Hierarchical Erasure Correction of Linear Codes[☆]

Netanel Raviv[a,*], Moshe Schwartz[b], Rami Cohen[c], Yuval Cassuto[c]

[a] *Department of Computer Science and Engineering, Washington University in St. Louis, One Brookings Dr., St. Louis, MO 63130*
[b] *School of Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer-Sheva, Israel*
[c] *Department of Electrical Engineering, Technion–Israel Institute of Technology, Haifa, Israel*

Abstract

Linear codes over finite extension fields have widespread applications in theory and practice. In some scenarios, the decoder has a sequential access to the codeword symbols, giving rise to a hierarchical erasure structure. In this paper we develop a mathematical framework for hierarchical erasures over extension fields, provide several bounds and constructions, and discuss potential applications in distributed storage and flash memories. Our results show intimate connection to Universally Decodable Matrices, as well as to Reed-Solomon and Gabidulin codes.

Keywords: linear codes, erasure-correcting codes, hierarchical erasures
2010 MSC: 11T71, 94B05

[☆] Parts of this paper were presented in the 2018 International Symposium on Information Theory (ISIT) [1]. Parts of Netanel Raviv's work were done when he was a postdoctoral researcher at the Department of Electrical Engineering, California Institute of Technology, Pasadena, California. The work of Moshe Schwartz was supported in part by a German Israeli Project Cooperation (DIP) grant under grant no. PE2398/1-1. The work of Yuval Cassuto was supported in part by the Israel Science Foundation and by the US-Israel Binational Science Foundation.

[*] Corresponding author.

Email addresses: `netanel.raviv@wustl.edu` (Netanel Raviv), `schwartz@ee.bgu.ac.il` (Moshe Schwartz), `rc@technion.ac.il` (Rami Cohen), `ycassuto@ee.technion.ac.il` (Yuval Cassuto)

1. Introduction

For a prime power q, let $\mathbb{F}_q$ be the finite field with q elements. For a positive integer α, let $\mathbb{F}_{q^\alpha}$ be its algebraic extension of degree α, that can be viewed as a vector space of dimension α over $\mathbb{F}_q$ by fixing an ordered basis $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_\alpha)$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$. For an integer n, a code $\mathcal{C} \subseteq \mathbb{F}_{q^\alpha}^n$ is called *linear over* $\mathbb{F}_{q^\alpha}$ (or linear, in short) if it is a linear subspace of $\mathbb{F}_{q^\alpha}^n$, in which case its dimension is denoted by k.

Traditionally, the coding-theoretic literature discusses encoding and decoding of linear codes under *erasures*, i.e., where codeword symbols are replaced by some symbol $*$ outside the field, and *errors*, where codeword symbols are replaced by arbitrary field elements. The mathematical framework for erasures and errors is very well understood, and bounds and matching constructions are well known in most cases.

However, in some scenarios, the decoder receives each codeword symbol *sequentially*, i.e., each codeword symbol is received in some gradual manner, rather than instantaneously. When these scenarios involve codes over $\mathbb{F}_{q^\alpha}$, codeword symbols are viewed as vectors over $\mathbb{F}_q$, and the decoder receives these vectors one $\mathbb{F}_q$ element after another. In this paper we study bounds and code constructions for this scenario. That is, codes that enable the decoder to complete the decoding process once sufficiently many $\mathbb{F}_q$ symbols are obtained regardless of their source, and in particular, even if $\mathbb{F}_{q^\alpha}$-symbols have not been obtained in full. Practical applications which present this behavior, for which our techniques are useful, are discussed in the sequel.

In the next section we lay the mathematical framework by which we study the problem, discuss potential applications, and summarize our contributions. Several constructions of codes capable of correcting hierarchical erasures are given in Section 3 while upper and lower bounds are discussed in Section 4.

2. Preliminaries

2.1. Framework and Problem Definition

Let $\mathbf{c} = (c_i)_{i=1}^n \in \mathbb{F}_{q^\alpha}^n$ be a codeword in a linear code. By fixing a basis[1] $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_\alpha)$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, consider each c_i as a vector in $\mathbb{F}_q^\alpha$, and denote (by abuse of notation) $c_i = (c_{i,1}, \ldots, c_{i,\alpha})$ where $c_i = \sum_{j=1}^{\alpha} c_{i,j}\omega_j$.

For an integer m, an m-hierarchical erasure in $\mathbf{c}$ amounts to erasing at most m *left-justified* entries of all c_i's. That is, for every m-hierarchical erasure in $\mathbf{c}$, there exists a tuple $(t_1, \ldots, t_n)$ of nonnegative integers whose sum is at most m such that $c_{1,1}, \ldots, c_{1,t_1}, c_{2,1}, \ldots, c_{2,t_2}, \ldots,\ c_{n,1}, \ldots, c_{n,t_n}$ are replaced by $*$. For example, for $\alpha = 3$, $n = 4$, and $m = 5$, all of the following are examples of m-hierarchical erasures in a codeword $\mathbf{c} \in \mathbb{F}_{q^3}^4$:

$$\begin{aligned}
&((*, c_{1,2}, c_{1,3}), (*, *, c_{2,3}), (*, c_{3,2}, c_{3,3}), (*, c_{4,2}, c_{4,3}))\\
&((c_{1,1}, c_{1,2}, c_{1,3}), (*, *, *), (*, c_{3,2}, c_{3,3}), (*, c_{4,2}, c_{4,3}))\\
&((*, *, c_{1,3}), (c_{2,1}, c_{2,2}, c_{2,3}), (*, *, c_{3,3}), (*, c_{4,2}, c_{4,3}))\,. \qquad (1)
\end{aligned}$$

In contrast, the following is *not* a hierarchical erasure, since the erasures are not left-justified:

$$((c_{1,1}, *, c_{1,3}), (*, c_{2,2}, c_{2,3}), (*, *, c_{3,3}), (*, c_{4,2}, c_{4,3}))\,.$$

Given a basis $\boldsymbol{\omega}$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, a linear code $\mathcal{C}$ is called *an m-correcting code over $\boldsymbol{\omega}$* if it is possible to correct any m-hierarchical erasure, where codeword symbols are represented in the basis $\boldsymbol{\omega}$. The goal of this paper is, given the parameters n, m, and α, to find a basis $\boldsymbol{\omega}$ and construct a linear m-correcting code over $\boldsymbol{\omega}$, with maximum dimension k and minimum base-field size q.

For positive integers α, n, and m let

$$\mathcal{N}_{\alpha,m}^n \triangleq \left\{ (t_1, t_2, \ldots, t_n) \,\middle|\, 0 \leqslant t_i \leqslant \alpha \text{ for all } i \text{ and } \sum_{i=1}^n t_i \leqslant m \right\}.$$

---

[1] Typically, bases are considered as sets, not as vectors. In this paper however, we consider bases of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$ as (row) vectors of length α over $\mathbb{F}_{q^\alpha}$, the entries of whom span $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$.

In the special case where $\alpha = m$ we use the shorthand notation $\mathcal{N}^n_\alpha$. An element $\mathbf{t} \in \mathcal{N}^n_{\alpha,m}$ is called *an erasure pattern*, and it uniquely determines the locations of the $*$ symbols in a hierarchical erasure. For instance, the erasure patterns which appear in (1) are $(1,2,1,1)$, $(0,3,1,1)$, and $(2,0,2,1)$, respectively. For a set $\mathcal{T} \subseteq \mathcal{N}^n_{\alpha,m}$, we say that $\mathcal{C} \subseteq \mathbb{F}^n_{q^\alpha}$ is $\mathcal{T}$-correcting over $\boldsymbol{\omega}$ if all erasure patterns in $\mathcal{T}$ can be corrected. An $\mathcal{N}^n_{\alpha,m}$-correcting code is called an m-correcting code.

We make repeated use of the following notations. For an integer ℓ let $[\ell] \triangleq \{1, 2, \ldots, \ell\}$. For $\mathbf{c} \in \mathbb{F}^\ell_{q^\alpha}$ and a basis $\boldsymbol{\omega}$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$ let

$$w_{\boldsymbol{\omega}}(\mathbf{c}) \triangleq \sum_{i\in[\ell]} \max\{j \in [\alpha] \mid c_{i,j} \neq 0\},$$

where the $c_{i,j}$'s are the coefficients of the entries $\mathbf{c}$ in the representation over $\boldsymbol{\omega}$, as explained above, and the subscript $\boldsymbol{\omega}$ is omitted if clear from the context.

Finally, we note that to the best of our knowledge, this paper is the first to study linear hierarchical erasure correcting codes. Yet, similar problems have been studied in the past. The closest one is [2], in which exactly the same erasure patterns have been studied, bounds formulated, and constructions provided. However, the codes there are linear *after* having each element from $\mathbb{F}_{q^\alpha}$ expanded to its coordinate vector of length α over $\mathbb{F}_q$ in some basis $\boldsymbol{\omega}$. But when considered as a code over $\mathbb{F}_{q^\alpha}$, the code is closed under addition and multiplication only by scalars from $\mathbb{F}_q$, and not necessarily under multiplication by scalars from $\mathbb{F}_{q^\alpha}$, namely, it is not necessarily linear. Such codes are sometimes referred to as *vector-linear* codes. This work was later generalized in [3], but still under the vector-linear coding framework. In another recent work [4], the decoder does not access the entire $\mathbb{F}_{q^\alpha}$ code symbol, but unlike our paper, it is allowed to freely choose the function to extract from the symbol.

2.2. Potential Applications

Linear codes have widespread applications in coding for distributed storage systems [5]. Normally, a database $\mathbf{x} \in \mathbb{F}^k_{q^\alpha}$ is mapped to a codeword $\mathbf{c} \in \mathbb{F}^n_{q^\alpha}$, and each codeword symbol is stored on a different storage server. Then, in cases where some servers might be unavailable due to hardware failures, the reconstruction of the entire database $\mathbf{x}$ by communicating with the storage servers corresponds to (ordinary) erasure correction.

However, it has been demonstrated recently that modern distributed systems are prone to the *stragglers* phenomenon [6], which are servers that respond much slower than the average. Moreover, communicating a large amount of data from a server does not occur instantaneously, but rather as an ordered sequence of bits or packets. Therefore, it is evident that our problem is directly applicable to storage systems that employ linear codes, and suffer from the straggler phenomenon. For applications of this sort, one might be more interested in the regime $\alpha \gg n$, since the number of storage servers in the systems is likely to be much smaller than the content of each individual server.

Additional applications can be found in flash storage devices that employ *low-density parity-check* (LDPC) codes. A flash memory cell can store 2^α distinct charge levels, each representing a stored binary vector of length α. Reading the cell can be done by applying a series of $2^\alpha - 1$ threshold tests, ordered in a way that recovers the α bits one after another[2]. In the event that this series of threshold tests discontinues abruptly due to hardware failures, the missing bits from the readout value correspond to a hierarchical erasure. A common and effective approach to decoding LDPC codes consists of *variable nodes*, representing the codeword symbols, and *check nodes*, which represent a linear combination of variable nodes. Then, decoding is performed in an iterative manner, where variable nodes communicate with check nodes and vice versa [7].

Each check node represents an equation $\sum_{i=1}^{n} h_i x_i = 0$, where each $x_i \in \mathbb{F}_{q^\alpha}$

---

[2] While a single cell may be tested using only α tests using a binary-search algorithm, in a typical flash memory a threshold test is administered to a large array of cells at once. Thus, typically, some cells in the array would test below the threshold and some above. To find out the charge levels in all the cells we would typically need to test all $2^\alpha - 1$ thresholds. Nonetheless, the thresholds may be ordered to test at 1/2-range, 1/4-range, 3/4-range, and so on, making the first test obtain the most-significant bit of each cell, the following two tests to obtain the second-most-significant bit, and so on.

is a variable node representing a value contained in a flash memory cell, and the h_i's are pre-determined coefficients in $\mathbb{F}_{q^\alpha}$. It is readily verified that if the right kernel of the row vector $\mathbf{h} = (h_i)_{i=1}^n$ is an m-correcting code, one can resolve any m-hierarchical erasure in the code symbols $(x_1, \ldots, x_n)$. For applications of this sort, one might be more interested in the regime $n \gg \alpha$, since the typical number of bits stored per cell is much smaller than a useful codeword length n. Decoding of LDPC codes with m-correcting check nodes was studied in [8, 9], which served as the main inspiration for the current paper.

2.3. Universally Decodable Matrices

The problems in this paper are intimately connected to *Universally Decodable Matrices* (UDMs) [10, 11], which are a useful tool in error correction of *slow-fading channels* [12].

Definition 1 ([10, Def. 1])**.** *For $m \geqslant \alpha$, matrices $A_1, \ldots, A_n \in \mathbb{F}_q^{\alpha\times m}$ are called Universally Decodable Matrices (UDMs) if for every $\mathbf{t} = (t_1, \ldots, t_n) \in \mathcal{N}_{\alpha,m}^n$ the following condition is satisfied: the matrix composed of the first t_1 rows of A_1, the first t_2 rows of A_2, ..., the first t_n rows of A_n, has full rank.*

In the following theorem let $I_{\alpha\times m}$ be the first α rows of an $m \times m$ identity matrix. Similarly, let $J_{\alpha\times m}$ be the first α rows in the anti-identity matrix, i.e., the matrix which contains 1's in its anti-diagonal, and zero elsewhere.

Theorem 1 ([11, Prop. 14])**.** *Let $n, m,$ and α be positive integers, let q be a prime power such that $q \geqslant n - 1$, and let γ be a primitive element in $\mathbb{F}_q$. Then, the following are $\alpha \times m$ UDMs over $\mathbb{F}_q$*

$$A_0 \triangleq I_{\alpha,m},\ A_1 \triangleq J_{\alpha,m}, A_2, \ldots, A_{n-1} \text{ where}$$
$$(A_{i+1})_{a,b} = \binom{b}{a}\gamma^{(i-1)(b-a)} \text{ for } (i,a,b) \in [n-2] \times [\alpha] \times [m].$$

UDMs will be used in Subsection 3.2 to define the parity check matrix of the constructed codes. Further, in Appendix A it is shown that the important special case $\alpha = m$ is tightly connected to the existence of UDMs with a certain mutual eigenvector.

2.4. Main Lemma

Most of the results in this paper are based on the following lemma. It is stated generally for $\mathcal{T}$-correcting codes for any $\mathcal{T} \subseteq \mathcal{N}^n_{\alpha,m}$, and specifies to m-correcting code by choosing $\mathcal{T} = \mathcal{N}^n_{\alpha,m}$. For an erasure pattern $\mathbf{t} \in \mathcal{N}^n_{\alpha,m}$ and a basis $\boldsymbol{\omega}$, denote

$$\begin{aligned}\mathcal{X}_{\mathbf{t}} = \mathcal{X}_{\mathbf{t}}(\boldsymbol{\omega}) \triangleq &\left\langle \{(\omega_i, 0, \ldots, 0)\}_{i\in[t_1]} \right\rangle \oplus \\ &\left\langle \{(0, \omega_i, 0, \ldots, 0)\}_{i\in[t_2]} \right\rangle \oplus \\ &\ldots \\ &\left\langle \{(0, \ldots, 0, \omega_i)\}_{i\in[t_n]} \right\rangle,\end{aligned} \tag{2}$$

where each vector in (2) is of length n, $\langle\cdot\rangle$ denotes span over $\mathbb{F}_q$, and $\oplus$ is the sum of subspaces that intersect trivially. For example, for $n = 3$, $m = 4$, and $\mathbf{t} = (2, 1, 1) \in \mathcal{N}^3_{2,4}$ we have $\mathcal{X}_{\mathbf{t}} = \langle(\omega_1, 0, 0), (\omega_2, 0, 0), (0, \omega_1, 0), (0, 0, \omega_1)\rangle$. Note that the elements of $\mathcal{X}_{\mathbf{t}}$ are precisely the ones that are indistinguishable from the zero vector under the erasure pattern $\mathbf{t}$.

Lemma 1. *For any $\mathcal{T} \subseteq \mathcal{N}^n_{\alpha,m}$, a linear code $\mathcal{C} \subseteq \mathbb{F}^n_{q^\alpha}$ is $\mathcal{T}$-correcting over $\boldsymbol{\omega}$ if and only if $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}} = \{0\}$ for every $\mathbf{t} \in \mathcal{T}$.*

Proof. To prove one direction, assume that $\mathcal{C}$ is $\mathcal{T}$-correcting. If $\mathcal{C}$ contains a nonzero codeword which belongs to $\mathcal{X}_{\mathbf{t}}$ for some $\mathbf{t} \in \mathcal{T}$, then this codeword is indistinguishable from the zero word under the erasure pattern $\mathbf{t}$, which implies that $\mathbf{t}$ is not correctable.

Conversely, assume that $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}} = \{0\}$ for every $\mathbf{t} \in \mathcal{T}$. If $\mathcal{C}$ is not $\mathcal{T}$-correcting, it follows that there exist two distinct words

$$\begin{aligned}\mathbf{c}^{(1)} &= \left((c^{(1)}_{1,1}, \ldots, c^{(1)}_{1,\alpha}), \ldots, (c^{(1)}_{n,1}, \ldots, c^{(1)}_{n,\alpha})\right)\\ \mathbf{c}^{(2)} &= \left((c^{(2)}_{1,1}, \ldots, c^{(2)}_{1,\alpha}), \ldots, (c^{(2)}_{n,1}, \ldots, c^{(2)}_{n,\alpha})\right)\end{aligned}$$

that are indistinguishable after some erasure pattern $\mathbf{t} = (t_i)^n_{i=1} \in \mathcal{T}$. This indistinguishability implies that $c^{(1)}_{i,j} = c^{(2)}_{i,j}$ for every $(i, j) \in [n] \times ([\alpha] \setminus [t_i])$; and since the code is linear, it follows that $\mathbf{d} \triangleq \mathbf{c}^{(1)} - \mathbf{c}^{(2)}$ belongs to $\mathcal{C}$ as well. However, it is readily verified that $\mathbf{d}$ is a nonzero codeword in $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}}$, a contradiction. $\square$

2.5. Our Contribution

We begin in Subsection 3.1 with a construction for the parameters $(n, k, m) = (2, 1, \alpha)$. The well-known trace operator is used in Subsection 3.2 to construct m-correcting codes that are better suited for the regime $n \gg \alpha$.

Since extending these two constructions to other parameters proved to be difficult, in Subsection 3.3 we resort to restricted types of erasure patterns called *balanced* and the important case $k = n-1$, which generalizes the prevalent parity code. In Subsection 3.4 we discuss *power* erasure patterns, that generalize the balanced ones, and provide a code construction for $k = n - 1$ at the price of a larger base field than for balanced patterns. We conclude the constructive part of the paper in Subsection 3.5, by showing that Gabidulin codes can correct yet another restricted type of erasure patterns. The parameters for all the constructions in this paper are given in Table 1. Finally, several simple upper bounds and an existential lower bound are given in Section 4.

3. Constructions

3.1. α-correcting codes of length two

Theorem 2. *For any prime power q and any even $\alpha \in \mathbb{N}$, the code*

$$\mathcal{C} \triangleq \left\{ \mathbf{c} \in \mathbb{F}_{q^\alpha}^2 \;\middle|\; (1, b) \cdot \mathbf{c}^\intercal = 0 \right\}$$

is α-correcting, where b is a root of an irreducible quadratic polynomial over $\mathbb{F}_q$.

To prove this theorem, the following lemmas are given. In what follows, for an element $b \in \mathbb{F}_{q^\alpha}$ and an even α, a basis $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_\alpha)$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$ is called *b-symmetric* if $\omega_{\alpha-i+1} = b\omega_i$ for all $i \in [\alpha/2]$; namely, if

$$\boldsymbol{\omega} = (\omega_1, \omega_2, \ldots, \omega_{\alpha/2}, b\omega_{\alpha/2}, \ldots, b\omega_2, b\omega_1).$$

Subsection	Field	Parameters	Patterns	Tool
3.1	Any	$n = 2$ $k = 1$ $m = \alpha$ even	$\mathcal{N}^2_\alpha$	Irreducible polynomial
3.2	$q \geqslant n - 1$	$k \geqslant n - m$	$\mathcal{N}^n_{\alpha,m}$	Trace, dual bases
3.3	$q \geqslant n - 1$ $\alpha = 2^\beta$	$k = n - 1$ $m = \alpha$	$\mathcal{N}^n_{\alpha\mid\text{bal}}$	Subfield independence
3.4	$q \geqslant \frac{\alpha}{2}n + 1$ $\alpha = 2^\beta$ $\frac{\alpha}{2} \mid q - 1$	$k = n - 1$ $m = \alpha$	$\mathcal{N}^n_{\alpha\mid\text{pow}}$	Determinant
3.5	Any	$k = n - r$ $\alpha \geqslant n \geqslant r$	$\mathcal{N}^n_{r,nr}$	Gabidulin codes

Table 1: Summary of constructions.

Lemma 2. *For any even $\alpha \in \mathbb{N}$ and any prime power q, there exists a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, where $b \in \mathbb{F}_{q^\alpha}$ is a root of an irreducible quadratic polynomial $P(x)$ over $\mathbb{F}_q$.*

Proof. Denote $\alpha = 2^t \ell$, where ℓ is odd and $t \geqslant 1$. We prove this claim by induction on t. For $t = 1$ let $\omega_1, \ldots, \omega_\ell$ be a basis of $\mathbb{F}_{q^\ell}$ over $\mathbb{F}_q$. Notice that $P(x)$ remains irreducible when seen as a polynomial over $\mathbb{F}_{q^\ell}$; otherwise, we have that $P(x)$ is a minimal polynomial of some element in $\mathbb{F}_{q^\ell}$, whose degree does not divide ℓ, a contradiction. Hence, we have that $b \notin \mathbb{F}_{q^\ell}$, and thus $(\omega_1, \ldots, \omega_\ell, b\omega_\ell, \ldots, b\omega_1)$ is a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$.

For $t > 1$, by the induction hypothesis there exists a b-symmetric basis $(\omega_1, \ldots, \omega_{\alpha/2})$ of $\mathbb{F}_{q^{\alpha/2}}$ over $\mathbb{F}_q$. By choosing any $\gamma \in \mathbb{F}_{q^\alpha} \setminus \mathbb{F}_{q^{\alpha/2}}$, it is readily verified that

$$\begin{aligned}\boldsymbol{\omega} &\triangleq (\gamma\omega_1, \omega_1, \ldots, \gamma\omega_{\alpha/4}, \omega_{\alpha/4}, \omega_{\alpha/4+1}, \gamma\omega_{\alpha/4+1}, \ldots, \omega_{\alpha/2}, \gamma\omega_{\alpha/2})\\ &= (\gamma\omega_1, \omega_1, \ldots, \gamma\omega_{\alpha/4}, \omega_{\alpha/4}, b\omega_{\alpha/4}, b\gamma\omega_{\alpha/4}, \ldots, b\omega_1, b\gamma\omega_1)\end{aligned}$$

is a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, where the last equality follows from the induction hypothesis. □

Lemma 3. *If* $\boldsymbol{\omega} = (\omega_i)_{i\in[\alpha]}$ *is a* b*-symmetric basis, with* $b \in \mathbb{F}_{q^\alpha}$ *being a root of an irreducible quadratic polynomial* $P(x) = x^2 + a_1x + a_0$ *over* $\mathbb{F}_q$*, then*

$$\langle b\omega_1, b\omega_2, \ldots, b\omega_t\rangle = \langle \omega_\alpha, \omega_{\alpha-1} \ldots, \omega_{\alpha-t+1}\rangle$$

for every $t \in [\alpha]$*.*

Proof. If $t \leqslant \alpha/2$, then the claim follows from the definition of a b-symmetric basis. If $t \geqslant \alpha/2 + 1$, we have that

$$\begin{aligned}\langle b\omega_1, \ldots, b\omega_t\rangle &= \left\langle \{b\omega_i\}_{i=1}^{\alpha/2}\right\rangle + \left\langle \{b\omega_i\}_{i=\alpha/2+1}^{t}\right\rangle\\ &= \left\langle \{\omega_i\}_{i=\alpha/2+1}^{\alpha}\right\rangle + \left\langle \{b^2\omega_{\alpha-i+1}\}_{i=\alpha/2+1}^{t}\right\rangle\\ &= \left\langle \{\omega_i\}_{i=\alpha/2+1}^{\alpha}\right\rangle + \left\langle \{(-a_1b - a_0)\omega_{\alpha-i+1}\}_{i=\alpha/2+1}^{t}\right\rangle\\ &= \left\langle \{\omega_i\}_{i=\alpha/2+1}^{\alpha}\right\rangle + \left\langle \{-a_1\omega_i - a_0\omega_{\alpha-i+1}\}_{i=\alpha/2+1}^{t}\right\rangle\\ &= \left\langle \{\omega_i\}_{i=\alpha/2+1}^{\alpha}\right\rangle + \left\langle \{\omega_i\}_{i=\alpha-t+1}^{\alpha/2}\right\rangle\\ &= \langle \omega_\alpha, \omega_{\alpha-1}, \ldots, \omega_{\alpha-t+1}\rangle.\end{aligned}$$

□

Lemma 2 and Lemma 3 imply Theorem 2 as follows.

Proof. (of Theorem 2) Let $\boldsymbol{\omega}$ be a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, as guaranteed by Lemma 2. According to Lemma 1, it suffices to prove that $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}} = \{0\}$ for every $\mathbf{t} \in \mathcal{N}_\alpha^2$. Assume to the contrary that there exists $\mathbf{t} \in \mathcal{N}_\alpha^2$ and a nonzero codeword $\mathbf{c} = (c_1, c_2) \in \mathcal{C}$ such that $\mathbf{c} \in \mathcal{X}_{\mathbf{t}}(\boldsymbol{\omega})$. This readily implies that

$$c_1 \in \langle \omega_1, \ldots, \omega_{t_1}\rangle, \tag{3}$$

$$c_2 \in \langle \omega_1, \ldots, \omega_{t_2}\rangle, \text{ and} \tag{4}$$

$$c_1 + bc_2 = 0. \tag{5}$$

Furthermore, Lemma 3 and Eq. (4) imply that bc_2 is in $\langle \omega_\alpha, \omega_{\alpha-1}, \ldots, \omega_{\alpha-t_2+1}\rangle$. Since $t_1 + t_2 < \alpha + 1$, it follows that $t_1 < \alpha - t_2 + 1$, and hence (3) implies that (5) is a sum of elements from trivially intersecting subspaces that results in zero, and hence c_1 and bc_2 must both be zero. Since b is nonzero, this implies that $(c_1, c_2) = (0, 0)$, a contradiction. $\square$

Remark 1. *An alternative proof for this construction can be obtained by viewing it as a pair of UDMs with the added property that they share an eigenvector whose entries span* $\mathbb{F}_{q^\alpha}$ *over* $\mathbb{F}_q$. *More details on this view (for general* $n \geqslant 2$*) are given in Appendix A.*

3.2. m-correcting codes from traces

In this section we make use of the *trace operator* Tr [13, Def. 2.22] and *dual bases* [13, Def. 2.30]. These are well-studied notions in the theory of finite fields, and are extensively used in coding techniques for distributed storage systems (e.g., [14, 15]).

The *trace* of an element $c \in \mathbb{F}_{q^\alpha}$ (with respect to $\mathbb{F}_q$) is defined as

$$\operatorname{Tr}(c) \triangleq c + c^q + c^{q^2} + \ldots + c^{q^{\alpha-1}}.$$

The trace function is linear over $\mathbb{F}_q$, i.e., $\operatorname{Tr}(\gamma a + \delta b) = \gamma \operatorname{Tr}(a) + \delta \operatorname{Tr}(b)$ for every $\gamma, \delta \in \mathbb{F}_q$ and $a, b \in \mathbb{F}_{q^\alpha}$. Two bases $\boldsymbol{\omega} = (\omega_i)_{i=1}^{\alpha}$ and $\boldsymbol{\mu} = (\mu_i)_{i=1}^{\alpha}$ are called dual if

$$\operatorname{Tr}(\omega_i \cdot \mu_j) = \begin{cases} 0 & \text{if } i \neq j, \\ 1 & \text{if } i = j, \end{cases}$$

and for every basis there exists a unique dual basis [13, Def. 2.30].

Theorem 3. *For positive integers* $m \geqslant \alpha$, *let* $\{A_i\}_{i\in[n]}$ *be* $\alpha \times m$ *UDMs over* $\mathbb{F}_q$, *and let* $\boldsymbol{\mu}$ *be a basis of* $\mathbb{F}_{q^\alpha}$ *over* $\mathbb{F}_q$. *Then, the code*

$$\mathcal{C} \triangleq \left\{(c_1, \ldots, c_n) \in \mathbb{F}_{q^\alpha}^n \,\middle|\, (A_1^\intercal \boldsymbol{\mu}^\intercal | \cdots | A_n^\intercal \boldsymbol{\mu}^\intercal) \cdot (c_1, \ldots, c_n)^\intercal = 0\right\}$$

is m*-correcting over the dual* $\boldsymbol{\omega}$ *of* $\boldsymbol{\mu}$, *and* $\dim \mathcal{C} \geqslant n - m$.

Proof. Assume to the contrary that there exists $\mathbf{t} \in \mathcal{N}_{\alpha,m}^n$ and a nonzero codeword $\mathbf{c} \in \mathcal{C}$ such that $\mathbf{c} \in \mathcal{X}_{\mathbf{t}}(\boldsymbol{\omega})$. Therefore, any codeword symbol c_i can be written as $c_i = \sum_{j\in[t_i]} c_{i,j}\omega_j$ for some coefficients $c_{i,j} \in \mathbb{F}_q$, and hence

$$\boldsymbol{\mu}^\intercal c_i = \begin{pmatrix} \sum_{j\in[t_i]} c_{i,j}\omega_j\mu_1 \\ \sum_{j\in[t_i]} c_{i,j}\omega_j\mu_2 \\ \vdots \\ \sum_{j\in[t_i]} c_{i,j}\omega_j\mu_\alpha \end{pmatrix}.$$

Thus, for every $\ell \in [m]$, the ℓ'th entry of the equation $\sum_{i\in[n]} A_i^\intercal \boldsymbol{\mu}^\intercal c_i = 0$ equals

$$\sum_{i\in[n]}\sum_{r\in[\alpha]} A_i^{(r,\ell)} \sum_{j\in[t_i]} c_{i,j}\omega_j \cdot \mu_r = 0,$$

where $A_i^{(r,\ell)}$ is the (r,ℓ)'th entry of A_i. Applying the trace function on both sides, and exploiting the linearity of the trace and the fact that $\boldsymbol{\omega}$ and $\boldsymbol{\mu}$ are dual, yields

$$\sum_{i\in[n]}\sum_{r\in[t_i]} A_i^{(r,\ell)} c_{i,r} = 0 \text{ for every } \ell \in [m].$$

In turn, this implies that the vector $(c_{1,1},\ldots,c_{1,t_1},\cdots,c_{n,1},\ldots,c_{n,t_n})$ is in the left kernel of

$$\begin{pmatrix} A_1^{(1:t_1)} \\ A_2^{(1:t_2)} \\ \vdots \\ A_n^{(1:t_n)} \end{pmatrix},$$

where $A_i^{(1:t_i)}$ is a matrix which contains the top t_i rows of A_i, which contradicts the definition of UDMs. The bound $\dim \mathcal{C} \geqslant n-m$ follows since $\mathcal{C}$ is the right kernel of an $m \times n$ matrix. □

In light of the bound $\dim \mathcal{C} \geqslant n-m$ that is given above, one might prefer to employ this construction in the regime $n \gg \alpha$. However, for the case of even $m = \alpha = n$, one can guarantee $\dim \mathcal{C} > 0$ by using techniques from Subsection 3.1. The proof is given in Appendix B.

Corollary 1. *For even $m = \alpha = n \in \mathbb{N}$, let $\{A_i\}_{i=1}^n$ be $\alpha \times \alpha$ UDMs such that A_1 is the identity matrix, and*

$$A_2 = \begin{pmatrix} & & & & & 1 \\ & & & & \iddots & \\ & & & 1 & & \\ & & -a_0 & -a_1 & & \\ & \iddots & & & \ddots & \\ -a_0 & & & & & -a_1 \end{pmatrix},$$

where $x^2 + a_1 x + a_0$ is an irreducible quadratic polynomial over $\mathbb{F}_q$ with a root $b \in \mathbb{F}_{q^\alpha}$. In addition, let $\boldsymbol{\mu}$ be a b-symmetric basis (see Lemma 2), and let $\boldsymbol{\omega}$ be its dual. Then, the code

$$\mathcal{C} \triangleq \left\{ (c_1, \ldots, c_n) \in \mathbb{F}_{q^\alpha}^n \,\middle|\, (A_1^\intercal \boldsymbol{\mu}^\intercal | \cdots | A_n^\intercal \boldsymbol{\mu}^\intercal) \cdot (c_1, \ldots, c_n)^\intercal = 0 \right\}$$

is an α-correcting code over $\boldsymbol{\omega}$ with $\dim \mathcal{C} \geqslant 1$.

3.3. Correcting balanced erasure patterns

The case $k = n-1$ and $m = \alpha$ is of particular importance, since it generalizes the widely used parity code (for storage applications), and corresponds to hierarchical erasure correction in check nodes of LDPC codes (see Subsection 2.2). This case is not handled well by previous subsections; in Subsection 3.1 it necessitates $n = 2$ (i.e., a short code), and in Subsection 3.2 one must have $m = 1$ (i.e., low erasure correction) to get $k = n - 1$. Hence, in this subsection we focus on this case, and show a code construction which protects against erasure patterns that we call *balanced*. This case is also addressed in Subsection 3.4 which follows, where a stronger erasure correction is guaranteed at the price of a larger base field, by using similar techniques.

Assume that $\alpha = 2^\beta$ for some integer β. An erasure pattern $\mathbf{t} \in \mathcal{N}_\alpha^n$ is called balanced if there exists an integer $0 \leqslant i \leqslant \min\{\beta, \log n\}$ (where the logarithm is to base 2) and a set $J \subseteq [n]$ with $|J| \leqslant 2^i$, such that for all $j \in [n]$,

$$\begin{cases} t_j \leqslant \frac{\alpha}{2^i} & \text{if } j \in J\text{; and} \\ t_j = 0 & \text{otherwise.} \end{cases}$$

For example, for $n = 4$ the erasure patterns

$$(\alpha/2, 0, \alpha/2, 0), \text{ and}$$
$$(\alpha/4, \alpha/4, \alpha/4, \alpha/4)$$

are balanced, whereas $(\alpha/2, \alpha/4, \alpha/4, 0)$ is not. The set of all balanced erasure patterns is denoted by $\mathcal{N}^n_{\alpha|\mathrm{bal}}$.

We consider bases $\boldsymbol{\omega} = (\omega_1, \ldots, \omega_\alpha)$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$ that we call *recursive*, i.e., bases such that $\langle \omega_1, \ldots, \omega_{\alpha/2^i} \rangle = \mathbb{F}_{q^{\alpha/2^i}}$ for all $0 \leqslant i \leqslant \beta$. For a vector $\mathbf{h} = (h_1, \ldots, h_n) \in \mathbb{F}^n_{q^\alpha}$ we define a code

$$\mathcal{C} = \mathcal{C}(\mathbf{h}) \triangleq \ker(\mathbf{h}) \triangleq \left\{ \mathbf{c} \in \mathbb{F}^n_{q^\alpha} \,\middle|\, \mathbf{h}\mathbf{c}^\intercal = 0 \right\}. \tag{6}$$

The ability of the code $\mathcal{C}$ to protect against balanced erasure patterns reduces to linear independence of some subsets of the h_i's over certain subfields of $\mathbb{F}_{q^\alpha}$, as we now show.

Lemma 4. *The code $\mathcal{C}$ (6) is $\mathcal{N}^n_{\alpha|\mathrm{bal}}$-correcting over a recursive basis $\boldsymbol{\omega}$ if and only if for every $1 \leqslant i \leqslant \min\{\beta, \log n\}$, every 2^i-subset of $\{h_j\}_{j\in[n]}$ is a linearly independent set over $\mathbb{F}_{q^{\alpha/2^i}}$.*

Proof. Assume that every 2^i-subset of $\{h_j\}_{j=1}^n$ is linearly independent over $\mathbb{F}_{q^{\alpha/2^i}}$ for every $0 \leqslant i \leqslant \min\{\beta, \log n\}$. According to Lemma 1, if $\mathcal{C}$ is not $\mathcal{N}^n_{\alpha|\mathrm{bal}}$-correcting, then there exists a nonzero $\mathbf{c} = (c_1, c_2, \ldots, c_n)$ in $\mathcal{C}$ and an erasure pattern $\mathbf{t} \in \mathcal{N}^n_{\alpha|\mathrm{bal}}$ such that $\mathbf{c} \in \mathcal{C} \cap \mathcal{X}_{\mathbf{t}}$. By the definition of $\mathcal{N}^n_{\alpha|\mathrm{bal}}$, it follows that there exists an integer i and a set $J \subseteq [n]$ of size at most 2^i such that $t_j \leqslant \alpha/2^i$ if $j \in J$, and $t_j = 0$ otherwise. Hence, we have that

$$c_j \in \langle \omega_1, \ldots, \omega_{\alpha/2^i} \rangle = \mathbb{F}_{q^{\alpha/2^i}} \text{ for all } j \in J,$$

which implies that $\sum_{j\in J} h_j c_j = 0$. However, this sum is a linear combination of a 2^i-subset of $\{h_j\}_{j\in[n]}$ over $\mathbb{F}_{q^{\alpha/2^i}}$, a contradiction. The proof of the inverse direction is similar. $\square$

In what follows we construct an $[n, n-1]_{q^\alpha}$ $\mathcal{N}^n_{\alpha|\text{bal}}$-correcting code, for any n and any α over a base field $\mathbb{F}_q$ with $q \geqslant n-1$. To this end, recall that $\alpha = 2^\beta$, and let $\{b_i\}_{i\in[\beta]} \subseteq \mathbb{F}_{q^\alpha}$ such that

$$\mathbb{F}_{q^{\alpha/2^{i-1}}} = \mathbb{F}_{q^{\alpha/2^i}}(b_i), \tag{7}$$

for all $i \in [\beta]$, i.e., we consider each subfield $\mathbb{F}_{q^{\alpha/2^{i-1}}}$ as a vector space of dimension two over $\mathbb{F}_{q^{\alpha/2^i}}$ by fixing the basis $\{1, b_i\}$.

For $0 \leqslant i \leqslant \beta$ and a $2^i \times n$ matrix M over $\mathbb{F}_{q^{\alpha/2^i}}$, let

$$\mathcal{H}_i(M) \triangleq \text{UH}(M) + b_i \text{LH}(M),$$

where UH and LH denote the upper half and lower half of M, respectively. Further, for an integer $1 \leqslant i \leqslant \beta$ and an $\alpha \times n$ matrix M over $\mathbb{F}_q$ let

$$\begin{aligned}
&\mathcal{H}^{(i)}(M) \triangleq \mathcal{H}_{\beta-i+1}(\cdots(\mathcal{H}_{\beta-1}(\mathcal{H}_\beta(M)))),\\
&\mathcal{H}^{(0)}(M) \triangleq M.
\end{aligned}$$

Throughout the remainder of this section we use a recursive basis induced by the $\{b_i\}_{i\in[\beta]}$ from (7). Namely, the basis is

$$\boldsymbol{\omega} \triangleq W_\beta, \text{ where } \quad W_0 \triangleq (1), \text{ and } \quad W_{i+1} \triangleq W_i | (b_{\beta-i} \cdot W_i), \tag{8}$$

and $|$ denotes concatenation. Alternatively,

$$\boldsymbol{\omega} \triangleq (1, b_1) \otimes (1, b_2) \otimes \cdots \otimes (1, b_\beta),$$

where $\otimes$ denotes the Kronecker product.

Finally, recall that a *Vandermonde* matrix defined by $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_n) \in \mathbb{F}_q^n$ is a matrix whose (i,j)'th entry equals ν_j^{i-1}. We say that a matrix V is a *generalized Vandermonde* (GV) matrix defined by $\boldsymbol{\nu}$ if $V = M \cdot \text{diag}(\mathbf{d})$ for some Vandermonde matrix M defined by $\boldsymbol{\nu}$ and some vector $\mathbf{d} = (d_1, \ldots, d_n)$

with nonzero entries. Note that a GV matrix $V \in \mathbb{F}_q^{r\times s}$ for some integers $s \geqslant r$, which is defined by s distinct field elements, is also an MDS matrix, i.e., all its $r \times r$ submatrices are invertible.

Theorem 4. *For an integer $\alpha = 2^\beta$ and an integer n, let q be a prime power such that $q \geqslant n$, and let $V \in \mathbb{F}_q^{\alpha\times n}$ be a Vandermonde matrix defined by distinct n elements. Then, for $\mathbf{h} = (h_1, h_2, \ldots, h_n) \triangleq \mathcal{H}^{(\beta)}(V)$, the code $\mathcal{C} \triangleq \ker(\mathbf{h})$ is a $\mathcal{N}^n_{\alpha|\text{bal}}$-correcting code over the basis $\boldsymbol{\omega}$ of (8).*

The proof of this theorem requires the following lemma.

Lemma 5. *Let $\alpha = 2^\beta$ and let V be an $\alpha \times n$ GV matrix defined by $\boldsymbol{\nu} = (\nu_1, \ldots, \nu_n) \in \mathbb{F}_q^n$. Then for all $0 \leqslant i \leqslant \beta$, the matrix $\mathcal{H}^{(i)}(V)$ is a GV matrix over $\mathbb{F}_{q^{2^i}}$ also defined by $\boldsymbol{\nu}$.*

Proof. We prove this claim by induction, in which the base case $i = 0$ is clear. For $i \geqslant 1$, assume that $V_i \triangleq \mathcal{H}^{(i)}(V) \in \mathbb{F}_{q^{2^i}}^{(\alpha/2^i)\times n}$ is a GV matrix, and let U_i and L_i be its upper and lower halves, respectively. Since V_i is a GV matrix, there exists a Vandermonde matrix $M \in \mathbb{F}_{q^{2^i}}^{(\alpha/2^i)\times n}$ defined by $\boldsymbol{\nu}$ and a vector $\mathbf{d} \in (\mathbb{F}^*_{q^{2^i}})^n$ such that $V_i = M \operatorname{diag}(\mathbf{d})$. Hence, it follows that $U_i = \text{UH}(M) \operatorname{diag}(\mathbf{d})$ and $L_i = \text{LH}(M) \operatorname{diag}(\mathbf{d})$, and therefore

$$\begin{aligned}
V_{i+i} &= \mathcal{H}^{(i+1)}(V) = \mathcal{H}_{\beta-i}(V_i) \\
&= U_i + b_{\beta-i} L_i \\
&= \text{UH}(M) \operatorname{diag}(\mathbf{d}) + b_{\beta-i}\text{LH}(M) \operatorname{diag}(\mathbf{d}).
\end{aligned}$$

Now, since M is a Vandermonde matrix, it is readily verified that $\text{LH}(M) = \text{UH}(M) \operatorname{diag}(\mathbf{x})$ for some $\mathbf{x} = (x_1, \ldots, x_n) \in (\mathbb{F}^*_{q^{2^i}})^n$, and thus

$$\begin{aligned}
V_{i+i} &= \text{UH}(M) \operatorname{diag}(\mathbf{d}) + b_{\beta-i}\text{UH}(M) \operatorname{diag}(\mathbf{x}) \operatorname{diag}(\mathbf{d}) \\
&= \text{UH}(M) \left(\operatorname{diag}(\mathbf{d}) + b_{\beta-i} \operatorname{diag}(\mathbf{x}) \operatorname{diag}(\mathbf{d})\right) \\
&= \text{UH}(M) \operatorname{diag}((\mathbb{1} + b_{\beta-i}\mathbf{x}) \odot \operatorname{diag}(\mathbf{d})),
\end{aligned}$$

where $\odot$ denotes the pointwise product of vectors (also called the *Hadamard product*), and $\mathbb{1}$ is the all 1's vector. Since $\text{UH}(M)$ is a Vandermonde matrix defined by $\boldsymbol{\nu}$, to finish the proof it suffices to show that the entries of $(\mathbb{1}+b_{\beta-i}\mathbf{x})\odot \operatorname{diag}(\mathbf{d})$ are nonzero. Assuming otherwise, it follows that $(1+b_{\beta-i}x_j)d_j = 0$ for some $j \in [n]$; and since $d_j \neq 0$ and $x_j \neq 0$, we have that $b_{\beta-i} = -x_j^{-1}$. However, $-x_j^{-1} \in \mathbb{F}_{q^{2^i}}$ and $b_{\beta-i} \notin \mathbb{F}_{q^{\frac{\alpha}{2^{\beta-i}}}} = \mathbb{F}_{q^{2^i}}$, a contradiction. $\square$

We are now ready to prove Theorem 4.

Proof. (of Theorem 4) According to Lemma 4, it suffices to show that for any $1 \leqslant i \leqslant \min\{\log n, \beta\}$, any 2^i-subset of $\{h_j\}_{j\in[n]}$ is linearly independent over $\mathbb{F}_{q^{\alpha/2^i}}$. For any such i, let $J \subseteq [n]$ be a subset of size 2^i, and let $H_J \in \mathbb{F}_{q^{\alpha/2^i}}^{2^i\times 2^i}$ be the matrix whose columns are the representations of all elements in $\{h_j\}_{j\in J}$ over the (ordered) basis W_i. Notice that $\{h_j\}_{j\in J}$ is a linearly independent set over $\mathbb{F}_{q^{\alpha/2^i}}$ if and only if H_J is invertible. However, H_J is a $2^i \times 2^i$ submatrix of $\mathcal{H}^{(\beta-i)}(V) \in \mathbb{F}_{q^{\alpha/2^i}}^{2^i\times n}$, which is a GV matrix defined by distinct elements according to Lemma 5, and hence also an MDS matrix. Thus, H_J is invertible, and the claim follows. $\square$

Remark 2. *According to Theorem 4 it follows that*

$$h_j = \prod_{i=1}^{\beta}\left(1+b_i a_j^{\alpha/2^i}\right) \text{ for all } j \in [n],$$

where $a_1,\ldots,a_n$ are the distinct $\mathbb{F}_q$-elements in the underlying Vandermonde matrix V.

Remark 3. *The above construction is closely related to a classical coding theoretic notion called* alternant codes *[16, Sec. 5.5]. An $[n,k]_q$* Generalized Reed-Solomon *(GRS) code is a linear code whose parity check matrix is an $(n-k)\times n$ GV matrix over $\mathbb{F}_q$. An alternant code $\mathcal{C}_{\text{alt}}$ is defined as $\mathcal{C}\cap F^n$, where $\mathcal{C}$ is an $[n,k]_q$ GRS code and F is a subfield of $\mathbb{F}_q$. Let $\alpha < n$, and for any $0 \leqslant i \leqslant \beta$ let $\mathcal{C}_i$ be the right kernel of $\mathcal{H}^{(i)}(V)$ over $\mathbb{F}_{q^{2^i}}$. Notice that Lemma 5 shows that $\mathcal{C}_i$ is an $[n, n-\alpha/2^i]_{q^{2^i}}$ GRS code. Furthermore, it is readily verified that $\mathcal{C}_j$ is an alternant code of $\mathcal{C}_i$ whenever $j \leqslant i$. Lemma 5 also implies that the codes we construct here have the property that all the alternant codes in the hierarchy* *are of maximum distance, and in cases where q is prime, these are all possible alternant codes.*

3.4. Correcting power erasure patterns

We generalize the results of the previous section by considering a larger family of erasure patterns, $\mathcal{N}^n_{\alpha|\mathrm{pow}}$, that includes balanced patterns, i.e., $\mathcal{N}^n_{\alpha|\mathrm{bal}} \subseteq \mathcal{N}^n_{\alpha|\mathrm{pow}}$. As before, let $\alpha = 2^\beta$ for some positive integer β. An erasure pattern $\mathbf{t} \in \mathcal{N}^n_\alpha$ is called a *power erasure pattern* if there exists $J \subseteq [n]$ such that

$$t_j = \begin{cases} \frac{\alpha}{2^{m_j}} & j \in J, \\ 0 & \text{otherwise}, \end{cases}$$

where $0 \leqslant m_j \leqslant \beta$ is an integer for all $j \in J$, and $\sum_{j\in J} 2^{-m_j} = 1$. Thus, for example, when $n = 4$, $(\alpha/2, \alpha/4, \alpha/4, 0)$ is a power erasure pattern but is not a balanced erasure pattern.

Theorem 5. *For an integer $\alpha = 2^\beta$, and an integer n, let q be a prime power such that $\frac{\alpha}{2}|q-1$. Let $\nu_1, \ldots, \nu_n \in \mathbb{F}_q$ be arbitrary non-zero scalars such that $\nu_j^{\alpha/2} \neq \nu_k^{\alpha/2}$ for all $j \neq k$. Let $V \in \mathbb{F}_q^{\alpha\times n}$ be a Vandermonde matrix defined by $(\nu_1, \ldots, \nu_n)$. Then, for $\mathbf{h} = (h_1, h_2, \ldots, h_n) \triangleq \mathcal{H}^{(\beta)}(V)$, the code $\mathcal{C} \triangleq \ker(\mathbf{h})$ is an $\mathcal{N}^n_{\alpha|\mathrm{pow}}$-correcting code over the basis $\boldsymbol{\omega}$ of (8).*

We shall require the following natural extension of Lemma 4.

Lemma 6. *The code $\mathcal{C}$ of (6) is $\mathcal{N}^n_{\alpha|\mathrm{pow}}$-correcting over a recursive basis $\boldsymbol{\omega}$ if and only if for every power erasure pattern $\mathbf{t} \in \mathcal{N}^n_{\alpha|\mathrm{pow}}$ (defined by the sets J and $\{m_j\}_{j\in J}$) the equation*

$$\sum_{j\in J} h_j c_j = 0,$$

has only the trivial solution when $c_j \in \mathbb{F}_{q^{\alpha/2^{m_j}}}$ for every $j \in J$.

Proof. If $\mathcal{C}$ is not $\mathcal{N}^n_{\alpha|\mathrm{pow}}$-correcting, then there exists a nonzero $\mathbf{c} = (c_1, c_2, \ldots, c_n)$ in $\mathcal{C}$ and a power erasure pattern $\mathbf{t} \in \mathcal{N}^n_{\alpha|\mathrm{pow}}$ such that $\mathbf{c} \in \mathcal{C} \cap \mathcal{X}_\mathbf{t}$. By the definition of $\mathcal{N}^n_{\alpha|\mathrm{pow}}$, it follows that there exist corresponding sets J and $\{m_j\}_{j\in J}$.

Hence, we have that

$$c_j \in \left\langle \omega_1, \dots, \omega_{\alpha/2^{m_j}} \right\rangle = \mathbb{F}_{q^{\alpha/2^{m_j}}} \text{ for all } j \in J,$$

as well as $\sum_{j \in J} h_j c_j = 0$, thus proving one direction of the claim. The proof of the inverse direction is similar. □

We now give the proof of Theorem 5.

Proof. (of Theorem 5) Let $\mathbf{t} \in \mathcal{N}^n_{\alpha|\mathrm{pow}}$ be a power erasure pattern, with corresponding sets J and $\{m_j\}_{j\in J}$. By applying Lemma 6 our goal is now to prove a solution to $\sum_{j\in J} h_j c_j = 0$ with $c_j \in \mathbb{F}_{q^{\alpha/2^{m_j}}}$ must be a trivial all-zero solution.

Let us denote by $\mathbf{v}_j^\intercal$, $j \in [n]$, the jth column of the Vandermonde matrix V. Additionally, recall the recursive basis $\boldsymbol{\omega} \triangleq W_\beta$ from (8). Thus, $\mathbf{v}_j^\intercal$ contains the coordinates (over $\mathbb{F}_q$) of h_j when using the basis $\boldsymbol{\omega}$.

If we define $\overline{\mathbf{v}}_j \triangleq (1, \nu_j, \dots, \nu_j^{\alpha/2^{m_j}-1})$ then

$$\mathbf{v}_j^\intercal = \begin{pmatrix} \overline{\mathbf{v}}_j^\intercal \\ \nu_j^{\alpha/2^{m_j}} \overline{\mathbf{v}}_j^\intercal \\ \vdots \\ \nu_j^{(2^{m_j}-1)\alpha/2^{m_j}} \overline{\mathbf{v}}_j^\intercal \end{pmatrix}.$$

Similarly, we define

$$\overline{\boldsymbol{\omega}}_j \triangleq W_{\beta - m_j} = (1, b_{m_j+1}) \otimes (1, b_{m_j+2}) \otimes \cdots \otimes (1, b_\beta),$$

which is the $\alpha/2^{m_j}$-prefix of $\boldsymbol{\omega}$. By the construction of the recursive basis $\boldsymbol{\omega}$ we have that $\overline{\boldsymbol{\omega}}_j$ is a basis for $\mathbb{F}_{q^{\alpha/2^{m_j}}}$. We now notice that

$$\begin{pmatrix} \overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \\ \nu_j^{\alpha/2^{m_j}} \overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \\ \vdots \\ \nu_j^{(2^{m_j}-1)\alpha/2^{m_j}} \overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \end{pmatrix},$$

is the coordinate vector of h_j when $\mathbb{F}_{q^\alpha}$ is viewed as a vector space over $\mathbb{F}_{q^{\alpha/2^{m_j}}}$ using the ordered basis

$$\hat{\boldsymbol{\omega}}_j \triangleq (1, b_1) \otimes (1, b_2) \otimes \cdots \otimes (1, b_{m_j}).$$

By rewriting $c_j = \sum_{i=1}^{\alpha/2^{m_j}} c_{j,i}\omega_i$, with $c_{j,i} \in \mathbb{F}_q$, our goal is equivalent to proving the set $\bigcup_{j\in J}\{h_j\omega_1, \dots, h_j\omega_{\alpha/2^{m_j}}\}$ is linearly independent over $\mathbb{F}_q$. For each $j \in J$, and for each $i \in [\alpha/2^{m_j}]$, we may write a column vector of the coordinates of $h_j\omega_i$ in $\mathbb{F}_{q^{\alpha/2^{m_j}}}$ using the basis $\hat{\boldsymbol{\omega}}$ as

$$\begin{pmatrix} \omega_i\overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \\ \nu_j^{\alpha/2^{m_j}}\omega_i\overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \\ \vdots \\ \nu_j^{(2^{m_j}-1)\alpha/2^{m_j}}\omega_i\overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}_j^\intercal \end{pmatrix},$$

where we note that both ω_i and $\overline{\boldsymbol{\omega}}_j \cdot \overline{\mathbf{v}}^\intercal$ are in $\mathbb{F}_{q^{\alpha/2^{m_j}}}$, and $\nu_j \in \mathbb{F}_q$. Now, viewing $\mathbb{F}_{q^{\alpha/2^{m_j}}}$ as a vector space over $\mathbb{F}_q$ using the basis $\overline{\boldsymbol{\omega}}_j$, multiplication by ω_i may be represented as a multiplication of the coordinates by $C_{j,i}$, an $\alpha/2^{m_j} \times \alpha/2^{m_j}$ matrix over $\mathbb{F}_q$ ($C_{i,j}$ can be made explicit using companion matrices, but this is immaterial to the rest of the proof). Thus, the coordinates of $h_j\omega_i$ over $\mathbb{F}_q$ using the basis $\boldsymbol{\omega}$ take on the simple form of

$$\mathbf{z}_{j,i}^\intercal \triangleq \begin{pmatrix} C_{j,i} & & & \\ & C_{j,i} & & \\ & & \ddots & \\ & & & C_{j,i} \end{pmatrix} \cdot \mathbf{v}_j^\intercal = \begin{pmatrix} C_{j,i}\overline{\mathbf{v}}_j^\intercal \\ \nu_j^{\alpha/2^{m_j}} C_{j,i}\overline{\mathbf{v}}_j^\intercal \\ \vdots \\ \nu_j^{(2^{m_j}-1)\alpha/2^{m_j}} C_{j,i}\overline{\mathbf{v}}_j^\intercal \end{pmatrix}$$

If we define the matrix $Z \in \mathbb{F}_q^{\alpha\times\alpha}$ to have as its columns $\{\mathbf{z}_{j,i}^\intercal\}$, $j \in J$, $i \in [\alpha/2^{m_j}]$, then it now suffices to prove $\det(Z) \neq 0$. Our strategy now is, for each $j \in J$, to take the $\alpha/2^{m_j}$ columns $\{\mathbf{z}_{j,i}^\intercal\}_{i\in[\alpha/2^{m_j}]}$ and replace them by using invertible column operations. The overall resulting matrix Z' will be shown to have $\det(Z') \neq 0$, implying $\det(Z) \neq 0$.

Fix any $j \in J$. Obviously the set $\{h_j\omega_i\}_{i\in[\alpha/2^{m_j}]}$ is linearly independent over $\mathbb{F}_q$ since $\{\omega_i\}_{i\in[\alpha/2^{m_j}]}$ is, and therefore also $\{\mathbf{z}_{j,i}^\intercal\}_{i\in[\alpha/2^{m_j}]}$. We now contend that this implies that the set $\{C_{j,i}\overline{\mathbf{v}}_j^\intercal\}_{i\in[\alpha/2^{m_j}]}$ is linearly independent over $\mathbb{F}_q$. Assuming to the contrary it is not, there exist $d_1, \dots, d_{\alpha/2^{m_j}} \in \mathbb{F}_q$, not all zero, such that $\sum_{i\in[\alpha/2^{m_j}]} d_i C_{j,i}\overline{\mathbf{v}}_j^\intercal = 0$, but then $\sum_{i\in[\alpha/2^{m_j}]} d_i \nu_j^{\ell/2^{m_j}} C_{j,i}\overline{\mathbf{v}}_j^\intercal = 0$ for any integer ℓ, implying $\sum_{i\in[\alpha/2^{m_j}]} \mathbf{z}_{j,i}^\intercal = 0$, a contradiction.

Let $\xi_j \in \mathbb{F}_q$ be an element of multiplicative order $o(\xi_j) = \alpha/2^{m_j}$, the existence of which is guaranteed by the requirement $\frac{\alpha}{2}|q-1$. Since we established that $\{C_{j,i}\overline{\mathbf{v}}_j^\intercal\}_{i\in[\alpha/2^{m_j}]}$ is linearly independent over $\mathbb{F}_q$, by invertible column operations we may map

$$\left(C_{j,1}\overline{\mathbf{v}}_j^\intercal \,\middle|\, C_{j,2}\overline{\mathbf{v}}_j^\intercal \,\middle|\, \dots \,\middle|\, C_{j,\alpha/2^{m_j}}\overline{\mathbf{v}}_j^\intercal\right)$$

$$\longmapsto \begin{pmatrix} 1 & 1 & \dots & 1 \\ \nu_j & \xi_j\nu_j & \dots & \xi_j^{\alpha/2^{m_j}-1}\nu_j \\ \nu_j^2 & (\xi_j\nu_j)^2 & \dots & (\xi_j^{\alpha/2^{m_j}-1}\nu_j)^2 \\ \vdots & \vdots & \ddots & \vdots \\ \nu_j^{\alpha/2^{m_j}-1} & (\xi_j\nu_j)^{\alpha/2^{m_j}-1} & \dots & (\xi_j^{\alpha/2^{m_j}-1}\nu_j)^{\alpha/2^{m_j}-1} \end{pmatrix},$$

i.e., the square Vandermonde matrix defined by $(\nu_j, \xi_j\nu_j, \xi_j^2\nu_j, \dots, \xi_j^{\alpha/2^{m_j}-1}\nu_j)$, which we denote by V_j for convenience. Using the same column operations on $\{\mathbf{z}_{j,i}^\intercal\}_{i\in[\alpha/2^{m_j}]}$ the mapping becomes

$$\left(\mathbf{z}_{j,1}^\intercal \,\middle|\, \mathbf{z}_{j,2}^\intercal \,\middle|\, \dots \,\middle|\, \mathbf{z}_{j,\alpha/2^{m_j}}^\intercal\right) \mapsto \begin{pmatrix} V_j \\ \nu_j^{\alpha/2^{m_j}} V_j \\ \vdots \\ \nu_j^{(2^{m_j}-1)\alpha/2^{m_j}} V_j \end{pmatrix}$$

$$= \begin{pmatrix} 1 & 1 & \dots & 1 \\ \nu_j & \xi_j\nu_j & \dots & \xi_j^{\alpha/2^{m_j}-1}\nu_j \\ \nu_j^2 & (\xi_j\nu_j)^2 & \dots & (\xi_j^{\alpha/2^{m_j}-1}\nu_j)^2 \\ \vdots & \vdots & \ddots & \vdots \\ \nu_j^{\alpha-1} & (\xi_j\nu_j)^{\alpha-1} & \dots & (\xi_j^{\alpha/2^{m_j}-1}\nu_j)^{\alpha-1} \end{pmatrix},$$

which is an $\alpha \times (\alpha/2^{m_j})$ Vandermonde matrix.

We repeat the above process for each $j \in J$ to obtain the matrix Z' which satisfies $\det(Z') = \xi\det(Z)$ for some $\xi \in \mathbb{F}_q$, $\xi \neq 0$, since only invertible column operations were used. Finally, we note that Z' is itself a Vandermonde matrix that is defined by (the multiset) $\bigcup_{j\in J}\{\xi_j^{i-1}\nu_j\}_{i\in[\alpha/2^{m_j}]}$ (in some order), and since $\nu_j^{\alpha/2} \neq \nu_k^{\alpha/2}$ for all $j \neq k$, we have $\det(Z') \neq 0$, as desired. $\square$

As a final note, we observe the field size requirements imposed by Theorem 5. We need to choose n distinct non-zero values from $\mathbb{F}_q$. However, each choice precludes some other elements from being chosen. More specifically, let $\xi \in \mathbb{F}_q$ be an element with multiplicative order $\frac{\alpha}{2}$, and let $\langle \xi \rangle$ be the multiplicative group spanned by it. Then we may choose at most one element from each of the cosets in $\mathbb{F}_q^*/\langle \xi \rangle$. Hence, $q \geqslant \frac{\alpha}{2}n + 1$.

3.5. Correcting bounded erasure patterns

In this subsection it is shown that Gabidulin codes, a well-known family of rank-metric codes, are capable of protecting against a large family of erasure patterns. In particular, for $\alpha \geqslant n$ and an integer $r \leqslant n$, the code $\mathrm{Gab}[n, n-r]_{q^\alpha}$, defined below, can protect against $\mathcal{T}_r \triangleq \mathcal{N}_{r,nr}^n = \{0, 1, \ldots, r\}^n$. Notice that $\mathcal{T}_r$ does not include full erasures of codeword symbols (unless the code is trivial), and yet Gabidulin codes can protect against erasures in the usual sense (see [17]).

For the next theorem, recall that a linearized polynomial is a polynomial over $\mathbb{F}_{q^\alpha}$ in which all nonzero coefficients correspond to monomials of the form x^{q^i} for some nonnegative integer i. For a linearized polynomial f, let its q-degree be $\deg_q(f) \triangleq \log_q(\deg f)$. It is widely known that any function from $\mathbb{F}_{q^\alpha}$ to itself, which is linear over $\mathbb{F}_q$, corresponds to a linearized polynomial. The following theorem applies over any basis $\boldsymbol{\omega}$.

Theorem 6. *For nonnegative integers r, n, and α such that $n \leqslant \alpha$ and $r < n$, the code*

$$\mathrm{Gab}[n, n-r]_{q^\alpha} \triangleq \left\{ (f(\omega_1), \ldots, f(\omega_n)) \,\middle|\, f \textit{ is linearized and } \deg_q(f) < n - r \right\}$$

is $\mathcal{T}_r$-correcting.

Proof. We show that $\mathrm{Gab}[n, n-r] \cap \mathcal{X}_{\mathbf{t}} = \{0\}$ for all $\mathbf{t} \in \mathcal{T}_r$. Assuming otherwise, we have a pattern $\mathbf{t} \in \mathcal{T}_r$ and a nonzero linearized polynomial f of q-degree less than $n - r$ such that

$$f(\omega_j) \in \left\langle \omega_1, \ldots, \omega_{t_j} \right\rangle, \text{ for all } j \in [n]. \tag{9}$$

Since f is a linearized polynomial and since $\mathbf{t} \in \mathcal{T}_r$, Eq. (9) implies that $f(\langle \omega_1, \ldots, \omega_n \rangle) \subseteq \langle \omega_1, \ldots, \omega_r \rangle$, which in turn implies that $\dim \ker(f) \geqslant n - r$. Thus, f has more roots than its degree, which is a contradiction. $\square$

Note that $n \leqslant \alpha$ is necessary, since the evaluation points $\omega_1, \ldots, \omega_n$ must be linearly independent over $\mathbb{F}_q$. Finally, we emphasize that this construction applies to any q.

4. Lower Bound

First, it is clear that any m-correcting code $\mathcal{C} \subseteq \mathbb{F}_{q^\alpha}^n$ can correct $m' \triangleq \lfloor m/\alpha \rfloor$ erasures in the usual sense. Therefore, the well-known Singleton bound implies that $m' \leqslant n - k$. Moreover, in cases where $m' = n - k$, namely, when $\mathcal{C}$ is an MDS code, the MDS conjecture (e.g., see [18], and its resolution in certain cases [19, 20]) implies $q^\alpha \geqslant n - 1$. In the remainder of this section a Gilbert-Varshamov type argument is used to prove the following existence theorem.

Theorem 7. *For all positive integers n, m, α, and r such that $m < \alpha(r-1)$, if*

$$q > \left((m+1) \binom{m+n-2}{n-2} \right)^{\frac{1}{\alpha(r-1)-m}}$$

then there exists an $[n, n-r]_{q^\alpha}$ m-correcting code $\mathcal{C}$.

Before proving the theorem, we prove an auxiliary claim, which applies for any basis $\boldsymbol{\omega}$. We say that a matrix over $\mathbb{F}_{q^\alpha}$ is *m-good* (good, in short) if its right kernel does not contain nonzero vectors $\mathbf{x}$ with $w(\mathbf{x}) \leqslant m$. In the proof of Theorem 7 we choose the columns of the parity-check matrix of the code one after another, while showing that there always exists an eligible column to add; the question of column eligibility boils down to the following lemma.

Lemma 7. *If $H_\ell \triangleq (\mathbf{g}_1^\intercal, \ldots, \mathbf{g}_\ell^\intercal) \in \mathbb{F}_{q^\alpha}^{r \times \ell}$ is good and*

$$\mathbf{g}_{\ell+1}^\intercal \notin \left\{ \gamma \cdot \sum_{i=1}^{\ell} x_i \mathbf{g}_i^\intercal \;\middle|\; \gamma \in \mathbb{F}_{q^\alpha} \text{ and } w(x_1, \ldots, x_\ell) \leqslant m \right\} \triangleq R_\ell \qquad (10)$$

then $H_{\ell+1} \triangleq (\mathbf{g}_1^\intercal, \ldots, \mathbf{g}_{\ell+1}^\intercal) \in \mathbb{F}_{q^\alpha}^{r \times (\ell+1)}$ is good.

Proof. Assume to the contrary that the right kernel of $H_{\ell+1}$ contains a nonzero vector $\mathbf{x} = (x_1, \ldots, x_{\ell+1}) \in \mathbb{F}_{q^\alpha}^{\ell+1}$ such that $w(\mathbf{x}) \leqslant m$, which implies that $-x_{\ell+1}\mathbf{g}_{\ell+1}^\intercal = \sum_{i=1}^{\ell} x_i \mathbf{g}_i^\intercal$ and that $w(x_1, \ldots, x_\ell) \leqslant m$. If $x_{\ell+1} = 0$, it follows that the vector $\mathbf{x}' \triangleq (x_1, \ldots, x_\ell)$ satisfies $H_\ell \mathbf{x}' = 0$ and $w(\mathbf{x}') \leqslant m$, in contradiction to H_ℓ being good. Otherwise, we have that $\mathbf{g}_{\ell+1}^\intercal = (-x_{\ell+1}^{-1}) \cdot \sum_{i=1}^{\ell} x_i \mathbf{g}_i^\intercal$, and hence $\mathbf{g}_{\ell+1}^\intercal \in R_\ell$ in contradiction with (10). □

The following two properties are easy to prove.

Lemma 8. *For the sets R_ℓ from* (10)*,*

1. $|R_{n-1}| \geqslant |R_\ell|$ *for all* $\ell \leqslant n-1$.
2. $|R_{n-1}| \leqslant q^\alpha \sum_{i=0}^{m} q^i \binom{i+n-2}{n-2} \leqslant (m+1) q^{\alpha+m} \binom{m+n-2}{n-2}$.

Proof. The first property is due to simple monotonicity. For the second property we upper bound the size of the set by assuming that all the linear combinations in the definition of the set are distinct. Then, we have q^α ways of choosing γ. Finally, the number of vectors $\mathbf{x} \in \mathbb{F}_{q^\alpha}^{n-1}$ with $w(\mathbf{x}) \leqslant m$ may be found using a standard balls-into-bins argument to be $\sum_{i=0}^{m} q^i \binom{i+n-2}{n-2}$. Since $q^i \binom{i+n-2}{n-2}$ is increasing in i we obtain the final inequality. □

We are now in a position to prove Theorem 7.

Proof. (of Theorem 7) We construct the parity check matrix of the code $\mathcal{C}$ column by column, starting from an $r \times r$ identity matrix. Clearly, it suffices to guarantee that all along this construction, the resulting matrices are good; this would guarantee that $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}} = \{0\}$ for every $\mathbf{t} \in \mathcal{N}_{\alpha,m}^{n}$, and thus that $\mathcal{C}$ is m-correcting by Lemma 1.

Assume that $H_\ell \in \mathbb{F}_{q^\alpha}^{r \times \ell}$ is good for some $\ell \geqslant r$ (for $\ell = r$ the goodness is satisfied since there are no nonzero vectors in the kernel). According to Lemma 7 and the above observations, it follows that if $|\mathbb{F}_{q^\alpha}^r| - |R_{n-1}| > 0$, then there exists a legitimate choice for the added column $\mathbf{g}_{\ell+1}^\intercal$. Hence, by the bound on $|R_{n-1}|$ from Lemma 8 we have

$$|\mathbb{F}_{q^\alpha}^r| - |R_{n-1}| \geqslant q^{\alpha r} - (m+1) q^{\alpha+m} \binom{m+n-2}{n-2}.$$

If that is strictly larger than zero, the desired code exists. Thus, it suffices to require

$$\begin{aligned} q^{\alpha r-\alpha-m} &> (m+1)\binom{m+n-2}{n-2} \\ q &> \left((m+1)\binom{m+n-2}{n-2}\right)^{\frac{1}{\alpha(r-1)-m}}. \end{aligned} \qquad \square$$

In the remainder of this section the bound on q in Theorem 7 is analyzed asymptotically in the two regimes of interest (see Subsection 2.2). In both regimes we focus on the practically important case where the dimension k (and hence r) is proportional to n, and the erasure correction capability m is proportional to αn; this corresponds to erasure correction of a constant fraction of the information symbols.

In the case $\alpha \gg n$ the parameter n is seen as constant and the parameter α tends to infinity. Say that $m = c_1\alpha$ and $\alpha(r-1) - m = c_2\alpha$ for some constants c_1, c_2, and then the condition on q from Theorem 7 becomes

$$q > \left((c_1\alpha+1)\binom{c_1\alpha+n-2}{n-2}\right)^{\frac{1}{c_2\alpha}} = \mathrm{poly}(\alpha)^{\frac{1}{\Theta(\alpha)}} \xrightarrow{\alpha\to\infty} 1.$$

In the case $n \gg \alpha$ we view α as constant and n as tending to infinity. Say that $m = c_1 n$ and $\alpha(r-1) - m = c_2 n$ for some c_1, c_2. By the well known approximation of the binomial coefficient (e.g., see [18, Lemma 7, p. 309]), the condition on q from Theorem 7 becomes

$$\begin{aligned} q &> \left((c_1 n+1)\binom{(1+c_1)n-2}{n-2}\right)^{\frac{1}{c_2 n}} \\ &= \left(2^{(1+c_1)nH\left(\frac{1}{1+c_1}\right)(1+o(1))}\right)^{\frac{1}{c_2 n}} \xrightarrow{n\to\infty} 2^{\frac{1+c_1}{c_2}H\left(\frac{1}{1+c_1}\right)}, \end{aligned}$$

where $H(x) \triangleq -x\log_2(x) - (1-x)\log_2(1-x)$ is the binary entropy function.

Acknowledgments

Moshe Schwartz was supported in part by a German Israeli Project Cooperation (DIP) grant under grant no. PE2398/1-1. Yuval Cassuto was supported in part by grants from the Israel Science Foundation and from the US-Israel Binational Science Foundation.

References

[1] N. Raviv, Y. Cassuto, R. Cohen, M. Schwartz, Erasure correction of scalar codes in the presence of stragglers, in: Proceedings of the 2018 IEEE International Symposium on Information Theory (ISIT2018), Vail, CO, USA, 2018, pp. 1983–1987.

[2] M. Y. Rosenbloom, M. A. Tsfasman, Codes for the m-metric, Problems of Information Transmission 33 (1) (1997) 45–52.

[3] A. Barg, W. Park, On linear ordered codes, Moscow Mathematical Journal 15 (4) (2015) 679–702.

[4] I. Tamo, M. Ye, A. Barg, Fractional decoding: error correction from partial information, in: Proceedings of the 2017 IEEE International Symposium on Information Theory (ISIT2017), Aachen, Germany, 2017, pp. 998–1002.

[5] J. S. Plank, T1: erasure codes for storage applications, in: Proc. of the 4th USENIX Conference on File and Storage Technologies, 2005, pp. 1–74.

[6] J. Dean, L. A. Barroso, The tail at scale, Communications of the ACM 56 (2) (2013) 74–80.

[7] R. G. Gallager, Low-Density Parity-Check Codes, Cambridge, MA: MIT Press, 1963.

[8] R. Cohen, Y. Cassuto, Iterative decoding of LDPC codes over the q-ary partial erasure channel, IEEE Trans. Inform. Theory 62 (5) (2016) 2658–2672.

[9] R. Cohen, N. Raviv, Y. Cassuto, LDPC codes over the q-ary multi-bit channel, IEEE Trans. Inform. Theory 65 (7) (2019) 4293–4306.

[10] A. Ganesan, P. O. Vontobel, On the existence of universally decodable matrices, IEEE Trans. Inform. Theory 53 (7) (2007) 2572–2575.

[11] P. O. Vontobel, A. Ganesan, On universally decodable matrices for space-time coding, Designs, Codes and Cryptography 41 (3) (2006) 325–342.

[12] S. Tavildar, P. Viswanath, Approximately universal codes over slow-fading channels, IEEE Trans. Inform. Theory 52 (7) (2006) 3233–3258.

[13] R. Lidl, H. Niederreiter, Finite Fields, Cambridge University Press, 1997.

[14] V. Guruswami, M. Wootters, Repairing reed-solomon codes, IEEE transactions on Information Theory 63 (9) (2017) 5684–5698.

[15] L. Jin, Y. Luo, C. Xing, Repairing algebraic geometry codes, IEEE Transactions on Information Theory 64 (2) (2017) 900–908.

[16] R. M. Roth, Introduction to Coding Theory, Cambridge Univ. Press, 2006.

[17] G. Richter, S. Plass, Fast decoding of rank-codes with rank errors and column erasures, in: Proceedings of the 2004 International Symposium on Information Theory (ISIT), IEEE, 2004, pp. 398–398.

[18] F. J. MacWilliams, N. J. A. Sloane, The Theory of Error-Correcting Codes, North-Holland, 1978.

[19] S. Ball, On sets of vectors of a finite vector space in which every subset of basis size is a basis, J. Eur. Math. Soc. 14 (3) (2012) 733–748.

[20] S. Ball, J. De Beule, On sets of vectors of a finite vector space in which every subset of basis size is a basis II, Designs, Codes and Cryptography 65 (1–2) (2012) 5–14.

Appendix A. α-correcting codes from mutual eigenvector of UDMs

For the case $m = \alpha$, there exists an intriguing connection between UDMs and α-correcting codes.

Theorem 8. *For $h_1, \ldots, h_n \in \mathbb{F}_{q^\alpha}$, a code $\mathcal{C} = \left\{ \mathbf{c} \in \mathbb{F}_{q^\alpha}^n \,\middle|\, (h_1, \ldots, h_n) \cdot \mathbf{c}^\intercal = 0 \right\}$ is an α-correcting code over an ordered basis $\boldsymbol{\omega} \triangleq (\omega_1, \ldots, \omega_\alpha)$ if and only if there exists a set $A_1, \ldots, A_n$ of UDMs over $\mathbb{F}_q$ such that for any $i \in [n]$, the element h_i is an eigenvalue of A_i with a corresponding eigenvector $\boldsymbol{\omega}^\intercal$.*

Proof. Let $A_1, \ldots, A_n \in \mathbb{F}_q^{\alpha \times \alpha}$ be UDMs with eigenvalues $h_1, \ldots, h_n \in \mathbb{F}_{q^\alpha}$, respectively, all of which correspond to the eigenvector $\boldsymbol{\omega}$, i.e.,

$$A_i \boldsymbol{\omega}^\intercal = h_i \boldsymbol{\omega}^\intercal \text{ for all } i \in [n]. \tag{A.1}$$

If $\mathcal{C}$ is not α-correcting, it follows that there exist $\mathbf{t} \in \mathcal{N}_\alpha^n$ and a nonzero codeword $\mathbf{c} = (c_1, c_2, \ldots, c_n) \in \mathcal{C}$ such that $c_i \in \langle \omega_1, \ldots, \omega_{t_i} \rangle$ for all $i \in [n]$, and therefore

$$h_i c_i \in \langle h_i \omega_1, \ldots, h_i \omega_{t_i} \rangle \overset{\text{(A.1)}}{=} \left\langle A_i^{(1)} \boldsymbol{\omega}^\intercal, \ldots, A_i^{(t_i)} \boldsymbol{\omega}^\intercal \right\rangle,$$

where $A_i^{(j)}$ denotes the j-th row of A_i. In turn, this implies that for all $i \in [n]$ there exists a nonzero vector $\mathbf{v}_i \in \mathbb{F}_q^{t_i}$ such that $\mathbf{v}_i A_i^{(1:t_i)} \boldsymbol{\omega}^\intercal = h_i c_i$, where for any positive integers r and s, the notation $A_i^{(s:r)}$ stands for the submatrix of A_i which consists of rows s through r. Thus, we have a nonzero vector $\mathbf{v} \triangleq (\mathbf{v}_1 | \mathbf{v}_2 | \ldots | \mathbf{v}_n) \in \mathbb{F}_q^\alpha$ that satisfies

$$\mathbf{v} \cdot \begin{pmatrix} A_1^{(1:t_1)} \\ A_2^{(1:t_2)} \\ \vdots \\ A_n^{(1:t_n)} \end{pmatrix} \cdot \boldsymbol{\omega}^\intercal = \sum_{i \in [n]} \mathbf{v}_i A_i^{(1:t_i)} \boldsymbol{\omega}^\intercal = \sum_{i \in [n]} h_i c_i = 0. \tag{A.2}$$

Now, since the entries of $\boldsymbol{\omega}$ are a basis, and since the A_i's and the $\mathbf{v}_i$'s are over $\mathbb{F}_q$, the expression $(\sum_{i \in [n]} \mathbf{v}_i A_i^{(1:t_1)}) \boldsymbol{\omega}^\intercal = 0$ implies that the vector $\sum_{i \in [n]} \mathbf{v}_i A_i^{(1:t_1)}$ is the zero vector. However, this implies that there exists a nonzero vector $\mathbf{v}$ in the left kernel of a matrix which consists of upper rows of UDMs, a contradiction.

Conversely, assume that $\mathcal{C}$ is α-correcting, and define matrices $A_1, \ldots, A_n \in \mathbb{F}_q^{\alpha \times \alpha}$ as follows. For every $i \in [n]$, let A_i be the matrix such that $A_i^{(j)}$ is the expansion of $h_i \omega_j$ over the basis $\boldsymbol{\omega}$, i.e., $h_i \omega_j = \sum_{\ell=1}^{\alpha} (A_i^{(j)})_\ell \omega_\ell$. Assuming to the contrary that $A_1, \ldots, A_n$ are not UDMs, we have an element $\mathbf{t} = (t_1, \ldots, t_n) \in \mathcal{N}_\alpha^n$ and a nonzero vector $\mathbf{v} \in \mathbb{F}_q^\alpha$ such that

$$\mathbf{v} \cdot \begin{pmatrix} A_1^{(1:t_1)} \\ A_2^{(1:t_2)} \\ \vdots \\ A_n^{(1:t_n)} \end{pmatrix} = 0.$$

Partition $\mathbf{v}$ to n consecutive parts $\mathbf{v}_1, \mathbf{v}_2, \ldots, \mathbf{v}_n$ of sizes $t_1, \ldots, t_n$, respectively, let $c_i \triangleq \mathbf{v}_i \cdot (\omega_1, \ldots, \omega_{t_i})^\intercal$ for all $i \in [n]$, and let $\mathbf{c} \triangleq (c_1, \ldots, c_n)$. Notice that $\mathbf{c} \in \mathcal{C}$, since:

$$\begin{aligned}
(h_1, \ldots, h_n)\mathbf{c}^\intercal &= \sum_{i=1}^n h_i \mathbf{v}_i (\omega_1, \ldots, \omega_{t_i})^\intercal = \sum_{i=1}^n \mathbf{v}_i (h_i\omega_1, \ldots, h_i\omega_{t_i})^\intercal \\
&= \sum_{i=1}^n \mathbf{v}_i \left(\sum_{\ell=1}^\alpha (A_i^{(1)})_\ell \omega_\ell, \ldots, \sum_{\ell=1}^\alpha (A_i^{(t_i)})_\ell \omega_\ell \right)^\intercal \\
&= \sum_{i=1}^n \mathbf{v}_i A_i^{(1:t_i)} \boldsymbol{\omega}^\intercal \\
&= \mathbf{v} \cdot \begin{pmatrix} A_1^{(1:t_1)} \\ A_2^{(1:t_2)} \\ \vdots \\ A_n^{(1:t_n)} \end{pmatrix} \cdot \boldsymbol{\omega}^\intercal = 0.
\end{aligned}$$

Moreover, since $\mathbf{c} \in \mathcal{X}$ by definition, it follows that $\mathbf{c}$ is a nonzero codeword in $\mathcal{C} \cap \mathcal{X}_{\mathbf{t}}$, a contradiction to $\mathcal{C}$ being an α-correcting code. □

Finally, we note that Theorem 2 can alternatively be proved by a direct application of Theorem 8, and the details are left to the curious reader.

Appendix B. An omitted proof

Proof. (of Corollary 1). First, we ought to show that such UDMs exist. Indeed, according to [11, Lemma 4], it follows that for any UDMs $\{B_i\}_{i=1}^n$ and any lower-triangular invertible matrices $\{C_i\}_{i=1}^n$, the matrices $\{A_i = C_i B_i\}_{i=1}^n$ are UDMs as well. The existence of suitable UDMs for our proof is then proved by letting $\{A_i\}_{i=1}^n$ be, say, the UDMs from Theorem 1 for the parameters at hand, letting C_i be an identity matrix for every $i \in [n] \setminus \{2\}$, and

$$C_2 = \begin{pmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & -a_1 & -a_0 & & \\ & \iddots & & & \ddots & \\ -a_1 & & & & & -a_0 \end{pmatrix}.$$

Now, observe that $\boldsymbol{\mu}^\intercal$ is an eigenvector for the eigenvalue 1 of A_1, and an eigenvector for the eigenvalue b of A_2 (see Appendix A for further implications of such mutual eigenvectors). Therefore, the square parity check matrix $(A_1^\intercal \boldsymbol{\mu}^\intercal | \cdots | A_n^\intercal \boldsymbol{\mu}^\intercal)$ has at least two dependent columns, which implies that $\dim \mathcal{C} \geqslant 1$. $\square$