}|\mathcal{E}_j|=n}} \frac{n!}{|\mathcal{E}_0|!\,|\mathcal{E}_1|!\ldots|\mathcal{E}_s|!} \prod_{j=0}^{s} \varepsilon_j^{|\mathcal{E}_j|} \cdot \left(1-\tilde{\psi}\right), \tag{9}$$

where $\tilde{\psi}$ *is either the lower bound of Lemma 3, or its exact value in the cases of Lemma 4 (in the latter cases, an equality is attained in* (9)*).*

Note that when all the partial-erasure sets are full erasures, the $\boldsymbol{o}^*$ of Lemma 3 is the all-s vector. In the BEC case ($s=1$), [3, Theorem 3.1] is obtained as a special case of Theorem 5. In Figure 2, we plot $\mathbb{E}_{\mathrm{SNBRE}}\left[P^{\mathrm{ML}}(\mathbf{H})\right]$ for a $q=4$ channel with $\varepsilon_2 = \varepsilon_1/10$ and different n values. This is compared to an asymptotically equivalent q-ary erasure channel (QEC), i.e., with $\varepsilon = \varepsilon_1/2 + \varepsilon_2$. It is demonstrated that the QMBC finite-length ML performance is orders of magnitude better, though the Shannon limit is the same.

B. LDPC ensembles

In this part we derive the probability of ML decoding failure for the regular (d_v, d_c) LDPC ensemble over the QMBC. We start with the following lemma, which will serve us later in calculating the probability that a certain check node is satisfied.

Lemma 6. *Consider a vector* $\boldsymbol{a}$ *of length* $m \geq 2$*, whose entries are i.i.d. random variables uniformly distributed on the non-zero GF(*$q=2^s$*) elements. Then*

$$\Pr\left(\sum_{i=1}^{m} a_i = 0\right) = \frac{1-(1-q)^{1-m}}{q} \leq \frac{1}{q-1}. \tag{10}$$

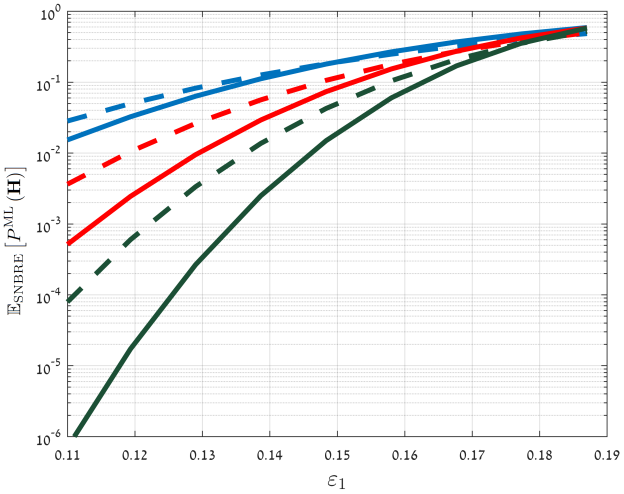

Fig. 2: Exact $\mathbb{E}_{\mathrm{SNBRE}}\left[P^{\mathrm{ML}}(\mathbf{H})\right]$ as a function of ε_1, for $\varepsilon_2 = \varepsilon_1/10$ and $q=4$ (solid lines). An asymptotically equivalent QEC with $\varepsilon = (3/5)\,\varepsilon_1$ is also shown (dashed lines). The codeword lengths are $n = 128, 256, 512$ (top to bottom) and the rate is $8/9$ (Shannon limit: 0.185).

The upper bound in (10) is tight, as it is attained for $m=2$. We now calculate the number of consistent vectors (see Definition 2) with a certain number of non-zero entries.

Lemma 7. *Given* $\mathcal{E} = \{\mathcal{E}_j\}_{j=1}^s$*, the number of vectors with* w *non-zero entries that are consistent with* $\mathcal{E}$ *is*

$$\eta(w) = \sum_{\substack{\boldsymbol{u}:\sum_{j=1}^{s} u_j = w,\\ u_j \leq |\mathcal{E}_j|}} \prod_{j=1}^{s} \binom{|\mathcal{E}_j|}{u_j} \left(2^j-1\right)^{u_j}. \tag{11}$$

If $q=2$ and all the partial-erasure sets are $\{0,1\}$ (i.e., BEC full-erasures), $\eta(w)$ becomes $\binom{|\mathcal{E}|}{w}$, the number of binary vectors of length $|\mathcal{E}|$ with Hamming weight w. Let us denote by $P^{\mathrm{ML}}(\mathcal{G})$ the probability of ML decoding failure for a certain Tanner graph $\mathcal{G}$ from the regular (d_v, d_c) ensemble. As in [9], we use polynomial characteristic functions to identify graph configurations leading to failure events, but now enumerating different configurations. We denote by $\mathrm{coef}\left(f(x), x^i\right)$ the ith coefficient f_i of x^i in the polynomial $f(x) = \sum_{i\geq 0} f_i x^i$. $\mathbb{E}_{\mathrm{LDPC}(d_v,d_c)}\left[P^{\mathrm{ML}}(\mathcal{G})\right]$ denotes the expected probability of ML decoding failure, over LDPC code graphs in the (d_v, d_c) ensemble. Recall that $\eta(w)$ is a function of $\{|\mathcal{E}_j|\}_{j=1}^s$.

Theorem 8.

$$\mathbb{E}_{\mathrm{LDPC}(d_v,d_c)}\left[P^{\mathrm{ML}}(\mathcal{G})\right] \leq \sum_{\substack{|\mathcal{E}_0|,|\mathcal{E}_1|,\ldots,|\mathcal{E}_s|:\\ \sum_{j=0}^{s}|\mathcal{E}_j|=n}} \frac{n!}{|\mathcal{E}_0|!\,|\mathcal{E}_1|!\ldots|\mathcal{E}_s|!} \prod_{j=0}^{s} \varepsilon_j^{|\mathcal{E}_j|} \cdot \min\left\{1, \sum_{w=1}^{|\mathcal{E}|} \eta(w) \frac{\mathrm{coef}\left(\left((1+y)^{d_c} - 1 - y d_c\right)^{n\frac{d_v}{d_c}}, y^{w d_v}\right)}{\binom{n d_v}{w d_v}} \cdot \left(\frac{1}{q-1}\right)^{w\frac{d_v}{d_c}}\right\}. \tag{12}$$

Proof. An ML decoder fails if and only if there is a non-trivial solution to the equation $\mathbf{H}_{\mathcal{E}}\boldsymbol{x}_{\mathcal{E}}^T = \mathbf{0}$, which is consistent with respect to $\{\mathcal{E}_j\}_{j=1}^{s}$:

$$\Pr\left(\exists \boldsymbol{x}_{\mathcal{E}} \neq \mathbf{0}, \boldsymbol{x}_{\mathcal{E}} \text{ is consistent} : \mathbf{H}_{\mathcal{E}}\boldsymbol{x}_{\mathcal{E}}^T = \mathbf{0}\right) \leq \sum_{\boldsymbol{x}_{\mathcal{E}} \neq \mathbf{0}, \boldsymbol{x}_{\mathcal{E}} \text{ is consistent}} \Pr\left(\mathbf{H}_{\mathcal{E}}\boldsymbol{x}_{\mathcal{E}}^T = \mathbf{0}\right), \tag{13}$$

where the upper bound follows by the union bound. Consider an arbitrary but fixed consistent vector $\boldsymbol{x}_{\mathcal{E}}$ and denote the number of its non-zero entries by $w(\boldsymbol{x}_{\mathcal{E}})$. There are $w(\boldsymbol{x}_{\mathcal{E}})d_v$ edges connected to variable nodes corresponding to the non-zero elements of $\boldsymbol{x}_{\mathcal{E}}$. For $\mathbf{H}_{\mathcal{E}}\boldsymbol{x}_{\mathcal{E}}^T = \mathbf{0}$ to hold, each neighbouring check of the $w(\boldsymbol{x}_{\mathcal{E}})$ non-zero variable nodes must be connected to these variable nodes at least twice. As the total number of check nodes is nd_v/d_c, there are $\text{coef}\left(\left((1+y)^{d_c} - 1 - d_c y\right)^{n\frac{d_v}{d_c}}, y^{w(\boldsymbol{x}_{\mathcal{E}})d_v}\right)$ such configurations out of the possible $\binom{nd_v}{w(\boldsymbol{x}_{\mathcal{E}})d_v}$ configurations of edges connected to $w(\boldsymbol{x}_{\mathcal{E}})$ partially-erased variable nodes. According to Lemma 6, the probability that a certain check node is satisfied is upper bounded by $1/(q-1)$, assuming uniform edge labels. The number of check nodes connected to $w(\boldsymbol{x}_{\mathcal{E}})$ variable nodes is at least $w(\boldsymbol{x}_{\mathcal{E}})d_v/d_c$. Thus, $(1/(q-1))^{w(\boldsymbol{x}_{\mathcal{E}})d_v/d_c}$ is an upper bound on the probability that all check nodes connected to the $w(\boldsymbol{x}_{\mathcal{E}})$ non-zero variable nodes are satisfied, such that $\mathbf{H}_{\mathcal{E}}\boldsymbol{x}_{\mathcal{E}}^T = \mathbf{0}$ holds. Finally, by summing over all the possible weights of consistent vectors (counted by $\eta(w)$ of Lemma 7) and taking into account the channel partial-erasure probabilities, (12) is obtained. The minimum in (12) is taken to tighten the upper bound. □

V. Simulation Results

To evaluate the performance of Algorithm 1, we used a regular $(2, 18)$ LDPC code (rate $8/9$), whose rate is of interest in practical flash memories. Two codeword lengths were considered: $n = 504$ and $n = 1008$. The average decoding performance is measured by symbol erasure rate (SER), where each variable node that remains partially erased when the decoder terminates contributes to this quantity.

A comparison of the decoding performance results is shown in Figure 3. The results show a considerable gain thanks to Algorithm 1, up to more than one order of magnitude in SER performance. Note that the performance gap increases with q for a fixed partial-erasure set. This is expected, as the number of resolvable edge labels increases with q (see Theorem 2). In Figure 3b, we compare a GF(8) code to a binary code, where each GF(8) symbol is considered as 3 bits. It is demonstrated that a GF(8) code with Algorithm 1 provides significantly better results.

VI. Conclusion

This work offers a study of the finite-length performance of iterative decoding of GF(q) LDPC codes over the QMBC. We showed that unlike the binary case, partially-erased stopping sets can be resolved by a wise setting of edge labels, and proposed an edge-labeling algorithm for improved finite-length decoding performance. For future work it will be interesting to devise an algorithm that optimizes the graph and edge labels jointly.

VII. Acknowledgement

This work was supported in part by the Israel Science Foundation and by a gift from Qualcomm.

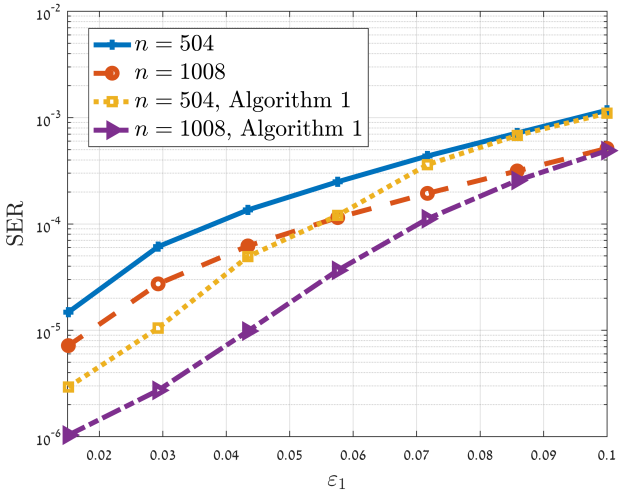

(a) $q = 4$, $j_{\max} = 1$ (decoding threshold 0.117).

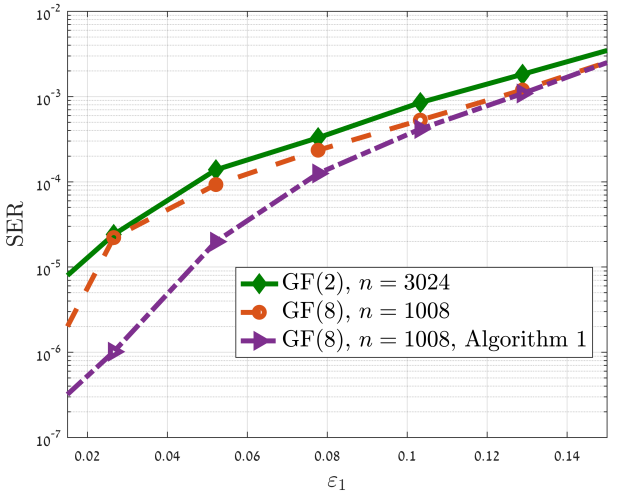

(b) $q = 8$, $j_{\max} = 1$ (decoding threshold 0.176).

Fig. 3: A comparison of symbol-erasure rate (SER) performance for GF(q) codes, between uniformly-distributed edge labels and labels optimized using Algorithm 1. The decoding thresholds are given for optimal edge-label distributions [1].

References

[1] R. Cohen and Y. Cassuto, "LDPC codes for the q-ary bit-measurement channel," *9th International Symposium on Turbo Codes & Iterative Information Processing*, pp. 261–265, September 2016.

[2] ——, "Iterative decoding of LDPC codes over the q-ary partial erasure channel," *IEEE Transactions on Information Theory*, vol. 62, no. 5, May 2016.

[3] C. Di, D. Proietti, I. Telatar, T. Richardson, and R. Urbanke, "Finite-length analysis of low-density parity-check codes on the binary erasure channel," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1570–1579, Jun 2002.

[4] A. Bazarsky, N. Presman, and S. Litsyn, "Design of non-binary quasi-cyclic LDPC codes by ACE optimization," in *2013 IEEE Information Theory Workshop (ITW)*, Sept 2013, pp. 1–5.

[5] B. Amiri, J. Kliewer, and L. Dolecek, "Analysis and enumeration of absorbing sets for non-binary graph-based codes," *IEEE Transactions on Communications*, vol. 62, no. 2, pp. 398–409, February 2014.

[6] M. Davey and D. MacKay, "Low-density parity check codes over GF(q)," *IEEE Communications Letters*, vol. 2, no. 6, pp. 165–167, June 1998.

[7] C. Poulliat, M. Fossorier, and D. Declercq, "Design of regular $(2, d_c)$-LDPC codes over GF(q) using their binary images," *IEEE Transactions on Communications*, vol. 56, no. 10, pp. 1626–1635, October 2008.

[8] A. McGregor and O. Milenkovic, "On the hardness of approximating stopping and trapping sets," *IEEE Transactions on Information Theory*, vol. 56, no. 4, pp. 1640–1650, April 2010.

[9] A. Orlitsky, K. Viswanathan, and J. Zhang, "Stopping set distribution of LDPC code ensembles," *IEEE Transactions on Information Theory*, vol. 51, no. 3, pp. 929–953, March 2005.