

Multi-Erasure Locally Recoverable Codes over Small Fields: A Tensor Product Approach

Pengfei Huang, *Member, IEEE*, Eitan Yaakobi, *Senior Member, IEEE*,
and Paul H. Siegel, *Life Fellow, IEEE*

Abstract—Erasure codes play an important role in storage systems to prevent data loss. In this work, we study a class of erasure codes called Multi-Erasure Locally Recoverable Codes (ME-LRCs) for storage arrays. Compared to previous related works, we focus on the construction of ME-LRCs over small fields. Our main contribution is a general construction of ME-LRCs based on generalized tensor product codes, and an analysis of their erasure-correcting properties. A decoding algorithm tailored for erasure recovery is given, and correctable erasure patterns are identified. We then prove that our construction yields optimal ME-LRCs with a wide range of code parameters, and present some explicit ME-LRCs over small fields. Next, we show that generalized integrated interleaving (GII) codes can be treated as a subclass of generalized tensor product codes, thus defining the exact relation between these codes. Finally, ME-LRCs are investigated in a probabilistic setting. We prove that ME-LRCs based upon a generalized tensor product construction can achieve the capacity of a compound erasure channel consisting of a family of erasure product channels.

Index Terms—Locally recoverable codes, small fields, tensor product codes, capacity-achieving, compound channel.

I. INTRODUCTION

Recently, erasure codes with both local and global erasure-correcting properties have received considerable attention [5], [13], [25]–[27], [29], thanks to their promising application in storage systems. The idea behind them is that when only a few erasures occur, these erasures can be corrected fast using only local parities. If the number of erasures exceeds

This work was supported in part by Western Digital Corporation, in part by NSF Grants CCF-1405119 and CCF-1619053, and in part by BSF Grant 2015816. E. Yaakobi was supported in part by the Center for Memory and Recording Research (CMRR) at the University of California San Diego. This paper was presented at the 55th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, USA, Oct. 2017 [16].

P. Huang was with the Center for Memory and Recording Research, Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093, USA. He is now with Western Digital Corporation, Milpitas, CA 95035, USA (e-mail: pehuangucsd@gmail.com).

P. H. Siegel is with the Center for Memory and Recording Research, Department of Electrical and Computer Engineering, University of California San Diego, La Jolla, CA 92093, USA (e-mail: psiegel@ucsd.edu).

E. Yaakobi is with the Department of Computer Science, Technion—Israel Institute of Technology, Haifa 32000, Israel (e-mail: yaakobi@cs.technion.ac.il). This work was done in part while E. Yaakobi was visiting the Center for Memory and Recording Research, University of California San Diego.

Copyright (c) 2017 IEEE. Personal use of this material is permitted. However, permission to use this material for any other purposes must be obtained from the IEEE by sending a request to pubs-permissions@ieee.org.

the local erasure-correcting capability, then the global parities are invoked.

In this paper, we consider erasure codes with both local and global erasure-correcting capabilities for a $\rho \times n_0$ storage array [5], where each row contains some local parities, and additional global parities are distributed in the array. The array structure is suitable for many storage applications. For example, a storage array can represent a large-scale distributed storage system consisting of a large number of storage nodes spread over different geographical locations. The storage nodes that are placed in the same location can form a local storage cluster. Thus, each row of the storage array can represent such a local storage cluster. Another example is a redundant array of independent disks (RAID) type of architecture for solid-state drives (SSDs) [5], [12]. In this scenario, a $\rho \times n_0$ storage array can represent a total of ρ SSDs, each of which contains n_0 flash memory chips. Within each SSD, an erasure code is applied to these n_0 chips for local protection. In addition, erasure coding is also done across all the SSDs for global protection of all the chips. More specifically, let us give the formal definition of this class of erasure codes as follows.

Definition 1. Consider a code \mathcal{C} over a finite field \mathbb{F}_q consisting of $\rho \times n_0$ arrays such that:

- 1) Each row in each array in \mathcal{C} belongs to a linear local code \mathcal{C}_0 with length n_0 and minimum distance d_0 over \mathbb{F}_q .
- 2) Reading the symbols of \mathcal{C} row-wise, \mathcal{C} is a linear code with length ρn_0 , dimension k , and minimum distance d over \mathbb{F}_q .

Then, we say that \mathcal{C} is a $(\rho, n_0, k; d_0, d)_q$ Multi-Erasure Locally Recoverable Code (**ME-LRC**).

Thus, a $(\rho, n_0, k; d_0, d)_q$ ME-LRC can locally correct $d_0 - 1$ erasures in each row, and is guaranteed to correct a total of $d - 1$ erasures anywhere in the array.

Our work is motivated by a recent work by Blaum and Hetzler [5]. In their work, the authors studied ME-LRCs where each row is a maximum distance separable (MDS) code, and gave code constructions with field size $q \geq \max\{\rho, n_0\}$ using generalized integrated interleaving (GII) codes [15], [32], [35]. Our Definition 1 follows from and generalizes the definition of the codes in [5] by not requiring each row to be an MDS code. There exist other related works. The ME-LRCs in Definition 1 can be seen as (r, δ) LRCs with disjoint repair sets. A code \mathcal{C} is called

an (r, δ) LRC [27] if for every coordinate there exists a punctured code (i.e., a repair set) of \mathcal{C} , with support containing this coordinate, whose length is at most $r + \delta - 1$, and whose minimum distance is at least δ . Although the existing constructions [27], [29] for (r, δ) LRCs with disjoint repair sets can generate ME-LRCs as in Definition 1, they use MDS codes as local codes and require a field size that is at least as large as the code length. A recent work [3] gives explicit constructions of (r, δ) LRCs over field \mathbb{F}_q derived from algebraic curves. These codes have disjoint repair sets with size $r + \delta - 1 = \sqrt{q}$ or $r + \delta - 1 = \sqrt{q} + 1$. Partial MDS (PMDS) codes [4] are also related to but different from ME-LRCs in Definition 1. In general, PMDS codes need to satisfy stricter requirements than ME-LRCs. A $\rho \times n_0$ array code is called an $(r; s)$ PMDS code if each row is an $[n_0, n_0 - r, r + 1]_q$ MDS code and whenever any r locations in each row are punctured, the resulting code is also an MDS code with minimum distance $s + 1$. A construction of $(r; s)$ PMDS codes for all r and s with field size $O(n_0^{\rho n_0})$ was given in [8]. More recently, a family of PMDS codes with field size $O(\max\{\rho, n_0^{r+s}\}^s)$ was presented in [10].

To the best of our knowledge, however, the construction of *optimal* ME-LRCs over any small field (e.g., over a field size less than the length of the local code, such as the binary field) has not been fully explored and solved. The goal of this paper is to study ME-LRCs over small fields. We propose a general construction based on generalized tensor product codes [20], [34], which were first utilized in [19] to construct binary single-erasure LRCs that had been considered in [13], [14], [17], [18], [26], [29], [31]. More specifically, the contributions of this paper are as follows:

1) We extend our previous construction in [19] to the scenario of multi-erasure LRCs over any field. As a result, the construction in [19] can be seen as a special case of the construction presented in this paper. In contrast to [5], our construction does not require field size $q \geq \max\{\rho, n_0\}$, and it can even generate binary ME-LRCs. We derive an upper bound on the minimum distance of ME-LRCs. For $2d_0 \geq d$, we show that our construction can produce ME-LRCs that are optimal with respect to (w.r.t.) the upper bound on the minimum distance.

2) We present an erasure decoding algorithm and its corresponding correctable erasure patterns, which include the pattern of any $d - 1$ erasures. We show that the ME-LRCs from our construction based on Reed-Solomon (RS) codes are optimal w.r.t. certain correctable erasure patterns.

3) So far the *exact* relation between GII codes [5], [32], [35] and generalized tensor product codes has not been fully investigated. We prove that GII codes are a subclass of generalized tensor product codes. As a result, the parameters of a GII code can be obtained by using the known properties of generalized tensor product codes.

4) We present a new interpretation of ME-LRCs from an information-theoretic perspective. Thanks to the locality property of ME-LRCs, it is quite natural to speculate that a $(\rho, n_0, k; d_0, d)_q$ ME-LRC might be suitable for an erasure product channel that has ρ parallel local erasure channels.

Indeed, we show that the generalized tensor product structure with some appropriate component codes (e.g., Reed-Muller codes and Bose-Chaudhuri-Hocquenghem (BCH) codes) can be used to obtain a sequence of ME-LRCs that achieve the capacity of a compound erasure channel which comprises a family of erasure product channels.

The remainder of the paper is organized as follows. In Section II, We introduce notation used in the paper and present bounds on the minimum distance of ME-LRCs. In Section III, we present a general construction of ME-LRCs. The erasure-correcting properties of these codes are studied and an erasure decoding algorithm is described. In Section IV, we present an optimal code construction and give several explicit optimal ME-LRCs over different fields. In Section V, we prove that GII codes are a subclass of generalized tensor product codes. In Section VI, we study capacity-achieving ME-LRCs for a compound erasure channel. Section VII concludes the paper.

II. PRELIMINARIES

In this section, we first introduce some notation that will be used throughout this paper, and then we derive field-size dependent bounds on the minimum distance of ME-LRCs. The upper bound obtained here will be used to prove the optimality of our construction for ME-LRCs in the following sections.

We use the notation $[n]$ to denote the set $\{1, \dots, n\}$. For a length- n vector v over \mathbb{F}_q and a set $\mathcal{I} \subseteq [n]$, the vector $v_{\mathcal{I}}$ denotes the restriction of the vector v to the coordinates in the set \mathcal{I} , and $w_q(v)$ represents the Hamming weight of the vector v over \mathbb{F}_q . The transpose of a matrix H is written as H^T . For a set \mathcal{S} , $|\mathcal{S}|$ represents the cardinality of the set. A linear code \mathcal{C} over \mathbb{F}_q of length n , dimension k , and minimum distance d will be denoted by $\mathcal{C} = [n, k, d]_q$. For a code with only one codeword, the minimum distance is defined as ∞ .

Now, we give an upper bound on the minimum distance of a $(\rho, n_0, k; d_0, d)_q$ ME-LRC by extending the shortening bound for LRCs in [7]. Bounds for other generalizations of LRCs can be found in [1], [3], [30].

Let $d_{opt}^{(q)}[n, k]$ denote the largest possible minimum distance of a linear code of length n and dimension k over \mathbb{F}_q , and let $k_{opt}^{(q)}[n, d]$ denote the largest possible dimension of a linear code of length n and minimum distance d over \mathbb{F}_q . Note that for large enough field size q , from the Singleton bound, $d_{opt}^{(q)}[n, k] = n - k + 1$ and $k_{opt}^{(q)}[n, d] = n - d + 1$.

Lemma 2. For any $(\rho, n_0, k; d_0, d)_q$ ME-LRC \mathcal{C} , the minimum distance d satisfies

$$d \leq \min_{0 \leq x \leq \lceil \frac{k}{k^*} \rceil - 1, x \in \mathbb{Z}} \left\{ d_{opt}^{(q)}[\rho n_0 - x n_0, k - x k^*] \right\}, \quad (1)$$

and the dimension satisfies

$$k \leq \min_{0 \leq x \leq \lceil \frac{k}{k^*} \rceil - 1, x \in \mathbb{Z}} \left\{ x k^* + k_{opt}^{(q)}[\rho n_0 - x n_0, d] \right\}, \quad (2)$$

where $k^* = k_{opt}^{(q)}[n_0, d_0]$.

Proof: The proof is based on the shortening argument used in [7]. See Appendix A. ■

An asymptotic lower bound for ME-LRCs with local MDS codes was given in [3]. Here, by simply adapting the Gilbert-Varshamov (GV) bound [28], we derive the following GV-like lower bound on ME-LRCs of finite length without specifying local codes.

Lemma 3. *A $(\rho, n_0, k; \geq d_0, \geq d)_q$ ME-LRC \mathcal{C} exists if*

$$\sum_{i=0}^{d-2} \binom{\rho(n_0 - \lceil \log_q(\sum_{j=0}^{d_0-2} \binom{n_0-1}{j} (q-1)^j) \rceil) - 1}{i} (q-1)^i < q^{\rho(n_0 - \lceil \log_q(\sum_{j=0}^{d_0-2} \binom{n_0-1}{j} (q-1)^j) \rceil) - k}. \quad (3)$$

Proof: See Appendix B. ■

III. ME-LRCs FROM GENERALIZED TENSOR PRODUCT CODES: CONSTRUCTION AND DECODING

Tensor product codes, first proposed by Wolf in [34], are a family of binary error-correcting codes defined by a parity-check matrix that is the tensor product of the parity-check matrices of two constituent codes. Later, they were generalized in [20]. In this section, we first introduce generalized tensor product codes over a finite field \mathbb{F}_q . Then, we give a general construction of ME-LRCs from generalized tensor product codes. We determine the minimum distance of the constructed ME-LRCs, describe a decoding algorithm tailored for erasure correction, and study the corresponding correctable erasure patterns.

A. Generalized Tensor Product Codes over \mathbb{F}_q

We start by presenting the tensor product operation of two matrices H' and H'' . Let H' be the parity-check matrix of a code with length n' and dimension $n' - v$ over \mathbb{F}_q . The matrix H' can be considered as a $v \times n'$ matrix over \mathbb{F}_q or as a $1 \times n'$ matrix of elements from \mathbb{F}_q^v . Let H' be the vector $H' = [h'_1 \ h'_2 \ \dots \ h'_{n'}]$, where h'_j , $1 \leq j \leq n'$, are elements of \mathbb{F}_q^v . Let H'' be the parity-check matrix of a code of length ℓ and dimension $\ell - \lambda$ over \mathbb{F}_q^v . We denote H'' by

$$H'' = \begin{bmatrix} h''_{11} & \dots & h''_{1\ell} \\ \vdots & \ddots & \vdots \\ h''_{\lambda 1} & \dots & h''_{\lambda\ell} \end{bmatrix},$$

where h''_{ij} , $1 \leq i \leq \lambda$ and $1 \leq j \leq \ell$, are elements of \mathbb{F}_q^v .

The tensor product of the matrices H'' and H' is defined as

$$H_{TP} = H'' \otimes H' = \begin{bmatrix} h''_{11}H' & \dots & h''_{1\ell}H' \\ \vdots & \ddots & \vdots \\ h''_{\lambda 1}H' & \dots & h''_{\lambda\ell}H' \end{bmatrix},$$

where $h''_{ij}H' = [h''_{ij}h'_1 \ h''_{ij}h'_2 \ \dots \ h''_{ij}h'_{n'}]$, $1 \leq i \leq \lambda$ and $1 \leq j \leq \ell$, and the products of elements are calculated

according to the rules of multiplication for elements over \mathbb{F}_q^v . The matrix H_{TP} will be considered as a $v\lambda \times n'\ell$ matrix of elements from \mathbb{F}_q , thus defining a tensor product code over \mathbb{F}_q .

Lemma 4. *The rank of the matrix H_{TP} over \mathbb{F}_q is $v\lambda$.*

Proof: Without loss of generality, assume that the first λ columns of H'' are linearly independent. Thus, we can transform H'' into the form:

$$\hat{H}'' = \begin{bmatrix} 1 & 0 & \dots & 0 & \hat{h}''_{1,\lambda+1} & \dots & \hat{h}''_{1\ell} \\ 0 & 1 & \dots & 0 & \hat{h}''_{2,\lambda+1} & \dots & \hat{h}''_{2\ell} \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & \hat{h}''_{\lambda,\lambda+1} & \dots & \hat{h}''_{\lambda\ell} \end{bmatrix},$$

where the first λ columns form the identity matrix. Then, by elementary row operations, the matrix $H_{TP} = H'' \otimes H'$ can be transformed into the form:

$$\hat{H}_{TP} = \begin{bmatrix} H' & 0 & \dots & 0 & \hat{h}''_{1,\lambda+1}H' & \dots & \hat{h}''_{1\ell}H' \\ 0 & H' & \dots & 0 & \hat{h}''_{2,\lambda+1}H' & \dots & \hat{h}''_{2\ell}H' \\ \vdots & \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & H' & \hat{h}''_{\lambda,\lambda+1}H' & \dots & \hat{h}''_{\lambda\ell}H' \end{bmatrix}.$$

Since the left part of \hat{H}_{TP} is a block diagonal matrix, the rank of \hat{H}_{TP} is $v\lambda$. The matrices H_{TP} and \hat{H}_{TP} have the same rank, so the rank of H_{TP} is $v\lambda$. ■

We provide an example to illustrate the tensor product operations described above.

Example 1. (cf. [34]) Let α be a primitive element of \mathbb{F}_4 . Let H'' be the following parity-check matrix over \mathbb{F}_4 for a $[5, 3, 3]_4$ code,

$$H'' = \begin{bmatrix} \alpha^0 & 0 & \alpha^0 & \alpha^0 & \alpha^0 \\ 0 & \alpha^0 & \alpha^0 & \alpha^1 & \alpha^2 \end{bmatrix}.$$

Let H' be the following parity-check matrix over \mathbb{F}_2 for a $[3, 1, 3]_2$ Hamming code,

$$H' = \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}.$$

Representing the elements of \mathbb{F}_4 as $\alpha^0 = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, $\alpha^1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$, $\alpha^2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, and $0 = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$, we have

$$H_{TP} = H'' \otimes H' = \begin{bmatrix} \alpha^0 \alpha^1 \alpha^2 & 0 & 0 & 0 & \alpha^0 \alpha^1 \alpha^2 & \alpha^0 \alpha^1 \alpha^2 & \alpha^0 \alpha^1 \alpha^2 \\ 0 & 0 & 0 & \alpha^0 \alpha^1 \alpha^2 & \alpha^0 \alpha^1 \alpha^2 & \alpha^1 \alpha^2 \alpha^0 & \alpha^2 \alpha^0 \alpha^1 \end{bmatrix}.$$

Using the same symbol-to-binary vector mapping, we represent H_{TP} over \mathbb{F}_2 as

$$H_{TP} = \begin{bmatrix} 1 & 0 & 1 & 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 \end{bmatrix},$$

which defines a binary $[15, 11, 3]_2$ code. □

Our construction of ME-LRCs is based on generalized tensor product codes [20]. Define the matrices H'_i and H''_i for $i = 1, 2, \dots, \mu$ as follows. The matrix H'_i is a $v_i \times n'$ matrix over \mathbb{F}_q such that the $(v_1 + v_2 + \dots + v_i) \times n'$ matrix

$$B_i = \begin{bmatrix} H'_1 \\ H'_2 \\ \vdots \\ H'_i \end{bmatrix} \quad (4)$$

is a parity-check matrix of an $[n', n' - v_1 - v_2 - \dots - v_i, d'_i]_q$ code C'_i , where $d'_1 \leq d'_2 \leq \dots \leq d'_i$. The matrix H''_i is a $\lambda_i \times \ell$ matrix over $\mathbb{F}_{q^{v_i}}$, which is a parity-check matrix of an $[\ell, \ell - \lambda_i, \delta_i]_{q^{v_i}}$ code C''_i .

We define a μ -level generalized tensor product code over \mathbb{F}_q as a linear code having a parity-check matrix over \mathbb{F}_q in the form of the following μ -level tensor product structure

$$H = \begin{bmatrix} H''_1 \otimes H'_1 \\ H''_2 \otimes H'_2 \\ \vdots \\ H''_\mu \otimes H'_\mu \end{bmatrix}. \quad (5)$$

As with the matrix H_{TP} , each level in the matrix H is obtained by tensor product operations over the extension field $\mathbb{F}_{q^{v_i}}$. In general, the tensor product computational complexity varies from one level to another, and the computational complexity is higher for the level with a larger field size q^{v_i} . As a result, a smaller v_i is preferred in practice. The rank of H over \mathbb{F}_q is $\sum_{i=1}^{\mu} v_i \lambda_i$, with level i contributing $v_i \lambda_i$. This can be shown similarly by using the proof technique of Lemma 4 for each level of H . We denote this code by C_{GTP}^μ . Its length is $n_t = n' \ell$ and the dimension is $k_t = n_t - \sum_{i=1}^{\mu} v_i \lambda_i$.

By adapting Theorem 2 in [20] from the field \mathbb{F}_2 to \mathbb{F}_q , we directly obtain the following theorem on the minimum distance of C_{GTP}^μ over \mathbb{F}_q .

Theorem 5. *The minimum distance d_t of a generalized tensor product code C_{GTP}^μ over \mathbb{F}_q satisfies*

$$d_t \geq \min\{\delta_1, \delta_2 d'_1, \delta_3 d'_2, \dots, \delta_\mu d'_{\mu-1}, d'_\mu\}.$$

Proof: The proof is given in Appendix C for completeness. ■

B. Construction of ME-LRCs

Now, we present a general construction of ME-LRCs based on generalized tensor product codes, and also determine the corresponding code parameters.

Construction A

Step 1: Choose $v_i \times n'$ matrices H'_i over \mathbb{F}_q and $\lambda_i \times \ell$ matrices H''_i over $\mathbb{F}_{q^{v_i}}$, for $i = 1, 2, \dots, \mu$, which satisfy the following two properties:

1) The parity-check matrix $H''_1 = I_{\ell \times \ell}$, i.e., the $\ell \times \ell$ identity matrix.

2) The matrices H'_i (equivalently, B_i), $1 \leq i \leq \mu$, and H''_j , $2 \leq j \leq \mu$, are chosen such that $d'_\mu \leq \delta_j d'_{j-1}$, for $j = 2, 3, \dots, \mu$.

Step 2: Generate a parity-check matrix H over \mathbb{F}_q according to (5) with the matrices H'_i and H''_i , for $i = 1, 2, \dots, \mu$. The constructed code corresponding to the parity-check matrix H is referred to as C_A .

Note that in Construction A, the parity-check matrix H''_1 is chosen to be the identity matrix to endow the code C_A with local erasure-correcting capability.

Theorem 6. *The code C_A is a $(\rho, n_0, k; d_0, d)_q$ ME-LRC with parameters $\rho = \ell$, $n_0 = n'$, $k = n' \ell - \sum_{i=1}^{\mu} v_i \lambda_i$, $d_0 = d'_1$, and $d = d'_\mu$.*

Proof: The code parameters ρ , n_0 , k , and d_0 can be easily determined directly from Construction A. We now prove that $d = d'_\mu$.

Since $\delta_1 = \infty$ (H''_1 is the identity matrix) and $d'_\mu \leq \delta_i d'_{i-1}$ for all $i = 2, 3, \dots, \mu$, Theorem 5 implies $d \geq d'_\mu$.

Now, we show that $d \leq d'_\mu$. For $i = 1, 2, \dots, \mu$, let $H'_i = [h'_1(i), \dots, h'_{n'}(i)]$ over $\mathbb{F}_{q^{v_i}}$, and let $[h''_{11}(i), \dots, h''_{\lambda_1}(i)]^T$ over $\mathbb{F}_{q^{v_i}}$ be the first column of H''_i . Since the code with the parity-check matrix B_μ has minimum distance d'_μ , there exist d'_μ columns of B_μ , say in the set of positions $J = \{b_1, b_2, \dots, b_{d'_\mu}\}$, which are linearly dependent; that is, $\sum_{j \in J} \alpha_j h'_j(i) = 0$, for some $\alpha_j \in \mathbb{F}_q$, for all $i = 1, 2, \dots, \mu$. Thus, we have $\sum_{j \in J} \alpha_j h''_{p1}(i) h'_j(i) = h''_{p1}(i) \left(\sum_{j \in J} \alpha_j h'_j(i) \right) = 0$, for $p = 1, 2, \dots, \lambda_i$ and $i = 1, 2, \dots, \mu$. That is, the columns in positions $b_1, b_2, \dots, b_{d'_\mu}$ of H are linearly dependent. ■

The following example illustrates Construction A.

Example 2. Let $H'_1 = [1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1]$ over \mathbb{F}_2 , and

$$H'_2 = \begin{bmatrix} 0 & 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 & 1 \end{bmatrix}$$

over \mathbb{F}_2 . The corresponding code C'_1 is the $[7, 6, 2]_2$ single parity-check code, and C'_2 is the even subcode of the $[7, 4, 3]_2$ Hamming code. We also choose

$$H''_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

over \mathbb{F}_2 and $H''_2 = [1 \ 1 \ 1]$ over \mathbb{F}_8 . Hence, in this construction, we use the following parameters: $n' = 7$, $\ell = 3$, $v_1 = 1$, $v_2 = 3$, $\lambda_1 = 3$, $\lambda_2 = 1$, $d'_1 = 2$, $d'_2 = 4$,

and $\delta_2 = 2$. The binary parity-check matrix H from (5) is

$$H = \begin{bmatrix} H_1'' \otimes H_1' \\ H_2'' \otimes H_2' \\ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \\ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \\ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \ 1 \\ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1 \ 0 \ 0 \ 0 \ 1 \ 1 \ 1 \ 1 \\ 0 \ 1 \ 1 \ 0 \ 0 \ 1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 0 \ 1 \ 1 \ 0 \ 1 \ 1 \ 0 \ 0 \ 1 \ 1 \\ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \ 1 \ 0 \ 1 \ 0 \ 1 \ 0 \ 1 \end{bmatrix}.$$

According to Construction A, the constructed binary code \mathcal{C}_A corresponding to the parity-check matrix H is a $(3, 7, 15; 2, 4)_2$ ME-LRC. \square

The multi-level structure of Construction A endows fine granularity and provides flexible code parameters. The following example shows that the three-level construction is more flexible than the two-level one.

Example 3. First, let us consider the *three-level construction* (i.e., $\mu = 3$). Consider a chain of nested binary extended BCH codes: $\mathcal{C}_3' = [32, 16, 8]_2 \subset \mathcal{C}_2' = [32, 26, 4]_2 \subset \mathcal{C}_1' = [32, 31, 2]_2$. Choose H_1'' to be the 5×5 identity matrix. To satisfy the condition in Construction A that $d_3'' \leq \delta_j d_{j-1}''$, for $j = 2, 3$, we have $\delta_2 \geq 4$ and $\delta_3 \geq 2$. Choose H_2'' to be a parity-check matrix of a $[5, 2, 4]_{25}$ code and H_3'' to be a parity-check matrix of a $[5, 4, 2]_{210}$ code. From Construction A, the resulting code \mathcal{C}_A is a $(\rho = 5, n_0 = 32, k = 130; d_0 = 2, d = 8)_2$ ME-LRC.

Now, for the *two-level construction* (i.e., $\mu = 2$), consider two nested binary extended BCH codes: $\mathcal{C}_2' = [32, 16, 8]_2 \subset \mathcal{C}_1' = [32, 31, 2]_2$. Note that here \mathcal{C}_2' is the same as \mathcal{C}_3' in the three-level construction above. Choose H_1'' to be the 5×5 identity matrix. To satisfy the condition in Construction A that $d_2'' \leq \delta_2 d_1''$, we have $\delta_2 \geq 4$. Choose H_2'' to be a parity-check matrix of a $[5, 2, 4]_{215}$ code. The resulting code \mathcal{C}_A from Construction A is a $(\rho = 5, n_0 = 32, k = 110; d_0 = 2, d = 8)_2$ ME-LRC.

The codes from the three-level construction and the two-level construction have the same code parameters except that the dimension of the latter one is smaller. However, in the three-level construction, if we replace H_3'' with a parity-check matrix of a $[5, 2, 4]_{210}$ code, then the resulting code \mathcal{C}_A becomes a $(\rho = 5, n_0 = 32, k = 110; d_0 = 2, d = 8)_2$ ME-LRC which has the same code parameters as the one obtained from the two-level construction. \square

C. Erasure Decoding and Correctable Erasure Patterns

We present a decoding algorithm for the ME-LRC \mathcal{C}_A obtained from Construction A, tailored for erasure correction.

Let the symbol ? represent an erasure and “e” denote a decoding failure. The erasure decoder $\mathcal{D}_A : (\mathbb{F}_q \cup \{?\})^{n'\ell} \rightarrow \mathcal{C}_A \cup \{“e”\}$ for an ME-LRC \mathcal{C}_A consists of two kinds of component decoders \mathcal{D}_i' and \mathcal{D}_i'' for $i = 1, 2, \dots, \mu$ described below.

1) The decoder for a coset of the code \mathcal{C}_i' with parity-check matrix B_i , $i = 1, 2, \dots, \mu$, is denoted by

$$\mathcal{D}_i' : (\mathbb{F}_q \cup \{?\})^{n'} \times (\mathbb{F}_q \cup \{?\})^{\sum_{j=1}^i v_j} \rightarrow (\mathbb{F}_q \cup \{?\})^{n'}.$$

It uses the following decoding rule: for a length- n' input vector \mathbf{y}' , and a length- $\sum_{j=1}^i v_j$ syndrome vector \mathbf{s}' without erasures, if \mathbf{y}' agrees with exactly one codeword $\mathbf{c}' \in \mathcal{C}_i' + \mathbf{e}$ on the non-erased entries with values in \mathbb{F}_q , where the vector \mathbf{e} is a coset leader determined by both the code \mathcal{C}_i' and the syndrome vector \mathbf{s}' , i.e., $\mathbf{s}' = \mathbf{e}B_i^T$, then $\mathcal{D}_i'(\mathbf{y}', \mathbf{s}') = \mathbf{c}'$; otherwise, $\mathcal{D}_i'(\mathbf{y}', \mathbf{s}') = \mathbf{y}'$. Therefore, if the length- n' input vector \mathbf{y}' is a codeword in $\mathcal{C}_i' + \mathbf{e}$ with no more than $d_i' - 1$ erasures and the syndrome vector \mathbf{s}' contains no erasures, then the decoder \mathcal{D}_i' can return the correct codeword.

2) The decoder for the code \mathcal{C}_i'' with parity-check matrix H_i'' , $i = 1, 2, \dots, \mu$, is denoted by

$$\mathcal{D}_i'' : (\mathbb{F}_{q^{v_i}} \cup \{?\})^\ell \rightarrow (\mathbb{F}_{q^{v_i}} \cup \{?\})^\ell.$$

It uses the following decoding rule: for a length- ℓ input vector \mathbf{y}'' , if \mathbf{y}'' agrees with exactly one codeword $\mathbf{c}'' \in \mathcal{C}_i''$ on the non-erased entries with values in $\mathbb{F}_{q^{v_i}}$, then $\mathcal{D}_i''(\mathbf{y}'') = \mathbf{c}''$; otherwise, $\mathcal{D}_i''(\mathbf{y}'') = \mathbf{y}''$. Therefore, if the length- ℓ input vector \mathbf{y}'' is a codeword in \mathcal{C}_i'' with no more than $\delta_i - 1$ erasures, then the decoder \mathcal{D}_i'' can successfully return the correct codeword.

Note that the decoders \mathcal{D}_i' and \mathcal{D}_i'' introduced above are maximum-likelihood (ML) decoders.

The erasure decoder \mathcal{D}_A for the code \mathcal{C}_A is summarized in Algorithm 1 below. Let the input word of length $n'\ell$ for the decoder \mathcal{D}_A be $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_\ell)$, where each component $\mathbf{y}_i \in (\mathbb{F}_q \cup \{?\})^{n'}$, $i = 1, \dots, \ell$. The vector \mathbf{y} is an erased version of a codeword $\mathbf{c} = (c_1, c_2, \dots, c_\ell) \in \mathcal{C}_A$.

Algorithm 1: Decoding Procedure of Decoder \mathcal{D}_A

Input: received word $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_\ell)$.

Output: codeword $\mathbf{c} \in \mathcal{C}_A$ or a decoding failure “e”.

1. Let $\mathbf{s}_j^1 = \mathbf{0}$, for $j = 1, 2, \dots, \ell$.
 2. $\hat{\mathbf{c}} = (\hat{c}_1, \dots, \hat{c}_\ell) = (\mathcal{D}_1'(\mathbf{y}_1, \mathbf{s}_1^1), \dots, \mathcal{D}_1'(\mathbf{y}_\ell, \mathbf{s}_\ell^1))$.
 3. Let $\mathcal{F} = \{j \in [\ell] : \hat{c}_j \text{ contains ?}\}$.
 4. **for** $i = 2, \dots, \mu$
 - If $\mathcal{F} \neq \emptyset$, do the following steps; otherwise go to step 5.
 - $(\mathbf{s}_1^i, \dots, \mathbf{s}_\ell^i) = \mathcal{D}_i''(\hat{c}_1 H_i^{i,T}, \dots, \hat{c}_\ell H_i^{i,T})$.
 - $\hat{c}_j = \mathcal{D}_i'(\hat{c}_j, (\mathbf{s}_j^1, \dots, \mathbf{s}_j^i))$ for $j \in \mathcal{F}$; \hat{c}_j remains the same for $j \in [\ell] \setminus \mathcal{F}$.
 - Update $\mathcal{F} = \{j \in [\ell] : \hat{c}_j \text{ contains ?}\}$.
 - end**
 5. If $\mathcal{F} = \emptyset$, let $\mathbf{c} = \hat{\mathbf{c}}$ and output \mathbf{c} ; otherwise return “e”.
-

In Algorithm 1, we use the following rules for the operations which involve the symbol ?:

- 1) Addition +: for any element $\gamma \in \mathbb{F}_q \cup \{?\}$, $\gamma + ? = ?$.
- 2) Multiplication \times : for any element $\gamma \in \mathbb{F}_q \cup \{?\} \setminus \{0\}$, $\gamma \times ? = ?$, and $0 \times ? = 0$.

3) If a length- n vector \mathbf{x} , $\mathbf{x} \in (\mathbb{F}_q \cup \{?\})^n$, contains an entry $?$, then \mathbf{x} is considered as the symbol $?$ in the set $\mathbb{F}_q^n \cup \{?\}$. Similarly, the symbol $?$ in the set $\mathbb{F}_q^n \cup \{?\}$ is treated as a length- n vector whose entries are all $?$.

In Algorithm 1, lines 1 to 3 correspond to correcting erasures locally, while the “for loop” at line 4 (i.e., when $i \geq 2$) corresponds to global erasure correction. We refer to the i th loop in the “for loop” as the i th level of decoding.

To describe correctable erasure patterns, we use the following notation. Let $w_e(\mathbf{v})$ denote the number of erasures $?$ in the vector \mathbf{v} . For a received word $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_\ell)$, let $N_\tau = |\{\mathbf{y}_m : w_e(\mathbf{y}_m) \geq d'_\tau, 1 \leq m \leq \ell\}|$ for $1 \leq \tau \leq \mu$.

Theorem 7. *The decoder \mathcal{D}_A for the $(\rho, n_0, k; d_0, d)_q$ ME-LRC \mathcal{C}_A can correct any received word \mathbf{y} that satisfies the following condition:*

$$N_\tau \leq \delta_{\tau+1} - 1, \forall 1 \leq \tau \leq \mu, \quad (6)$$

where $\delta_{\mu+1}$ is defined to be 1.

Proof: See Appendix D. ■

The following corollary follows from Theorem 7.

Corollary 8. *The decoder \mathcal{D}_A for the $(\rho, n_0, k; d_0, d)_q$ ME-LRC \mathcal{C}_A can correct any received word \mathbf{y} with less than d erasures.*

Proof: See Appendix E. ■

We use the following example to illustrate Algorithm 1.

Example 4. Consider the binary $(\rho = 3, n_0 = 7, k = 15; d_0 = 2, d = 4)_2$ ME-LRC \mathcal{C}_A constructed in Example 2. It consists of three sub-blocks, each of length 7. It can locally correct 1 erasure in each sub-block, and is guaranteed to correct 3 erasures globally.

Moreover, some erasure patterns with more than 3 (i.e., $d - 1$) erasures can be corrected with Algorithm 1. For instance, consider the codeword \mathbf{c} of \mathcal{C}_A : $\mathbf{c} = (1 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0, 0 \ 0 \ 0 \ 0 \ 1 \ 1 \ 0, 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0)$. The erased word $\mathbf{y} = (\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3) = (1 \ ? \ 0 \ 0 \ 0 \ 0 \ 0, ? \ 0 \ ? \ 0 \ ? \ 1 \ 0, 0 \ 0 \ ? \ 0 \ 0 \ 0 \ 0)$ with 5 erasures can be decoded by Algorithm 1. More specifically, in Algorithm 1, in step 1, we have $\mathbf{s}_j^1 = \mathbf{0}$ for all $j = 1, 2, 3$. In step 2, we locally correct \mathbf{y}_1 and \mathbf{y}_3 , but \mathbf{y}_2 cannot be recovered and is left for global correction. Thus, $\mathcal{F} = \{2\}$ in step 3. In step 4, we first obtain $\mathbf{s}_2^2 = (0 \ 1 \ 1)$ due to $\mathbf{s}_1^2 = (0 \ 1 \ 1)$, $\mathbf{s}_3^2 = (0 \ 0 \ 0)$, and $\mathbf{s}_1^2 + \mathbf{s}_2^2 + \mathbf{s}_3^2 = \mathbf{0}$. Then, using $\mathbf{s}_1^2 = \mathbf{0}$ and $\mathbf{s}_2^2 = (0 \ 1 \ 1)$, we can correct \mathbf{y}_2 which has 3 erasures. □

IV. OPTIMAL CONSTRUCTION AND EXPLICIT ME-LRCS OVER SMALL FIELDS

In this section, we study the optimality of Construction A, and also present several explicit ME-LRCS that are optimal over different fields.

A. Optimal Construction

We show how to construct ME-LRCS which are optimal w.r.t. the bound (1) by adding more constraints to Construction A. To this end, we specify the choice of the matrices in Construction A. This specification, referred to as **Design I**, is as follows.

Design I: Matrix Specifications

- 1) H_1' is the parity-check matrix of an $[n', n' - v_1, d_1']_q$ code which satisfies $k_{opt}^{(q)}[n', d_1'] = n' - v_1$.
- 2) B_μ is the parity-check matrix of an $[n', n' - \sum_{i=1}^\mu v_i, d'_\mu]_q$ code with $d_{opt}^{(q)}[n', n' - \sum_{i=1}^\mu v_i] = d'_\mu$.
- 3) The minimum distances satisfy $d'_\mu \leq 2d'_1$.
- 4) H_i'' is an all-one vector of length ℓ over $\mathbb{F}_{q^{v_i}}$, i.e., the parity-check matrix of a linear code with minimum distance $\delta_i = 2$, for all $i = 2, \dots, \mu$.

Theorem 9. *The code \mathcal{C}_A from Construction A with Design I is a $(\rho = \ell, n_0 = n', k = n'\ell - v_1\ell - \sum_{i=2}^\mu v_i; d_0 = d'_1, d = d'_\mu)_q$ ME-LRC, which is optimal with respect to the bound (1).*

Proof: From Theorem 6, the code parameters ρ, n_0, k, d_0 , and d can be determined. We have $k^* = k_{opt}^{(q)}[n', d'_1] = n' - v_1$. Setting $x = \ell - 1$, we get

$$\begin{aligned} d &\leq \min_{0 \leq x \leq \lfloor \frac{k}{k^*} \rfloor - 1} \left\{ d_{opt}^{(q)}[\rho n_0 - x n_0, k - x k^*] \right\} \\ &\leq d_{opt}^{(q)}[\ell n' - (\ell - 1)n', k - (\ell - 1)k^*] \\ &= d_{opt}^{(q)}[n', n' - \sum_{i=1}^\mu v_i] = d'_\mu. \end{aligned}$$

This proves that \mathcal{C}_A achieves the bound (1). ■

Theorem 9 shows that by properly specifying the (short) component codes used in Construction A, the optimality of the resulting (long) ME-LRC can be guaranteed.

B. Explicit ME-LRCS from Construction A

Our construction is very flexible and can generate many ME-LRCS over different fields. In the following, we present several examples.

1) ME-LRCS with local extended BCH codes over \mathbb{F}_2

From the structure of BCH codes [28], there exists a chain of nested binary extended BCH codes: $\mathcal{C}_3 = [2^m, 2^m - 1 - 3m, 8]_2 \subset \mathcal{C}_2 = [2^m, 2^m - 1 - 2m, 6]_2 \subset \mathcal{C}_1 = [2^m, 2^m - 1 - m, 4]_2$, for $m \geq 5$.

Let the matrices B_1, B_2 , and B_3 be the parity-check matrices of $\mathcal{C}_1, \mathcal{C}_2$, and \mathcal{C}_3 , respectively.

Example 5. For $\mu = 3$, in Construction A, we use the above matrices B_1, B_2 , and B_3 . We also choose H_2'' and H_3'' to be the all-one vector of length ℓ over \mathbb{F}_{2^m} .

From Theorem 6, the corresponding $(\ell, n_0, k; d_0, d)_2$ ME-LRC \mathcal{C}_A has parameters $n_0 = 2^m, k = 2^m\ell - (m + 1)\ell -$

$2m$, $d_0 = 4$, and $d = 8$. Since the extended Hamming code \mathcal{C}_1 has an optimal dimension and the extended triple-error-correcting BCH code \mathcal{C}_3 has an optimal minimum distance, this code satisfies the requirements of Design I. Thus, from Theorem 9, it is optimal w.r.t. the bound (1).

The code \mathcal{C}_A has ℓ sub-blocks. Any sub-block with less than 4 erasures can be corrected locally, and any 7 (i.e., $d - 1$) erasures are guaranteed to be recovered. Moreover, any erasure pattern is correctable if it satisfies: 1) the number of sub-blocks that have more than 3 erasures is less than 2, and 2) no sub-block has more than 7 erasures. For instance, consider the erasure pattern where only the first three sub-blocks have erasures, and the first, second, and third sub-blocks have 3, 3, and 6 erasures, respectively. This erasure pattern has a total of 12 erasures. Although it has more than 7 (i.e., $d - 1$) erasures, it is still correctable. In addition, for some erasure patterns, the decoding can be completed at the second level in Algorithm 1. One such erasure pattern is where only the first three sub-blocks have erasures and the erasure numbers are 2, 2, and 5, respectively.

The advantage of using the three-level (or similarly multi-level) instead of the two-level construction is that the three-level structure endows fine granularity; that is, it increases the erasure-correcting capability as well as the decoding complexity gradually. Thanks to this property, some erasure patterns as presented above can be corrected during the second level of decoding in Algorithm 1 and thus the decoding is terminated earlier (i.e., skipping the third level of decoding), resulting in a smaller decoding latency. \square

2) ME-LRCs with local algebraic geometry codes over \mathbb{F}_4

Algebraic geometry codes usually have large minimum distance and often possess a nested structure [33]. We use a class of algebraic geometry codes called Hermitian codes [36] to construct ME-LRCs.

From the construction of Hermitian codes [36], there exists a chain of nested 4-ary Hermitian codes: $\mathcal{C}_H(1) = [8, 1, 8]_4 \subset \mathcal{C}_H(2) = [8, 2, 6]_4 \subset \mathcal{C}_H(3) = [8, 3, 5]_4 \subset \mathcal{C}_H(4) = [8, 4, 4]_4 \subset \mathcal{C}_H(5) = [8, 5, 3]_4 \subset \mathcal{C}_H(6) = [8, 6, 2]_4 \subset \mathcal{C}_H(7) = [8, 7, 2]_4$.

Now, let the matrices B_1, B_2, B_3 , and B_4 be the parity-check matrices of $\mathcal{C}_H(4), \mathcal{C}_H(3), \mathcal{C}_H(2)$, and $\mathcal{C}_H(1)$, respectively. Let H_i'' , $i = 2, 3, 4$, be the all-one vector of length ℓ over \mathbb{F}_4 .

Example 6. For $\mu = 2$, in Construction A, we use the above matrices B_1, B_2 , and H_2'' . From Theorem 6, the corresponding $(\ell, n_0, k; d_0, d)_4$ ME-LRC \mathcal{C}_A has parameters $n_0 = 8$, $k = 4\ell - 1$, $d_0 = 4$, and $d = 5$.

For $\mu = 3$, in Construction A, we use the above matrices B_1, B_2, B_3, H_2'' , and H_3'' . From Theorem 6, the corresponding $(\ell, n_0, k; d_0, d)_4$ ME-LRC \mathcal{C}_A has parameters $n_0 = 8$, $k = 4\ell - 2$, $d_0 = 4$, and $d = 6$.

For $\mu = 4$, in Construction A, we use the above matrices $B_i, i = 1, \dots, 4$, and H_j'' , $j = 2, 3, 4$. From Theorem 6, the corresponding $(\ell, n_0, k; d_0, d)_4$ ME-LRC \mathcal{C}_A has parameters $n_0 = 8$, $k = 4\ell - 3$, $d_0 = 4$, and $d = 8$.

All of the above three families of ME-LRCs over \mathbb{F}_4 are optimal w.r.t. the bound (1). \square

3) ME-LRCs with local singly-extended Reed-Solomon codes over \mathbb{F}_q

Let $n' \leq q$ and α be a primitive element of \mathbb{F}_q . We choose H_1' to be the parity-check matrix of an $[n', n' - d_1' + 1, d_1']_q$ singly-extended RS code, namely

$$H_1' = \begin{bmatrix} 1 & 1 & \cdots & 1 & 1 \\ 1 & \alpha & \cdots & \alpha^{n'-2} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \alpha^{d_1'-2} & \cdots & \alpha^{(n'-2)(d_1'-2)} & 0 \end{bmatrix}.$$

For $i = 2, 3, \dots, \mu$, we choose H_i' to be

$$H_i' = \begin{bmatrix} 1 & \alpha^{d_{i-1}'-1} & \cdots & \alpha^{(n'-2)(d_{i-1}'-1)} & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 1 & \alpha^{d_i'-2} & \cdots & \alpha^{(n'-2)(d_i'-2)} & 0 \end{bmatrix},$$

where $d_1' < d_2' < \cdots < d_\mu'$. We also require that

$$\delta_i = \lceil \frac{d_\mu'}{d_{i-1}'} \rceil = \lceil \frac{d_\mu'}{d_{i-1}' + 1} \rceil = \cdots = \lceil \frac{d_\mu'}{d_i' - 1} \rceil, \forall i = 2, \dots, \mu$$

and $\delta_2 > \delta_3 > \cdots > \delta_\mu$.

For $i = 2, 3, \dots, \mu$, let H_i'' be the parity-check matrix of an $[\ell, \ell - \delta_i + 1, \delta_i = \lceil \frac{d_\mu'}{d_{i-1}'} \rceil]_{q^{v_i}}$ MDS code, which exists whenever $\ell \leq q^{v_i}$, where $v_i = d_i' - d_{i-1}'$. Note that for an MDS code with minimum distance 2, the code length can be arbitrarily long.

Example 7. We use the above chosen matrices H_i' and H_i'' for Construction A. The corresponding $(\ell, n', k; d_0, d)_q$ ME-LRC \mathcal{C}_A has parameters $k = (n' - d_1' + 1)\ell - \sum_{i=2}^\mu (\lceil \frac{d_\mu'}{d_{i-1}'} \rceil - 1)(d_i' - d_{i-1}')$, $d_0 = d_1'$, and $d = d_\mu'$; the field size q satisfies $q \geq \max\{q', n'\}$, where $q' = \max_{i=2, \dots, \mu} \{ \lceil \frac{d_\mu'}{d_{i-1}' - 1} \rceil \}$.

When $\mu = 2$ and $d_1' < d_2' \leq 2d_1'$, the corresponding $(\ell, n', k; d_0, d)_q$ ME-LRC \mathcal{C}_A has parameters $k = (n' - d_1' + 1)\ell - (d_2' - d_1')$, $d_0 = d_1'$, and $d = d_2'$; the field size q needs to satisfy $q \geq n'$. Since \mathcal{C}_A satisfies the requirements of Design I, from Theorem 9, it is optimal w.r.t. the bound (1). \square

The following theorem shows that the μ -level ME-LRC \mathcal{C}_A constructed in Example 7 is optimal in the sense of possessing the largest possible dimension among all codes with the same erasure-correcting capability.

Theorem 10. Let \mathcal{C} be a code of length $\ell n'$ and dimension k over \mathbb{F}_q . Each codeword in \mathcal{C} consists of ℓ sub-blocks, each of length n' . Assume that \mathcal{C} corrects all erasure patterns satisfying the condition in (6), where $\delta_\tau = \lceil \frac{d_\mu'}{d_{\tau-1}'} \rceil$ for $2 \leq \tau \leq \mu$. Then, the dimension satisfies $k \leq (n' - d_1' + 1)\ell - \sum_{i=2}^\mu (\lceil \frac{d_\mu'}{d_{i-1}'} \rceil - 1)(d_i' - d_{i-1}')$.

Proof: The proof is by contradiction.

Let each codeword in \mathcal{C} correspond to an $\ell \times n'$ array. We index the coordinates of the array from 1 to $\ell n'$, proceeding from left to right within each row, and taking the rows from top to bottom. Let \mathcal{I}_1 be the set of coordinates defined by $\mathcal{I}_1 = \{(i-1)n' + j : \delta_2 - 1 < i \leq \ell, 1 \leq j \leq d'_1 - 1\}$. For $2 \leq \tau \leq \mu$, let \mathcal{I}_τ be the set of coordinates given by $\mathcal{I}_\tau = \{(i-1)n' + j : \delta_{\tau+1} - 1 < i \leq \delta_\tau - 1, 1 \leq j \leq d'_\tau - 1\}$, where $\delta_{\mu+1}$ is defined to be 1. Let \mathcal{I} be the set of all the coordinates of the array.

By calculation, we have $|\mathcal{I} \setminus (\mathcal{I}_1 \cup \mathcal{I}_2 \cup \dots \cup \mathcal{I}_\mu)| = (n' - d'_1 + 1)\ell - \sum_{i=2}^\mu (\lceil \frac{d'_\mu}{d'_{i-1}} \rceil - 1)(d'_i - d'_{i-1})$. Now, assume that $k > (n' - d'_1 + 1)\ell - \sum_{i=2}^\mu (\lceil \frac{d'_\mu}{d'_{i-1}} \rceil - 1)(d'_i - d'_{i-1})$. Then, there exist at least two distinct codewords c' and c'' in \mathcal{C} that agree on the coordinates in the set $\mathcal{I} \setminus (\mathcal{I}_1 \cup \mathcal{I}_2 \cup \dots \cup \mathcal{I}_\mu)$. We erase the values on the coordinates in the set $\mathcal{I}_1 \cup \mathcal{I}_2 \cup \dots \cup \mathcal{I}_\mu$ of both c' and c'' . This erasure pattern satisfies the condition in (6). Since c' and c'' are distinct, this erasure pattern is uncorrectable. Thus, our assumption that $k > (n' - d'_1 + 1)\ell - \sum_{i=2}^\mu (\lceil \frac{d'_\mu}{d'_{i-1}} \rceil - 1)(d'_i - d'_{i-1})$ is violated. ■

Remark 1. The construction by Blaum and Hertzler [5] based on GII codes cannot generate the ME-LRCs constructed in Examples 5 and 6. For the ME-LRC in Example 7, since the local code is the singly-extended RS code, the construction in [5] can also be used to produce an ME-LRC that has the same code parameters ρ, n_0, k, d_0 and d as those of the ME-LRC \mathcal{C}_A from our construction. However, the construction in [5] requires the field size q to satisfy $q \geq \max\{\ell, n'\}$, which in general is larger than that in our construction.

V. RELATION TO GENERALIZED INTEGRATED INTERLEAVING CODES

Integrated interleaving (II) codes were first introduced in [15] as a two-level error-correcting scheme for data storage applications, and were then extended in [32] and more recently in [35] as generalized integrated interleaving (GII) codes for multi-level data protection. In [5], [37], GII codes were utilized for local erasure recovery.

The main difference between GII codes and generalized tensor product codes is that a generalized tensor product code over \mathbb{F}_q is defined by operations over the base field \mathbb{F}_q and its extension fields, as shown in (5); in contrast, a GII code over \mathbb{F}_q is defined by operations only over the field \mathbb{F}_q . As a result, generalized tensor product codes are more flexible than GII codes, and generally GII codes cannot be used to construct ME-LRCs over very small fields, e.g., the binary field.

The goal of this section is to study the exact relation between generalized tensor product codes and GII codes. We will show that GII codes are in fact a subclass of generalized tensor product codes. The idea is to reformulate the parity-check matrix of a GII code into the form of a parity-check matrix of a generalized tensor product code. Establishing this relation allows some code properties of GII codes to be obtained directly from known results about generalized

tensor product codes. We start by considering II codes, the two-level case of GII codes, to illustrate our idea.

A. Integrated Interleaving Codes

We take our definition of II codes from [15]. Let $\mathcal{C}_i, i = 1, 2$, be $[n, k_i, d_i]_q$ linear codes over \mathbb{F}_q such that $\mathcal{C}_2 \subset \mathcal{C}_1$ and $d_2 > d_1$. An II code \mathcal{C}_{II} is defined as follows:

$$\mathcal{C}_{II} = \left\{ c = (c_0, c_1, \dots, c_{m-1}) : c_i \in \mathcal{C}_1, 0 \leq i < m, \right. \\ \left. \text{and } \sum_{i=0}^{m-1} \alpha^{bi} c_i \in \mathcal{C}_2, b = 0, 1, \dots, \gamma - 1 \right\}, \quad (7)$$

where α is a primitive element of \mathbb{F}_q and $\gamma < m \leq q - 1$.

According to the above definition, it is known that the parity-check matrix of \mathcal{C}_{II} is

$$H_{II} = \begin{bmatrix} I & \otimes & H_1 \\ \Gamma_2 & \otimes & H_2 \end{bmatrix}, \quad (8)$$

where \otimes denotes the Kronecker product. The matrices H_1 and $\begin{bmatrix} H_1 \\ H_2 \end{bmatrix}$ over \mathbb{F}_q are the parity-check matrices of \mathcal{C}_1 and \mathcal{C}_2 , respectively; the matrix I over \mathbb{F}_q is the $m \times m$ identity matrix; and the matrix Γ_2 over \mathbb{F}_q is the parity-check matrix of an $[m, m - \gamma, \gamma + 1]_q$ code with the following form

$$\Gamma_2 = \begin{bmatrix} 1 & 1 & \dots & 1 \\ 1 & \alpha & \dots & \alpha^{m-1} \\ 1 & \alpha^2 & \dots & \alpha^{2(m-1)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^{(\gamma-1)} & \dots & \alpha^{(\gamma-1)(m-1)} \end{bmatrix}. \quad (9)$$

Remark 2. The parity-check matrix H_{II} over \mathbb{F}_q in (8) of \mathcal{C}_{II} is obtained by operations over the same field \mathbb{F}_q . In contrast, the parity-check matrix H over \mathbb{F}_q in (5) of a generalized tensor product code is obtained by operations over both the base field \mathbb{F}_q and its extension fields.

Remark 3. In general, the codes \mathcal{C}_1 and \mathcal{C}_2 in (7) are chosen to be RS codes [15]. If \mathcal{C}_1 and \mathcal{C}_2 are chosen to be binary codes, then m can only be $m = 1$.

To see the relation between II codes and generalized tensor product codes, we reformulate H_{II} in (8) by *splitting* the rows of H_2 , producing the following form,

$$H_{II} = \begin{bmatrix} I & \otimes & H_1 \\ \Gamma_2 & \otimes & H_2(1) \\ \Gamma_2 & \otimes & H_2(2) \\ \vdots & \vdots & \vdots \\ \Gamma_2 & \otimes & H_2(k_1 - k_2) \end{bmatrix}. \quad (10)$$

Here, the matrix H_1 over \mathbb{F}_q is the parity-check matrix of \mathcal{C}_1 and is treated as a vector over the extension field $\mathbb{F}_{q^{n-k_1}}$; correspondingly, the matrix I is treated as the $m \times m$ identity matrix over $\mathbb{F}_{q^{n-k_1}}$. For $1 \leq i \leq k_1 - k_2$, $H_2(i)$ over \mathbb{F}_q represents the i th row of H_2 , and Γ_2 over \mathbb{F}_q is the matrix in (9).

Now, referring to the matrix in (5), the matrix in (10) can be interpreted as a parity-check matrix of a $(1 + k_1 - k_2)$ -level generalized tensor product code over \mathbb{F}_q . Thus, we conclude that an II code is a generalized tensor product code. Using the properties of generalized tensor product codes, we can directly obtain the following result, which was proved in [15] in an alternative manner.

Lemma 11. *The code \mathcal{C}_{II} is a linear code over \mathbb{F}_q of length $N = nm$, dimension $K = (m - \gamma)k_1 + \gamma k_2$, and minimum distance $D \geq \min\{(\gamma + 1)d_1, d_2\}$.*

Proof: For $1 \leq i \leq k_1 - k_2$, let the following parity-check matrix

$$\begin{bmatrix} H_1 \\ H_2(1) \\ \vdots \\ H_2(i) \end{bmatrix}$$

define an $[n, k_1 - i, d_{2,i}]_q$ code. It is clear that $d_1 \leq d_{2,1} \leq d_{2,2} \leq \dots \leq d_{2,k_1 - k_2} = d_2$.

From the properties of generalized tensor product codes, the redundancy is $N - K = nm - K = (n - k_1)m + \gamma(k_1 - k_2)$; that is, the dimension is $K = k_1(m - \gamma) + k_2\gamma$. Using Theorem 5, the minimum distance is $D \geq \min\{d_1(\gamma + 1), d_{2,1}(\gamma + 1), \dots, d_{2,k_1 - k_2 - 1}(\gamma + 1), d_{2,k_1 - k_2}\} = \min\{(\gamma + 1)d_1, d_2\}$. ■

B. Generalized Integrated Interleaving Codes

We extend the idea used in the previous subsection to the more general case of GII codes. We use the definition of GII codes from [35] for consistency.

Let \mathcal{C}_i , $i = 0, 1, \dots, \gamma$, be $[n, k_i, d_i]_q$ codes over \mathbb{F}_q such that

$$\begin{aligned} \mathcal{C}_{i_s} &= \dots = \mathcal{C}_{i_{s-1}+1} \subset \mathcal{C}_{i_{s-1}} = \dots = \mathcal{C}_{i_{s-2}+1} \\ &\subset \dots \subset \mathcal{C}_{i_1} = \dots = \mathcal{C}_1 \subset \mathcal{C}_0, \end{aligned} \quad (11)$$

where $i_0 = 0$, $i_s = \gamma$, and $i_0 \leq i_1 \leq \dots \leq i_s$. The minimum distances satisfy $d_0 \leq d_1 \leq \dots \leq d_\gamma$. A GII code \mathcal{C}_{GII} is defined as:

$$\begin{aligned} \mathcal{C}_{GII} &= \left\{ \mathbf{c} = (c_0, c_1, \dots, c_{m-1}) : c_i \in \mathcal{C}_0, 0 \leq i < m, \right. \\ &\quad \left. \text{and } \sum_{i=0}^{m-1} \alpha^{bi} c_i \in \mathcal{C}_{\gamma-b}, b = 0, 1, \dots, \gamma - 1 \right\}, \end{aligned} \quad (12)$$

where α is a primitive element of \mathbb{F}_q and $\gamma < m \leq q - 1$.

We will use several matrices in the representation of the parity-check matrix of \mathcal{C}_{GII} . Let the matrix I over \mathbb{F}_q be the $m \times m$ identity matrix. Let H_0 over \mathbb{F}_q be the parity-check matrix of \mathcal{C}_0 . For $1 \leq j \leq s$, let the matrix $\begin{bmatrix} H_0 \\ H_{i_j} \end{bmatrix}$ over \mathbb{F}_q

represent the parity-check matrix of \mathcal{C}_{i_j} , where H_{i_j} has the form

$$H_{i_j} = \begin{bmatrix} H_{i_1 \setminus i_0} \\ H_{i_2 \setminus i_1} \\ \vdots \\ H_{i_j \setminus i_{j-1}} \end{bmatrix}.$$

The size of the submatrix $H_{i_j \setminus i_{j-1}}$ in H_{i_j} is $(k_{i_{j-1}} - k_{i_j}) \times n$. For any $i \leq j$, let the matrix $\Gamma(i, j; \alpha)$ over \mathbb{F}_q be the parity-check matrix of an $[m, m - (j - i + 1), j - i + 2]_q$ code with the following form

$$\Gamma(i, j; \alpha) = \begin{bmatrix} 1 & \alpha^i & \dots & \alpha^{i(m-1)} \\ 1 & \alpha^{i+1} & \dots & \alpha^{(i+1)(m-1)} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & \alpha^j & \dots & \alpha^{j(m-1)} \end{bmatrix}. \quad (13)$$

Now, according to the definition in (12), using the matrices introduced above, the parity-check matrix of \mathcal{C}_{GII} is

$$H_{GII} = \begin{bmatrix} I & \otimes & H_0 \\ \Gamma(0, i_s - i_{s-1} - 1; \alpha) & \otimes & H_{i_s} \\ \Gamma(i_s - i_{s-1}, i_s - i_{s-2} - 1; \alpha) & \otimes & H_{i_{s-1}} \\ \vdots & \vdots & \vdots \\ \Gamma(i_s - i_2, i_s - i_1 - 1; \alpha) & \otimes & H_{i_2} \\ \Gamma(i_s - i_1, i_s - i_0 - 1; \alpha) & \otimes & H_{i_1} \end{bmatrix}, \quad (14)$$

which, after rearranging the rows, can be simplified into

$$H_{GII} = \begin{bmatrix} I & \otimes & H_0 \\ \Gamma(0, i_s - i_0 - 1; \alpha) & \otimes & H_{i_1 \setminus i_0} \\ \Gamma(0, i_s - i_1 - 1; \alpha) & \otimes & H_{i_2 \setminus i_1} \\ \vdots & \vdots & \vdots \\ \Gamma(0, i_s - i_{s-2} - 1; \alpha) & \otimes & H_{i_{s-1} \setminus i_{s-2}} \\ \Gamma(0, i_s - i_{s-1} - 1; \alpha) & \otimes & H_{i_s \setminus i_{s-1}} \end{bmatrix}. \quad (15)$$

Note that since the matrix of (15) is obtained from the matrix of (14) by a permutation of rows, they define the same code. For the sake of convenience, we use the same notation H_{GII} for both of them by a slight abuse of notation.

To make a connection between GII codes and generalized tensor product codes, we further reformulate the matrix H_{GII}

in (15) as follows,

$$H_{GII} = \left[\begin{array}{c|c} I & \otimes H_0 \\ \hline \Gamma(0, i_s - i_0 - 1; \alpha) & \otimes H_{i_1 \setminus i_0}(1) \\ \vdots & \vdots \\ \Gamma(0, i_s - i_0 - 1; \alpha) & \otimes H_{i_1 \setminus i_0}(k_{i_0} - k_{i_1}) \\ \hline \Gamma(0, i_s - i_1 - 1; \alpha) & \otimes H_{i_2 \setminus i_1}(1) \\ \vdots & \vdots \\ \Gamma(0, i_s - i_1 - 1; \alpha) & \otimes H_{i_2 \setminus i_1}(k_{i_1} - k_{i_2}) \\ \hline \vdots & \vdots \\ \Gamma(0, i_s - i_{s-2} - 1; \alpha) & \otimes H_{i_{s-1} \setminus i_{s-2}}(1) \\ \vdots & \vdots \\ \Gamma(0, i_s - i_{s-2} - 1; \alpha) & \otimes H_{i_{s-1} \setminus i_{s-2}}(k_{i_{s-2}} - k_{i_{s-1}}) \\ \hline \Gamma(0, i_s - i_{s-1} - 1; \alpha) & \otimes H_{i_s \setminus i_{s-1}}(1) \\ \vdots & \vdots \\ \Gamma(0, i_s - i_{s-1} - 1; \alpha) & \otimes H_{i_s \setminus i_{s-1}}(k_{i_{s-1}} - k_{i_s}) \end{array} \right] \quad (16)$$

where, in the first level, the matrix H_0 over \mathbb{F}_q is treated as a vector over the extension field $\mathbb{F}_{q^{n-k_0}}$, and correspondingly the matrix I is treated as the $m \times m$ identity matrix over $\mathbb{F}_{q^{n-k_0}}$. For $1 \leq x \leq s$ and $1 \leq y \leq k_{i_{x-1}} - k_{i_x}$, $H_{i_x \setminus i_{x-1}}(y)$ over \mathbb{F}_q represents the y th row of the matrix $H_{i_x \setminus i_{x-1}}$.

Now, referring to the matrix in (5), the matrix in (16) can be seen as a parity-check matrix of a $(1 + k_0 - k_{i_s})$ -level generalized tensor product code over \mathbb{F}_q . As a result, we can directly obtain the following lemma, which was also proved in [35] in a different way.

Lemma 12. *The code \mathcal{C}_{GII} is a linear code over \mathbb{F}_q of length $N = nm$, dimension $K = \sum_{x=1}^s k_x + (m - \gamma)k_0 = \sum_{j=1}^s (i_j - i_{j-1})k_j + (m - \gamma)k_0$, and minimum distance $D \geq \min \{(\gamma + 1)d_0, (\gamma - i_1 + 1)d_{i_1}, \dots, (\gamma - i_{s-1} + 1)d_{i_{s-1}}, d_{i_s}\}$.*

Proof: For $1 \leq x \leq s$ and $1 \leq y \leq k_{i_{x-1}} - k_{i_x}$, let the following parity-check matrix

$$\left[\begin{array}{c} H_0 \\ \hline H_{i_1 \setminus i_0}(1) \\ \vdots \\ H_{i_1 \setminus i_0}(k_{i_0} - k_{i_1}) \\ \hline \vdots \\ H_{i_x \setminus i_{x-1}}(1) \\ \vdots \\ H_{i_x \setminus i_{x-1}}(y) \end{array} \right]$$

define an $[n, k_{i_{x-1}} - y, d_{i_x, y}]_q$ code, so we have $d_{i_{x-1}} \leq d_{i_x, 1} \leq d_{i_x, 2} \leq \dots \leq d_{i_x, (k_{i_{x-1}} - k_{i_x})} = d_{i_x}$. From the properties of generalized tensor product codes, it is easy to obtain the dimension $K = \sum_{j=1}^s (i_j - i_{j-1})k_j + (m - \gamma)k_0$.

From Theorem 5, the minimum distance satisfies

$$\begin{aligned} D &\geq \min \left\{ (\gamma + 1)d_0, (\gamma + 1)d_{i_1, 1}, \right. \\ &\quad \dots, (\gamma + 1)d_{i_1, k_{i_0} - k_{i_1} - 1}, (\gamma - i_1 + 1)d_{i_1}, \\ &\quad \dots, (\gamma - i_{s-1} + 1)d_{i_{s-1}}, (\gamma - i_{s-1} + 1)d_{i_s, 1}, \\ &\quad \left. \dots, (\gamma - i_{s-1} + 1)d_{i_s, k_{i_{s-1}} - k_{i_s} - 1}, d_{i_s} \right\} \\ &= \min \left\{ (\gamma + 1)d_0, (\gamma - i_1 + 1)d_{i_1}, \right. \\ &\quad \left. \dots, (\gamma - i_{s-1} + 1)d_{i_{s-1}}, d_{i_s} \right\}. \end{aligned}$$

■

Remark 4. In some prior works, generalized tensor product codes are called generalized error-location (GEL) codes [6], [24]. Recently, in [35], the similarity between GII codes and GEL codes was observed. However, the exact relation between them was not studied. In [35], the author also proposed a new generalized integrated interleaving scheme over binary BCH codes, called GII-BCH codes. These codes can also be seen as a special case of generalized tensor product codes.

Remark 5. Construction A for generalized tensor product codes is also related to the well-known $|\mathbf{u}|\mathbf{u} + \mathbf{v}|$ construction [23, Ch. 2.9] which is defined as follows. Given an $[n, k_1, d_1]_q$ code \mathcal{C}_1 and an $[n, k_2, d_2]_q$ code \mathcal{C}_2 , we can form a new code \mathcal{C}_3 consisting of all vectors: $|\mathbf{u}|\mathbf{u} + \mathbf{v}|$, $\mathbf{u} \in \mathcal{C}_1$ and $\mathbf{v} \in \mathcal{C}_2$. Then \mathcal{C}_3 is a $[2n, k_1 + k_2, \min\{2d_1, d_2\}]_q$ code.

If we assume that \mathcal{C}_2 is a subcode of \mathcal{C}_1 , then Construction A corresponds to the $|\mathbf{u}|\mathbf{u} + \mathbf{v}|$ construction. More specifically, let \mathcal{C}_1 and \mathcal{C}_2 have parity-check matrices H'_1 and $\begin{bmatrix} H'_1 \\ H'_2 \end{bmatrix}$ respectively. Choose $H''_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ and $H''_2 = [-1 \ 1]$. Then, Construction A generates the matrix in (5) as

$$H = \begin{bmatrix} H'_1 & 0 \\ 0 & H'_1 \\ -H'_2 & H'_2 \end{bmatrix},$$

which is a parity-check matrix of the code \mathcal{C}_3 obtained from the $|\mathbf{u}|\mathbf{u} + \mathbf{v}|$ construction.

VI. CAPACITY-ACHIEVING ME-LRCs FOR A COMPOUND ERASURE PRODUCT CHANNEL

In this section, we turn to a probabilistic setting and interpret ME-LRCs from an information-theoretic perspective. Specifically, we construct ME-LRCs that are universally good for a family of erasure product channels defined as a compound erasure channel, i.e., that achieve the compound capacity. Since we will not make explicit reference to minimum distances d_0 and d in the following discussion of ME-LRCs, we will simplify the notation $(\rho, n_0, k; d_0, d)$ to $(\rho, n_0, k)_q$ when referring to ME-LRC code parameters.

A. Information-Theoretic Motivation

Consider the memoryless q -ary erasure channel (QEC) $W: \mathcal{X} \rightarrow \mathcal{Y}$, with input alphabet \mathcal{X} , output alphabet \mathcal{Y} , and transition probabilities $W(y|x)$, $x \in \mathcal{X}$, $y \in \mathcal{Y}$. The input alphabet \mathcal{X} is \mathbb{F}_q , and the output alphabet \mathcal{Y} is $\mathbb{F}_q \cup \{?\}$ (of size $q+1$), where $?$ represents an erasure symbol. For every pair consisting of a transmitted symbol $x \in \mathbb{F}_q$ and a received symbol $y \in \mathbb{F}_q \cup \{?\}$, the transition probability $W(y|x)$ is:

$$W(y|x) = \begin{cases} 1 - \varepsilon & \text{if } y = x \\ \varepsilon & \text{if } y = ? \\ 0 & \text{otherwise,} \end{cases}$$

where ε is called the erasure probability. The capacity of this QEC W is denoted by $C(W)$ and is attained by the uniform input distribution [28], i.e.,

$$\begin{aligned} C(W) &= \max_{p(x)} I(X; Y) \\ &= \max_{p(x)} \sum_{x \in \mathcal{X}} \sum_{y \in \mathcal{Y}} \frac{1}{q} W(y|x) \log_q \frac{W(y|x)}{\sum_{x' \in \mathcal{X}} \frac{1}{q} W(y|x')} \\ &= 1 - \varepsilon. \end{aligned}$$

Note that the base of the logarithm is q . For a linear code $\mathcal{C} = [n, k, d]_q$ over a QEC W , let $P_e^{(n)}(x)$ denote the conditional block probability of error, assuming that x was sent, $x \in \mathcal{C}$. Let $P_e^{(n)}(\mathcal{C})$ denote the average probability of error of this code. Assuming equiprobable codewords, it is clear that, by symmetry,

$$P_e^{(n)}(\mathcal{C}) = \frac{1}{|\mathcal{C}|} \sum_{x \in \mathcal{C}} P_e^{(n)}(x) = P_e^{(n)}(x).$$

The communication scenario in which the ME-LRCs will be applied is the following. Consider a channel consisting of a parallel bank of ρ independent *local* QECs. The erasure probabilities $\varepsilon_1, \dots, \varepsilon_\rho$ of the local QECs are not precisely known at the encoder, but it is known that they are a permutation of a specific vector of erasure probabilities $\boldsymbol{\varepsilon}$. Given a $(\rho, n_0, k)_q$ ME-LRC, each of the ρ sub-blocks (i.e., local codewords) is transmitted over a corresponding local QEC. Roughly speaking, if the local channel has a small enough erasure probability, then it can be decoded locally. If, on the other hand, the local channel has too large an erasure probability, the local code needs to resort to some global parities to help decoding. In the remainder of this section, we formalize this scenario and present a construction of ME-LRCs that achieve the compound capacity of the set of parallel banks of channels corresponding to a given erasure probability vector $\boldsymbol{\varepsilon}$.

B. Erasure Product Channel and Compound Channel

We now formally define an erasure product channel $W_{pd}(\mathbf{y}_\ell | \mathbf{x}_\ell; \sigma)$ consisting of ℓ parallel QECs W_1, W_2, \dots, W_ℓ in a certain fixed order which is determined by a permutation σ . Without loss of generality, we assume that ℓ is some fixed value, and write $\ell = \sum_{i=1}^\mu \ell_i$, for some $1 \leq \mu \leq \ell$. The first ℓ_1 QECs W_1, \dots, W_{ℓ_1}

are the same, with erasure probability ε_1 ; similarly, for $2 \leq i \leq \mu$, the QECs $W_{\sum_{j=1}^{i-1} \ell_j + 1}, \dots, W_{\sum_{j=1}^i \ell_j}$ are the same, with erasure probability ε_i . We also assume that $\varepsilon_1 < \varepsilon_2 < \dots < \varepsilon_\mu$, so the capacities of these QECs satisfy $C(W_{\ell_1}) > C(W_{\ell_2}) > \dots > C(W_{\ell_\mu})$.

Consider a permutation given by a bijective mapping $\sigma: [\ell] \rightarrow [\ell]$. The erasure product channel is defined as:

$$W_{pd}(\mathbf{y}_\ell | \mathbf{x}_\ell; \sigma): \mathcal{X}_1 \times \dots \times \mathcal{X}_\ell \rightarrow \mathcal{Y}_1 \times \dots \times \mathcal{Y}_\ell$$

with each input alphabet $\mathcal{X}_i = \mathbb{F}_q$ and each output alphabet $\mathcal{Y}_i = \mathbb{F}_q \cup \{?\}$ for $i = 1, \dots, \ell$, and transition probability

$$\begin{aligned} W_{pd}(\mathbf{y}_\ell | \mathbf{x}_\ell; \sigma) &= W_{pd}(y_1, \dots, y_\ell | x_1, \dots, x_\ell; \sigma) \\ &= \prod_{i=1}^\ell p(y_i | x_i), \end{aligned}$$

where $x_i \in \mathcal{X}_i$, $y_i \in \mathcal{Y}_i$, and the probability $p(y_i | x_i)$ is equal to $W_{\sigma(i)}(y_i | x_i)$, for $i = 1, \dots, \ell$.

The capacity of the channel $W_{pd}(\mathbf{y}_\ell | \mathbf{x}_\ell; \sigma)$, denoted by C_σ , is given by

$$C_\sigma = \max_{p(x_1, \dots, x_\ell)} I_\sigma(X_1, \dots, X_\ell; Y_1, \dots, Y_\ell),$$

where $I_\sigma(X_1, \dots, X_\ell; Y_1, \dots, Y_\ell)$ represents the mutual information between the input vector (X_1, \dots, X_ℓ) and the output vector (Y_1, \dots, Y_ℓ) under the permutation σ of the component channels. It is known that the capacity C_σ is the sum of the capacities of the parallel channels [9, p. 41], i.e.,

$$C_\sigma = \sum_{i=1}^\ell C(W_i) = \sum_{i=1}^\mu \ell_i C(W_{\ell_i}) = \sum_{i=1}^\mu \ell_i (1 - \varepsilon_i). \quad (17)$$

Now, we consider the *compound* erasure channel W_{cc} that is the collection of erasure product channels $\{W_{pd}(\mathbf{y}_\ell | \mathbf{x}_\ell; \sigma), \sigma \in \Sigma\}$, where the set Σ represents all the $\frac{\ell!}{\ell_1! \ell_2! \dots \ell_\mu!}$ permutations of the multiset $\mathcal{T}(\mu)$ consisting of $\ell_1, \ell_2, \dots, \ell_\mu$ repetitions of the integers $1, 2, \dots, \mu$, respectively. During the code transmission over the compound channel W_{cc} , the permutation σ is fixed. However, neither the encoder nor the decoder knows which permutation σ is used. They only know the set of possible ℓ parallel QECs.

For each message $M \in \{1, \dots, q^{nR}\}$, the encoder generates a length- ℓn sequence $(x_1^n, x_2^n, \dots, x_\ell^n)$, which consists of ℓ length- n subsequences x_i^n , $1 \leq i \leq \ell$. Then, the encoder transmits these ℓ subsequences over the ℓ parallel channels simultaneously. The decoder receives a corresponding length- ℓn sequence $(y_1^n, y_2^n, \dots, y_\ell^n)$, and produces an estimate $\hat{M} \in \{1, \dots, q^{nR}\}$ or an error message. We assume the message M is uniformly distributed over $\{1, \dots, q^{nR}\}$. Under the permutation σ , the average probability of error is defined as $P_{e,\sigma}^n = \Pr\{\hat{M} \neq M | \sigma \text{ is selected}\}$. A rate $R > 0$ is said to be achievable if there exists a sequence of (q^{nR}, n) codes such that $\lim_{n \rightarrow \infty} P_{e,\sigma}^n = 0$ for all $\sigma \in \Sigma$. The capacity C_{cc} of the compound channel W_{cc} is the supremum over all achievable rates. The following result is from [9].

Proposition 13. (cf. [9, p. 170]) The capacity C_{cc} of the compound channel W_{cc} with no information about permutation σ available at either the encoder or the decoder is

$$C_{cc} = \max_{p(x_1, \dots, x_\ell)} \min_{\sigma \in \Sigma} I_\sigma(X_1, \dots, X_\ell; Y_1, \dots, Y_\ell).$$

We have the following upper bound on the capacity C_{cc} .

Lemma 14. The capacity C_{cc} of the compound channel W_{cc} satisfies $C_{cc} \leq \bar{C}_{cc} = \sum_{i=1}^{\mu} \ell_i(1 - \varepsilon_i)$.

Proof: By changing the order of max and min operations, we obtain

$$\begin{aligned} C_{cc} &= \max_{p(x_1, \dots, x_\ell)} \min_{\sigma \in \Sigma} I_\sigma(X_1, \dots, X_\ell; Y_1, \dots, Y_\ell) \\ &\leq \min_{\sigma \in \Sigma} \max_{p(x_1, \dots, x_\ell)} I_\sigma(X_1, \dots, X_\ell; Y_1, \dots, Y_\ell) \\ &= \min_{\sigma \in \Sigma} C_\sigma = \sum_{i=1}^{\mu} \ell_i(1 - \varepsilon_i), \end{aligned}$$

where the last step follows from (17) that the capacity $C_\sigma = \sum_{i=1}^{\mu} \ell_i(1 - \varepsilon_i)$ for every fixed permutation σ . ■

C. Capacity-Achieving ME-LRCs for the Compound Erasure Channel W_{cc}

We now show that the upper bound \bar{C}_{cc} in Lemma 14 can be achieved by a sequence of *deterministic* codes obtained from an explicit algebraic construction. In the following, we will use a generalized tensor product structure to construct a sequence of ME-LRCs that achieve the upper bound \bar{C}_{cc} on the capacity C_{cc} of the compound erasure channel W_{cc} . To this end, we first present a lemma on the existence of nested capacity-achieving linear codes over a set of QECs.

Lemma 15. Consider a set of μ QECs W_1, W_2, \dots, W_μ with erasure probabilities $\varepsilon_1 < \varepsilon_2 < \dots < \varepsilon_\mu$ and corresponding capacities $C(W_1) > C(W_2) > \dots > C(W_\mu)$. For any rates $R_1 > R_2 > \dots > R_\mu$ such that $R_i < C(W_i) = 1 - \varepsilon_i$, there exists a sequence of nested linear codes $\mathcal{C}_1^\mu = [n, k_\mu = R_\mu n]_q \subset \mathcal{C}_1^{\mu-1} = [n, k_{\mu-1} = R_{\mu-1} n]_q \subset \dots \subset \mathcal{C}_1^1 = [n, k_1 = R_1 n]_q$ such that the decoding error probability of each \mathcal{C}_1^i over the channel W_i , under maximum-likelihood (ML) decoding, satisfies $P_e^{(n)}(\mathcal{C}_1^i) \rightarrow 0$, as n goes to infinity.

Proof: See Appendix F. ■

Remark 6. Recently, it was shown in [22] that generalized Reed-Muller codes and primitive narrow-sense BCH codes over \mathbb{F}_q achieve the capacity of the QEC under block-ML decoding. Since these codes have nested structure [28], they are examples of nested capacity-achieving linear codes over a set of QECs. Polar codes also have the nested capacity-achieving property under successive cancellation (SC) decoding [2], [21].

Now, we present Construction B, which is based on the generalized tensor product structure in (5), to generate ME-LRCs that can achieve the capacity of the compound channel

W_{cc} . The key part of Construction B is to choose a set of appropriate nested capacity-achieving codes as component codes, whose existence is guaranteed by Lemma 15.

Construction B

Step 1: (Rate splitting) For any given rate $R < \bar{C}_{cc} = \sum_{i=1}^{\mu} \ell_i(1 - \varepsilon_i)$, choose a set of rates R_1, R_2, \dots, R_μ such that $\sum_{i=1}^{\mu} \ell_i R_i = R$ and $R_i < C(W_{\ell_i}) = 1 - \varepsilon_i$.

Step 2: Choose a set of $v_i \times n$ matrices H'_i , $i = 1, \dots, \mu$, over \mathbb{F}_q such that the corresponding parity-check matrices B_i in (4) generate a set of component codes \mathcal{C}'_i satisfying the following properties:

- 1) *Nested structure:* $\mathcal{C}'_\mu \subset \mathcal{C}'_{\mu-1} \subset \dots \subset \mathcal{C}'_1$, where $\mathcal{C}'_i = [n, n - \sum_{m=1}^i v_m]_q$ for $1 \leq i \leq \mu$.
- 2) *Capacity-achieving:* the code \mathcal{C}'_i has the required rate $R_i = \frac{n - \sum_{m=1}^i v_m}{n} < C(W_{\ell_i}) = 1 - \varepsilon_i$, and on the channel W_{ℓ_i} , the error probability satisfies $P_e^{(n)}(\mathcal{C}'_i) \rightarrow 0$, as n goes to infinity.

Step 3: Choose a parity-check matrix $H''_1 = I_{\ell \times \ell}$, i.e., the $\ell \times \ell$ identity matrix. For $i = 2, \dots, \mu$, choose a $\lambda_i \times \ell$ matrix H''_i over $\mathbb{F}_{q^{v_i}}$, which is a parity-check matrix of an $[\ell, \ell - \lambda_i, \delta_i = \lambda_i + 1]_{q^{v_i}}$ MDS code \mathcal{C}''_i , where the value of λ_i is chosen as $\lambda_i = \ell - \sum_{m=1}^{i-1} \ell_m$.

Step 4: Generate a parity-check matrix H over \mathbb{F}_q according to (5) with the matrices H'_i and H''_i , for $i = 1, 2, \dots, \mu$. The constructed code corresponding to the parity-check matrix H is referred to as \mathcal{C}_B .

Remark 7. In Step 1 of Construction B, there always exists a rate splitting. For example, for any $\delta > 0$, $R = \sum_{i=1}^{\mu} \ell_i(1 - \varepsilon_i) - \delta$, we can choose $R_i = (1 - \varepsilon_i) - \frac{\delta}{\ell}$ for $1 \leq i \leq \mu$.

In Step 2 of Construction B, explicit codes such as generalized Reed-Muller codes, BCH codes, and polar codes can be used as the component codes.

In Step 3 of Construction B, for a fixed ℓ , there always exists an $[\ell, \ell - \lambda_i, \delta_i = \lambda_i + 1]_{q^{v_i}}$ MDS code for a sufficiently large n . This is because such an MDS code exists whenever $\ell \leq q^{v_i} = q^{(R_{i-1} - R_i)n}$, so we only need $n \geq \lceil \frac{\log_q(\ell)}{R_{i-1} - R_i} \rceil$. When we analyze the capacity-achieving property of \mathcal{C}_B below, n is considered to go to infinity, so those MDS codes exist and Construction B is valid even when the underlying field size q is very small, e.g., $q = 2$.

On the contrary, in general the generalized integrated interleaving codes in (12) cannot be used to construct ME-LRCs that achieve the capacity of the compound channel W_{cc} , since they require the underlying field size $q > \ell$, which is not always satisfied, for example, as in the case of a binary compound channel W_{cc} with $q = 2$ and $\ell = 10$.

Note that, in contrast to Construction A, Construction B only specifies the rate and capacity-achieving properties of the component codes, with no specific reference to the minimum distance properties. The following theorem shows that the ME-LRC obtained from Construction B can achieve the capacity of the compound erasure channel W_{cc} .

Theorem 16. *The code \mathcal{C}_B is a $(\rho, n_0, k)_q$ ME-LRC with parameters $\rho = \ell$, $n_0 = n$, and $k = (n - v_1)\ell - \sum_{i=2}^{\mu} v_i \lambda_i$. Moreover, the ME-LRC is capacity-achieving over the compound channel W_{cc} , i.e., the error probability $P_e^n(\mathcal{C}_B) \rightarrow 0$, as n goes to infinity.*

Proof: We first verify that the parallel code rate $R_B = \frac{k}{n}$ of the constructed code \mathcal{C}_B equals R . To see this, we have

$$\begin{aligned} \frac{k}{n} &= \frac{(n - v_1)\ell - \sum_{i=2}^{\mu} v_i \lambda_i}{n} \\ &= \frac{(n - v_1)\ell - \sum_{i=2}^{\mu} v_i (\ell - \sum_{j=1}^{i-1} \ell_j)}{n} \\ &= \frac{(n - v_1)\ell_1}{n} + \frac{(n - \sum_{j=1}^2 v_j)\ell_2}{n} + \\ &\quad \dots + \frac{(n - \sum_{j=1}^{\mu} v_j)\ell_{\mu}}{n} \\ &= \sum_{i=1}^{\mu} \ell_i R_i. \end{aligned}$$

In the Step 1 of Construction B, we require $R = \sum_{i=1}^{\mu} \ell_i R_i$, so we have $R_B = R$.

Second, we prove that the code \mathcal{C}_B is capacity-achieving by showing that the decoding error probability $P_e^n(\mathcal{C}_B) \rightarrow 0$, as n goes to infinity.

We use Algorithm 1 to decode \mathcal{C}_B , where the component decoders \mathcal{D}'_i and \mathcal{D}''_i are chosen to be maximum-likelihood (ML) decoders as in Section III-C. From Algorithm 1, the decoding for \mathcal{C}_B has a total of μ levels. Let us consider a successful decoding event for \mathcal{C}_B over W_{cc} , denoted by E_s , and calculate its probability $P(E_s)$.

For the first level, we use the correct syndrome vector $(s_1^1, \dots, s_{\ell}^1) = \mathbf{0}$ to decode all the sub-blocks over the ℓ QECs using the ML erasure decoding. For each sub-block, the capacity of its corresponding QEC is unknown. However, the sub-blocks over the ℓ_1 QECs (each with capacity $1 - \varepsilon_1$) will be decoded successfully with a high probability which can be expressed as $P_1 = (1 - P_e^n(\mathcal{C}'_1))^{\ell_1}$.

For the second level, since the ℓ_1 sub-blocks have been corrected in the first level, the number of uncorrected sub-blocks is at most $\sum_{i=2}^{\mu} \ell_i$. These uncorrected sub-blocks can be detected, because the ML erasure decoder does not produce any miscorrections. As a result, the correct syndrome vector $(s_1^2, \dots, s_{\ell}^2)$ can be obtained. Using the correct syndrome vectors $(s_1^i, \dots, s_{\ell}^i)$, $i = 1, 2$, the sub-blocks over the ℓ_2 QECs (each with capacity $1 - \varepsilon_2$) are corrected. The probability associated with this is $P_2 \geq (1 - P_e^n(\mathcal{C}'_2))^{\ell_2}$; we use a lower bound here since according to Algorithm 1, some of the sub-blocks over the ℓ_2 QECs may have already been corrected in the first level.

Similarly, for the m th level, $3 \leq m \leq \mu$, since $\sum_{i=1}^{m-1} \ell_i$ sub-blocks have been corrected in the previous levels, the number of uncorrected sub-blocks is at most $\sum_{i=m}^{\mu} \ell_i$. As a result, the correct syndrome vector $(s_1^m, \dots, s_{\ell}^m)$ can be obtained. Using the correct syndrome vectors $(s_1^i, \dots, s_{\ell}^i)$, $i = 1, 2, \dots, m$, the sub-blocks over the ℓ_m QECs (each with

capacity $1 - \varepsilon_m$) are corrected. The corresponding probability is $P_m \geq (1 - P_e^n(\mathcal{C}'_m))^{\ell_m}$.

Thus, the probability of successful decoding $P_s^n(\mathcal{C}_B)$ of \mathcal{C}_B can be lower bounded as

$$P_s^n(\mathcal{C}_B) \geq P(E_s) = \prod_{i=1}^{\mu} P_i \geq \prod_{i=1}^{\mu} (1 - P_e^n(\mathcal{C}'_i))^{\ell_i}.$$

Correspondingly, we can upper bound the decoding error probability $P_e^n(\mathcal{C}_B)$ of \mathcal{C}_B as

$$\begin{aligned} P_e^n(\mathcal{C}_B) &= 1 - P_s^n(\mathcal{C}_B) \\ &\leq 1 - \prod_{i=1}^{\mu} (1 - P_e^n(\mathcal{C}'_i))^{\ell_i}. \end{aligned} \quad (18)$$

From the capacity-achieving property of the chosen component codes, we already have $P_e^n(\mathcal{C}'_i) \rightarrow 0$ as n goes to infinity, so in (18), $P_e^n(\mathcal{C}_B) \rightarrow 0$ as n goes to infinity. Thus, we conclude that \mathcal{C}_B can achieve the capacity of the compound channel W_{cc} . ■

VII. CONCLUSION

In this work, we presented a general construction for ME-LRCs over small fields. This construction yields optimal ME-LRCs with respect to an upper bound on the minimum distance for a wide range of code parameters. Then, an erasure decoder was proposed and corresponding correctable erasure patterns were identified. ME-LRCs based on Reed-Solomon codes were shown to be optimal among all codes having the same erasure-correcting capability. In addition, generalized integrated interleaving codes were proved to be a subclass of generalized tensor product codes, thus giving the exact relation between the two classes of codes. Finally, we investigated ME-LRCs over a compound erasure product channel, and we showed that a generalized tensor product structure can be employed to construct capacity-achieving ME-LRCs for such a channel.

APPENDIX A PROOF OF LEMMA 2

Proof: For the case of $x = 0$, it is trivial. For $1 \leq x \leq \lceil \frac{k}{k^*} \rceil - 1$, $x \in \mathbb{Z}^+$, let \mathcal{I} represent the set of the coordinates of the first x rows in the array. Thus, $|\mathcal{I}| = xn_0$. First, consider the code $\mathcal{C}_{\mathcal{I}} = \{c_{\mathcal{I}} : c \in \mathcal{C}\}$ whose dimension is denoted by $k_{\mathcal{I}}$, which satisfies $k_{\mathcal{I}} \leq xk^*$. Then, we consider the code $\mathcal{C}_{\mathcal{I}}^0 = \{c_{[\rho n_0] \setminus \mathcal{I}} : c_{\mathcal{I}} = \mathbf{0} \text{ and } c \in \mathcal{C}\}$. Since the code \mathcal{C} is linear, the size of the code $\mathcal{C}_{\mathcal{I}}^0$ is $q^{k - k_{\mathcal{I}}}$ and it is a linear code as well. Moreover, the minimum distance \hat{d} of the code $\mathcal{C}_{\mathcal{I}}^0$ is at least d , i.e., $\hat{d} \geq d$.

Thus, we get an upper bound on the minimum distance d ,

$$\begin{aligned} d &\leq \hat{d} \leq d_{opt}^{(q)}[\rho n_0 - |\mathcal{I}|, k - k_{\mathcal{I}}] \\ &\leq d_{opt}^{(q)}[\rho n_0 - xn_0, k - xk^*]. \end{aligned}$$

Similarly, we also get an upper bound on the dimension k ,

$$k - k_{\mathcal{I}} \leq k_{opt}^{(q)}[\rho n_0 - |\mathcal{I}|, \hat{d}] \leq k_{opt}^{(q)}[\rho n_0 - xn_0, d].$$

Therefore, we conclude that

$$k \leq k_{opt}^{(q)}[\rho n_0 - x n_0, d] + k_{\mathcal{I}} \leq k_{opt}^{(q)}[\rho n_0 - x n_0, d] + x k^*. \quad \blacksquare$$

APPENDIX B PROOF OF LEMMA 3

Proof: We can construct a $(\rho, n_0, k; \geq d_0, \geq d)_q$ ME-LRC in two steps, and use the GV bound [28] twice. First, there exists a $[\rho(n_0 - r_0), k, \geq d]_q$ array code \mathcal{G}_1 of size $\rho \times (n_0 - r_0)$ where r_0 is an integer $0 \leq r_0 < n_0$, if the parameters satisfy

$$\sum_{i=0}^{d-2} \binom{\rho(n_0 - r_0) - 1}{i} (q-1)^i < q^{\rho(n_0 - r_0) - k}. \quad (19)$$

Second, there exists a length- n_0 code \mathcal{G}_2 with minimum distance at least d_0 , if its redundancy r_0 satisfies

$$r_0 > \log_q \left(\sum_{i=0}^{d_0-2} \binom{n_0-1}{i} (q-1)^i \right). \quad (20)$$

Now, we encode each row of the code \mathcal{G}_1 using the code \mathcal{G}_2 by adding r_0 more redundancy symbols. The resulting code is a $(\rho, n_0, k; \geq d_0, \geq d)_q$ ME-LRC. Let $r_0 = \lceil \log_q \left(\sum_{i=0}^{d_0-2} \binom{n_0-1}{i} (q-1)^i \right) \rceil$, and substitute it into (19), producing (3). \blacksquare

APPENDIX C PROOF OF THEOREM 5

Proof: A codeword \mathbf{x} in C_{GTP}^μ is an $n'\ell$ -dimensional vector over \mathbb{F}_q , denoted by $\mathbf{x} = (x_1, x_2, \dots, x_\ell)$, where x_i in \mathbf{x} is an n' -dimensional vector, for $i = 1, 2, \dots, \ell$. Let $\mathbf{s}_i^j = x_i H_j^{T'}$, for $i = 1, 2, \dots, \ell$ and $j = 1, 2, \dots, \mu$. Thus, \mathbf{s}_i^j is a v_j -dimensional vector over \mathbb{F}_q , and is considered as an element in $\mathbb{F}_{q^{v_j}}$. Let $\mathbf{s}^j = (s_1^j, s_2^j, \dots, s_\ell^j)$ be the ℓ -dimensional vector over $\mathbb{F}_{q^{v_j}}$ whose components are s_i^j , $i = 1, 2, \dots, \ell$.

To prove the theorem, we consider separately the two possibilities:

- 1) $\mathbf{s}^j \neq \mathbf{0}$ for some $1 \leq j \leq \mu$.
- 2) $\mathbf{s}^j = \mathbf{0}$ for all $1 \leq j \leq \mu$.

First, note that the condition $\mathbf{x} H^T = \mathbf{0}$ implies that $\mathbf{s}^j H_j^{T''} = \mathbf{0}$ for all $1 \leq j \leq \mu$.

Now, consider the first possibility, namely that $\mathbf{s}^j \neq \mathbf{0}$ for some $1 \leq j \leq \mu$, and let $1 \leq j \leq \mu$ be the smallest positive integer such that $\mathbf{s}^j \neq \mathbf{0}$. If $j = 1$, then $\mathbf{s}^1 \neq \mathbf{0}$, and the condition $\mathbf{s}^1 H_1^{T''} = \mathbf{0}$ means that \mathbf{s}^1 is a codeword in the code \mathcal{C}_1'' defined by H_1'' . This implies that $w_{q^{v_1}}(\mathbf{s}^1) \geq \delta_1$. Since $w_q(\mathbf{x}) \geq w_{q^{v_1}}(\mathbf{s}^1)$, we conclude that $w_q(\mathbf{x}) \geq \delta_1$.

Next, suppose that $2 \leq j \leq \mu$. Then $\mathbf{s}^i = \mathbf{0}$ in $(\mathbb{F}_{q^{v_i}})^\ell$, for $i = 1, \dots, j-1$. This means that $x_i B_{j-1}^T = \mathbf{0}$ for $i = 1, 2, \dots, \ell$; that is, x_i is a codeword in the code \mathcal{C}'_{j-1} defined by the parity-check matrix B_{j-1} , whose minimum distance is d'_{j-1} . Therefore, we have $w_q(x_i) \geq d'_{j-1}$ if

$x_i \neq \mathbf{0}$, $i = 1, 2, \dots, \ell$. Now, since $\mathbf{s}^j \neq \mathbf{0}$, the condition $\mathbf{s}^j H_j^{T''} = \mathbf{0}$ means that \mathbf{s}^j is a codeword in the code \mathcal{C}_j'' defined by H_j'' . Therefore, $w_{q^{v_j}}(\mathbf{s}^j) \geq \delta_j$. It follows that $w_q(\mathbf{x}) \geq w_{q^{v_j}}(\mathbf{s}^j) d'_{j-1} \geq \delta_j d'_{j-1}$.

From consideration of the first possibility, therefore, we conclude that

$$d_t \geq \min\{\delta_1, \delta_2 d'_1, \dots, \delta_\mu d'_{\mu-1}\}. \quad (21)$$

Now, we turn to the second possibility, namely that $\mathbf{s}^j = \mathbf{0}$ for all $1 \leq j \leq \mu$. This means that $x_i B_\mu^T = \mathbf{0}$ for $i = 1, 2, \dots, \ell$; that is, x_i is a codeword in the code \mathcal{C}'_μ defined by the parity-check matrix B_μ , whose minimum distance is d'_μ . Therefore, we have $w_q(x_i) \geq d'_\mu$ if $x_i \neq \mathbf{0}$, $i = 1, 2, \dots, \ell$. Since $\mathbf{x} \neq \mathbf{0}$, some $x_i \neq \mathbf{0}$, so we conclude that

$$w_q(\mathbf{x}) \geq d'_\mu. \quad (22)$$

Combining (21) and (22), we conclude that

$$d_t \geq \min\{\delta_1, \delta_2 d'_1, \dots, \delta_\mu d'_{\mu-1}, d'_\mu\}.$$

This completes the proof. \blacksquare

APPENDIX D PROOF OF THEOREM 7

Proof: The proof follows from the decoding procedure of the decoder \mathcal{D}_A . The ME-LRC \mathcal{C}_A has $d_0 = d'_1$ and $d = d'_\mu$. For a received word $\mathbf{y} = (y_1, y_2, \dots, y_\ell)$, each vector \mathbf{y}_i , $1 \leq i \leq \ell$, corresponds to a row in the array.

For the first level, since $\delta_1 = \infty$, the correct syndrome vector (s_1^1, \dots, s_ℓ^1) is the all-zero vector, i.e., $(s_1^1, \dots, s_\ell^1) = \mathbf{0}$. Thus, the rows with at most $d'_1 - 1$ erasures are corrected.

For the second level, the remaining uncorrected row $\hat{\mathbf{c}}_j$, $j \in \mathcal{F}$, has at least d'_1 erasures. The total number of such uncorrected rows with indices in \mathcal{F} is less than δ_2 , because we require $N_1 \leq \delta_2 - 1$ in the condition (6). Thus, the correct syndrome vector (s_1^2, \dots, s_ℓ^2) can be obtained. As a result, the rows with at most $d'_2 - 1$ erasures are corrected.

Similarly, by induction, if the decoder runs until the μ th level, the remaining uncorrected row $\hat{\mathbf{c}}_j$, $j \in \mathcal{F}$, has at least $d'_{\mu-1}$ erasures. The total number of such uncorrected rows with indices in \mathcal{F} is less than δ_μ , because we require $N_{\mu-1} \leq \delta_\mu - 1$ in the condition (6). Therefore, all the correct syndrome vectors (s_1^i, \dots, s_ℓ^i) , $i = 1, 2, \dots, \mu$, are obtained. On the other hand, the remaining uncorrected row $\hat{\mathbf{c}}_j$, $j \in \mathcal{F}$, has at most $d'_\mu - 1$ erasures, since we also require $N_\mu \leq 0$ in the condition (6). Thus, all of these uncorrected rows can be corrected in this step using all these correct syndromes. \blacksquare

APPENDIX E PROOF OF COROLLARY 8

Proof: The ME-LRC \mathcal{C}_A has $d_0 = d'_1$ and $d = d'_\mu$. We only need to show that the received word \mathbf{y} with any $d'_\mu - 1$ erasures satisfies the condition in Theorem 7. We prove it by contradiction. If the condition is not satisfied, there exists an integer i , $1 \leq i \leq \mu$, such that $N_i \geq \delta_{i+1}$. Therefore,

we have $w_e(\mathbf{y}) \geq d'_i \delta_{i+1} \geq d'_\mu$, where the last inequality is from the requirement of Construction A. Thus, we get a contradiction of the assumption that the received word \mathbf{y} has $d'_\mu - 1$ erasures. ■

APPENDIX F
PROOF OF LEMMA 15

Proof: To prove the lemma, we will use the following result for the QEC, which is a consequence of Theorem 6.2.1 of Gallager [11, p. 206].

Lemma 17. *For the QEC W with erasure probability ε , let n and nR be integers such that $R < C(W) = 1 - \varepsilon$. Let $P_e^{(n)}(\mathcal{C})$ denote the average of $P_e^{(n)}(\mathcal{C})$ over all linear $[n, nR]_q$ codes \mathcal{C} under maximum-likelihood decoding. Then,*

$$\overline{P_e^{(n)}(\mathcal{C})} \leq q^{-nE_q(\varepsilon, R)},$$

where $E_q(\varepsilon, R)$ is the random coding error exponent of Gallager and $E_q(\varepsilon, R) > 0$ for all R satisfying $0 \leq R < C(W)$.

The following lemma is a direct consequence of Lemma 17.

Lemma 18. *For every $\rho \in (0, 1]$, at least a fraction $1 - \rho$ (i.e., $\geq 1 - \rho$) of all linear $[n, nR]_q$ codes \mathcal{C} satisfy*

$$P_e^{(n)}(\mathcal{C}) \leq (1/\rho)q^{-nE_q(\varepsilon, R)}.$$

Proof: The proof is based on contradiction. Consider the set \mathcal{S} of codes \mathcal{C} for which $P_e^{(n)}(\mathcal{C}) > (1/\rho)q^{-nE_q(\varepsilon, R)}$. Assume that \mathcal{S} forms more than a fraction ρ of all linear $[n, nR]_q$ codes \mathcal{C} . Then, we have

$$\overline{P_e^{(n)}(\mathcal{C})} > \frac{\rho}{|\mathcal{S}|} \sum_{\mathcal{C} \in \mathcal{S}} P_e^{(n)}(\mathcal{C}) > q^{-nE_q(\varepsilon, R)},$$

which contradicts Lemma 17. Therefore, \mathcal{S} only forms at most a fraction ρ of all linear $[n, nR]_q$ codes \mathcal{C} . ■

With the above two lemmas, we are ready to prove Lemma 15.

Consider an ensemble \mathcal{G}_1 of all $k_1 \times n$ full rank matrices over \mathbb{F}_q . The size of \mathcal{G}_1 is $|\mathcal{G}_1| = (q^n - 1)(q^n - q) \cdots (q^n - q^{k_1 - 1})$. Now, for each matrix $G_i^1 \in \mathcal{G}_1$, $1 \leq i \leq |\mathcal{G}_1|$, take the last k_2 rows to form a new matrix G_i^2 . All these new matrices form a new ensemble \mathcal{G}_2 , including possible repetitions. It is clear that $|\mathcal{G}_2| = |\mathcal{G}_1|$ and in \mathcal{G}_2 , each $k_2 \times n$ full rank matrix over \mathbb{F}_q appears $(q^n - q^{k_2})(q^n - q^{k_2 + 1}) \cdots (q^n - q^{k_1 - 1})$ times. Similarly, for each matrix $G_i^1 \in \mathcal{G}_1$, $1 \leq i \leq |\mathcal{G}_1|$, take the last k_j , $3 \leq j \leq \mu$, rows to form a new matrix G_i^j . All these new matrices form a new ensemble \mathcal{G}_j . It is clear that $|\mathcal{G}_j| = |\mathcal{G}_1|$ and in \mathcal{G}_j , each $k_j \times n$ full rank matrix over \mathbb{F}_q appears $(q^n - q^{k_j})(q^n - q^{k_j + 1}) \cdots (q^n - q^{k_1 - 1})$ times.

Note that the number of generator matrices of a linear $[n, k]_q$ code is the same for all such codes. Therefore, from Lemma 18, in each ensemble \mathcal{G}_j for $1 \leq j \leq \mu$, at least

a fraction x of all matrices in this ensemble will generate linear codes \mathcal{C} such that the error probability $P_e^{(n)}(\mathcal{C}) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_j, R_j)}$.

Now, choose x to be a certain value satisfying $\frac{1}{2} < x < 1$. Let \mathcal{S}_1 be the subset of the ensemble \mathcal{G}_1 such that $\frac{|\mathcal{S}_1|}{|\mathcal{G}_1|} \geq x$ and each matrix in \mathcal{S}_1 generates a linear code \mathcal{C}_1^1 with the error probability $P_e^{(n)}(\mathcal{C}_1^1) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_1, R_1)}$. Let \mathcal{S}_2 be the subset of the ensemble \mathcal{G}_1 such that $\frac{|\mathcal{S}_2|}{|\mathcal{G}_1|} \geq x$ and for each matrix in \mathcal{S}_2 , its last k_2 rows generate a linear code \mathcal{C}_1^2 with the error probability $P_e^{(n)}(\mathcal{C}_1^2) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_2, R_2)}$. Then, using basic properties of set operations, we have

$$\begin{aligned} \frac{|\mathcal{S}_1 \cap \mathcal{S}_2|}{|\mathcal{G}_1|} &= \frac{|\mathcal{S}_1|}{|\mathcal{G}_1|} + \frac{|\mathcal{S}_2|}{|\mathcal{G}_1|} - \frac{|\mathcal{S}_1 \cup \mathcal{S}_2|}{|\mathcal{G}_1|} \\ &\geq \frac{|\mathcal{S}_1|}{|\mathcal{G}_1|} + \frac{|\mathcal{S}_2|}{|\mathcal{G}_1|} - 1 \\ &\geq 2x - 1 > 0. \end{aligned}$$

Thus, we find a non-empty subset $\mathcal{S}_{12} = \mathcal{S}_1 \cap \mathcal{S}_2$ in the ensemble \mathcal{G}_1 such that: 1) \mathcal{S}_{12} has at least a fraction $2x - 1 > 0$ of all the matrices in \mathcal{G}_1 , and 2) each matrix in \mathcal{S}_{12} generates a linear code \mathcal{C}_1^1 with the error probability $P_e^{(n)}(\mathcal{C}_1^1) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_1, R_1)}$, and its last k_2 rows generate a linear code \mathcal{C}_1^2 with the error probability $P_e^{(n)}(\mathcal{C}_1^2) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_2, R_2)}$.

Similarly, arguing as above, it is not hard to see that for any x satisfying $\frac{\mu-1}{\mu} < x < 1$, in the ensemble \mathcal{G}_1 , we can find a non-empty subset $\overline{\mathcal{G}}_1 \subseteq \mathcal{G}_1$ such that: 1) $\overline{\mathcal{G}}_1$ has at least a fraction $\mu x - (\mu - 1) > 0$ of all the matrices in \mathcal{G}_1 , and 2) for each matrix \overline{G}_1 in $\overline{\mathcal{G}}_1$, for each j , $1 \leq j \leq \mu$, the last k_j rows of \overline{G}_1 will generate a linear code \mathcal{C}_1^j with the error probability $P_e^{(n)}(\mathcal{C}_1^j) \leq (\frac{1}{1-x})q^{-nE_q(\varepsilon_j, R_j)}$.

Thus, there exists a sequence of nested linear codes $\mathcal{C}_1^\mu = [n, k_\mu = R_\mu n]_q \subset \mathcal{C}_1^{\mu-1} = [n, k_{\mu-1} = R_{\mu-1} n]_q \subset \cdots \subset \mathcal{C}_1^1 = [n, k_1 = R_1 n]_q$ such that for all $1 \leq i \leq \mu$, the error probability $P_e^{(n)}(\mathcal{C}_1^i) \rightarrow 0$, as n goes to infinity. ■

REFERENCES

- [1] A. Agarwal, A. Barg, S. Hu, A. Mazumdar, and I. Tamo, "Combinatorial alphabet-dependent bounds for locally recoverable codes," *IEEE Trans. Inf. Theory*, vol. 64, no. 5, pp. 3481–3492, May 2018.
- [2] E. Arıkan, "Channel polarization: A method for constructing capacity-achieving codes for symmetric binary-input memoryless channels," *IEEE Trans. Inf. Theory*, vol. 55, no. 7, pp. 3051–3073, July 2009.
- [3] A. Barg, I. Tamo, and S. Vlăduț, "Locally recoverable codes on algebraic curves," *IEEE Trans. Inf. Theory*, vol. 63, no. 8, pp. 4928–4939, Aug. 2017.
- [4] M. Blaum, J. L. Hafner, and S. Hetzler, "Partial-MDS codes and their application to RAID type of architectures," *IEEE Trans. Inf. Theory*, vol. 59, no. 7, pp. 4510–4519, July 2013.
- [5] M. Blaum and S. R. Hetzler, "Integrated interleaved codes as locally recoverable codes: Properties and performance," *International Journal of Information and Coding Theory*, vol. 3, no. 4, pp. 324–344, 2016.
- [6] M. Bossert, H. Griefßer, J. Maucher, and V. V. Zyablov, "Some results on generalized concatenation of block codes," in *Applied Algebra, Algebraic Algorithms and Error-Correcting Codes (AAECC), M. Fossorier, H. Imai, S. Lin, and A. Poli, Eds., (Lecture Notes in Computer Science, vol 1719)*. Berlin: Springer, 1999, pp. 181–190.

- [7] V. Cadambe and A. Mazumdar, "An upper bound on the size of locally recoverable codes," in *Proc. IEEE Int. Symp. Network Coding (NetCod)*, Calgary, AB, Canada, June 2013, pp. 1–5.
- [8] G. Calis and O. O. Koyluoglu, "A general construction for PMDS codes," *IEEE Commun. Lett.*, vol. 21, no. 3, pp. 452–455, Mar. 2017.
- [9] A. El Gamal and Y.-H. Kim, *Network Information Theory*. New York: Cambridge University Press, 2011.
- [10] R. Gabrys, E. Yaakobi, M. Blaum, and P. H. Siegel, "Constructions of partial MDS codes over small fields," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Aachen, Germany, June 2017, pp. 1–5.
- [11] R. G. Gallager, *Information Theory and Reliable Communication*. New York: Wiley, 1968.
- [12] G. A. Gibson, *Redundant Disk Arrays: Reliable, Parallel Secondary Storage*. Cambridge, MA: MIT Press, 1992.
- [13] P. Gopalan, C. Huang, H. Simitci, and S. Yekhanin, "On the locality of codeword symbols," *IEEE Trans. Inf. Theory*, vol. 58, no. 11, pp. 6925–6934, Nov. 2012.
- [14] S. Goparaju and R. Calderbank, "Binary cyclic codes that are locally repairable," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Honolulu, HI, USA, June–July 2014, pp. 676–680.
- [15] M. Hassner, K. Abdel-Ghaffar, A. Patel, R. Koetter, and B. Trager, "Integrated interleaving – a novel ECC architecture," *IEEE Trans. Magn.*, vol. 37, no. 2, pp. 773–775, Mar. 2001.
- [16] P. Huang, E. Yaakobi, and P. H. Siegel, "Multi-erasure locally recoverable codes over small fields," in *Proc. 55th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Monticello, IL, USA, Oct. 2017, pp. 1123–1130.
- [17] P. Huang, E. Yaakobi, H. Uchikawa, and P. H. Siegel, "Cyclic linear binary locally repairable codes," in *Proc. IEEE Inf. Theory Workshop (ITW)*, Jerusalem, Israel, Apr.–May 2015, pp. 1–5.
- [18] —, "Linear locally repairable codes with availability," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Hong Kong, China, June 2015, pp. 1871–1875.
- [19] —, "Binary linear locally repairable codes," *IEEE Trans. Inf. Theory*, vol. 62, no. 11, pp. 6268–6283, Nov. 2016.
- [20] H. Imai and H. Fujiya, "Generalized tensor product codes," *IEEE Trans. Inf. Theory*, vol. 27, no. 2, pp. 181–187, Mar. 1981.
- [21] S. B. Korada, "Polar codes for channel and source coding," Ph.D. dissertation, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland, 2009.
- [22] S. Kudekar, S. Kumar, M. Mondelli, H. D. Pfister, E. Şaşıoğlu, and R. L. Urbanke, "Reed–Muller codes achieve capacity on erasure channels," *IEEE Trans. Inf. Theory*, vol. 63, no. 7, pp. 4298–4316, July 2017.
- [23] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*. New York: Elsevier, 1977.
- [24] J. Maucher, V. V. Zyablov, and M. Bossert, "On the equivalence of generalized concatenated codes and generalized error location codes," *IEEE Trans. Inf. Theory*, vol. 46, no. 2, pp. 642–649, Mar. 2000.
- [25] F. Oggier and A. Datta, "Self-repairing homomorphic codes for distributed storage systems," in *Proc. IEEE Int. Conf. Comput. Commun. (INFOCOM)*, Shanghai, China, Apr. 2011, pp. 1215–1223.
- [26] D. S. Papailiopoulos and A. G. Dimakis, "Locally repairable codes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, MA, USA, July 2012, pp. 2771–2775.
- [27] N. Prakash, G. Kamath, V. Lalitha, and P. Kumar, "Optimal linear codes with a local-error-correction property," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Cambridge, MA, USA, July 2012, pp. 2776–2780.
- [28] R. Roth, *Introduction to Coding Theory*. New York: Cambridge University Press, 2006.
- [29] I. Tamo and A. Barg, "A family of optimal locally recoverable codes," *IEEE Trans. Inf. Theory*, vol. 60, no. 8, pp. 4661–4676, Aug. 2014.
- [30] I. Tamo, A. Barg, and A. Frolov, "Bounds on the parameters of locally recoverable codes," *IEEE Trans. Inf. Theory*, vol. 62, no. 6, pp. 3070–3083, June 2016.
- [31] I. Tamo, A. Barg, S. Goparaju, and R. Calderbank, "Cyclic LRC codes and their subfield subcodes," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Hong Kong, China, June 2015, pp. 1262–1266.
- [32] X. Tang and R. Koetter, "A novel method for combining algebraic decoding and iterative processing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, Seattle, WA, USA, July 2006, pp. 474–478.
- [33] M. Tsfasman, S. Vladut, and D. Nogin, *Algebraic Geometric Codes: Basic Notions*. American Mathematical Society, 2007, no. 139.
- [34] J. Wolf, "On codes derivable from the tensor product of check matrices," *IEEE Trans. Inf. Theory*, vol. 11, no. 2, pp. 281–284, Apr. 1965.
- [35] Y. Wu, "Generalized integrated interleaved codes," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 1102–1119, Feb. 2017.
- [36] K. Yang and P. Kumar, "On the true minimum distance of Hermitian codes," in *Coding Theory and Algebraic Geometry, H. Stichtenoth and M.A. Tsfasman, Eds. (Lecture Notes in Mathematics, vol. 1518)*. Berlin: Springer, 1992, pp. 99–107.
- [37] X. Zhang, "Modified generalized integrated interleaved codes for local erasure recovery," *IEEE Commun. Lett.*, vol. 21, no. 6, pp. 1241–1244, June 2017.

Pengfei Huang (S'13–M'20) received the B.E. degree in electrical engineering from Zhejiang University, Hangzhou, China, in 2010, the M.S. degree in electrical engineering from Shanghai Jiao Tong University, Shanghai, China, in 2013, and the Ph.D. degree in electrical engineering from the University of California San Diego, CA, USA, in 2018. From 2014 to 2018, he was associated with the Center for Memory and Recording Research. Since 2018, he has been with Western Digital Corporation conducting research and development on enterprise solid-state drives. His current research interests are coding for distributed storage systems and non-volatile memories.

Eitan Yaakobi (S'07–M'12–SM'17) is an Associate Professor at the Computer Science Department at the Technion — Israel Institute of Technology. He received the B.A. degrees in computer science and mathematics, and the M.Sc. degree in computer science from the Technion — Israel Institute of Technology, Haifa, Israel, in 2005 and 2007, respectively, and the Ph.D. degree in electrical engineering from the University of California, San Diego, in 2011. Between 2011–2013, he was a postdoctoral researcher in the department of Electrical Engineering at the California Institute of Technology and at the Center for Memory and Recording Research at the University of California, San Diego. His research interests include information and coding theory with applications to non-volatile memories, associative memories, DNA storage, data storage and retrieval, and private information retrieval. He received the Marconi Society Young Scholar in 2009 and the Intel Ph.D. Fellowship in 2010–2011.

Paul H. Siegel (M'82–SM'90–F'97–LF'19) received the S.B. and Ph.D. degrees in mathematics from Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 1975 and 1979, respectively. He held a Chaim Weizmann Postdoctoral Fellowship with the Courant Institute, New York University, New York, NY, USA. He was with the IBM Research Division, San Jose, CA, USA, from 1980 to 1995. He joined the faculty of the University of California, San Diego, CA, USA, in July 1995, where he is currently a Distinguished Professor of Electrical and Computer Engineering in the Jacobs School of Engineering. He is affiliated with the Center for Memory and Recording Research where he holds an Endowed Chair and served as Director from 2000 to 2011. His research interests include information theory and communications, particularly coding and modulation techniques, with applications to digital data storage and transmission. He was a Member of the Board of Governors of the IEEE Information Theory Society from 1991 to 1996 and again from 2009 to 2014. He served as Co-Guest Editor of the May 1991 Special Issue on

“Coding for Storage Devices” of the IEEE Transactions on Information Theory. He served the same Transactions as Associate Editor for Coding Techniques from 1992 to 1995, and as Editor-in-Chief from July 2001 to July 2004. He was also Co-Guest Editor of the May/September 2001 two-part issue on “The Turbo Principle: From Theory to Practice” and the February 2016 issue on “Recent Advances in Capacity Approaching Codes” of the IEEE Journal on Selected Areas in Communications. He is a member of the National Academy of Engineering. He was the 2015 Padovani Lecturer of the IEEE Information Theory Society. He was the corecipient of the 2007 Best Paper Award in Signal Processing and Coding for Data Storage from the Data Storage Technical Committee of the IEEE Communications Society. He was the corecipient of the 1992 IEEE Information Theory Society Paper Award and the 1993 IEEE Communications Society Leonard G. Abraham Prize Paper Award.