Erasure Correction of Scalar Codes in the Presence of Stragglers

Netanel Raviv[⋆], **Yuval Cassuto**[†], **Rami Cohen**[†], and **Moshe Schwartz**[‡]
[⋆]Department of Electrical Engineering, California Institute of Technology, Pasadena 91125, CA, USA
[†]Faculty of Electrical Engineering, Technion - Israel Institute of Technology, Haifa 3200003, Israel
[‡]Electrical and Computer Engineering, Ben-Gurion University of the Negev, Beer Sheva 8410501, Israel
netanel.raviv@gmail.com, ycassuto@ee.technion.ac.il, rc@technion.ac.il, schwartz@ee.bgu.ac.il

***Abstract*—Recent advances in coding for distributed storage systems have reignited the interest in scalar codes over extension fields. In parallel, the rise of large-scale distributed systems has motivated the study of computing in the presence of *stragglers*, i.e., servers that are slow to respond or unavailable.**

This paper addresses storage systems that employ linear codes over extension fields. A common task in such systems is the reconstruction of the entire dataset using sequential symbol transmissions from multiple servers, which are received concurrently at a central data collector. However, a key bottleneck in the reconstruction process is the possible presence of stragglers, which may result in excessive latency. To mitigate the straggler effect, the reconstruction should be possible given *any* sufficiently large set of sequentially received symbols, regardless of their source. In what follows, an algebraic framework for this scenario is given, and a number of explicit constructions are provided. Our main result is a construction that uses a recursive composition of generalized Reed-Solomon codes over smaller fields. In addition, we show links of this problem to Gabidulin codes and to universally decodable matrices.

I. Introduction

Scalar codes over extension fields are prevalent in many real-world distributed storage systems as a means to combat node failures. However, even when no node failures occur, distributed systems are prone to the detrimental effect of *stragglers*, i.e., nodes that are slow to respond. In this paper, we provide a mathematical modeling for the straggler effect in systems which employ scalar codes, and suggest several code constructions to mitigate this effect.

For a prime power q let $\mathbb{F}_q$ be a field with q elements, and for an integer α let $\mathbb{F}_{q^\alpha}$ be its algebraic extension of degree α. Assume that a dataset $x \in \mathbb{F}_{q^\alpha}^k$ is coded across n storage nodes by using a (scalar) linear code over $\mathbb{F}_{q^\alpha}$, where each node contains a symbol in $\mathbb{F}_{q^\alpha}$.

Upon the reconstruction of x, each node transmits its content to the data collector *sequentially*, i.e., one q-ary symbol after the other[1]. The data collector receives these transmissions concurrently, and wishes to reconstruct the dataset as soon as possible without having to wait for all transmissions to complete. Note that due to the stochastic nature of the system and the nodes' workload, the amount of downloadable bits from each node is unknown a priori neither to the nodes nor to the data collector. Our objective in this paper is to construct scalar codes that minimize this reconstruction time.

To this end, consider an *ordered* basis $\boldsymbol{\omega} \triangleq (\omega_1, \ldots, \omega_\alpha)$ of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, where each element of $\mathbb{F}_{q^\alpha}$ is identified by the vector of coefficients in its representation by $\boldsymbol{\omega}$. Consequently, for any $i \in [n] \triangleq \{1, \ldots, n\}$, the transmission of a codeword symbol $c_i = \sum_{j=1}^{\alpha} c_{i,j}\omega_j \in \mathbb{F}_{q^\alpha}$ to the data collector is assumed to contain the q-ary coefficients $c_{i,j}$, which arrive free of errors and in a decreasing order of their j index. Notice that the sequential arrival of the information bits translates to *left-justified* erasures of the coefficients in the representation of every codeword symbol by the fixed basis $\boldsymbol{\omega}$; that is, some (possibly empty) prefix of the coefficient vector is missing from the response of every server.

Clearly, elementary information-theoretic arguments imply that at least αk symbols over $\mathbb{F}_q$ must be obtained at the data collector for a unique reconstruction to be possible. For a given positive integer m, we are interested in constructing linear codes, of maximum possible dimension, such that x can be reconstructed from *any* set of $\alpha n - m$ symbols over $\mathbb{F}_q$ that arrive sequentially at the data collector.

This problem may be seen as a generalization of erasure correction in the usual sense. For example, note that when the erased symbols are confined to $n - k$ servers and $m = \alpha(n - k)$, then this problem reduces to ordinary $(n - k)$-erasure correction; and clearly, this can be obtained by employing any $[n, k]_{q^\alpha}$ MDS code (i.e., a code of length n and dimension k over $\mathbb{F}_{q^\alpha}$, which attains the Singleton bound with equality). Furthermore, it is also possible to use an $[\alpha n, \alpha n - m]_q$ MDS code, which can correct *any* m-erasures (i.e., not necessarily sequential). However, this approach does not yield linear codes over $\mathbb{F}_{q^\alpha}$, and requires a fairly large base field size.

This paper focuses on linear $[n, n - 1]_{q^\alpha}$ codes that are capable of alleviating the straggler effect (formal definitions are given in the sequel). In cases where a code construction remains elusive, we suggest a relaxation of the problem by restricting the set of possible *patterns* of missing symbols. Yet, we are able to prove that codes in the stronger model exist for all parameters n and α, as long as q is sufficiently large.

This paper is structured as follows. Previous work regarding the use of scalar codes for distributed storage systems is discussed in Section II, and formal definitions are given in Section III. Code constructions are given in Section IV, which includes short codes for any erasure pattern, and long codes for more restricted types of patterns. The existence of α-correcting codes over large fields is proved in Section V. Connections to GRS codes, Gabidulin codes, and universally decodable matrices are discussed in Section VI, and open problems are discussed in Section VII.

The work of Netanel Raviv was supported in part by the postdoctoral fellowship of the Center for the Mathematics of Information (CMI), Caltech, and in part by the Lester-Deutsch postdoctoral fellowship. This research was supported in part by the Israel Science Foundation under grant no. 130/14, 1676/14 and the US-Israel Binational Science Foundation.

[1]In the underlying system q-ary symbols may be generally mapped to bits.

978-1-5386-4780-6/18/$31.00©2018 IEEE

II. Previous work

Scalar codes for distributed storage, which were studied in a series of recent works [1], [2], [5], [7], [9], [13], often outperform their vector counterparts by utilizing the extension field structure to speed up encoding and decoding.[2]

Yet, all of the aforementioned works have focused on *node repair* (i.e., reconstruction of one or more codeword symbols in the case of node failure), rather than on reconstructing the entire dataset, and did not address the straggler effect.

The paper [10] addresses the reconstruction of an entire dataset that is stored by a scalar MDS code. In [10], the number of downloadable field elements from each node is assumed to be known prior to transmission, and hence, each node computes a function with that many output symbols. In our paper however, the nodes compose their transmission without any knowledge regarding its arrival, other than it being sequential.

Hierarchical erasures (also known as partial erasures) also emerge in the study of flash memories [3], [4]. A flash memory cell can store 2^α distinct charge levels, that are read by applying a series of α threshold tests. It often happens that this series of threshold tests discontinues abruptly due to hardware failures, a scenario which clearly correspond to a hierarchical erasure. The work of [3], [4] studies LDPC coding under these erasures, analyzes the optimal edge-label distributions and decoding threshold, and discusses the finite length case. In particular, these works have laid the grounds for the algebraic ideas that are developed hereafter.

Finally, it should be noted that *fountain codes* address a similar scenario, in which packets arrive in an unordered manner. However, the emphasis in fountain codes is given on probabilistic constructions and large packets, whereas the present paper discusses worst-case constructions and small packets.

III. Preliminaries

Informally speaking, a linear code $\mathcal{C} \subseteq \mathbb{F}_{q^\alpha}^n$ is called an *m-hierarchical erasure correcting code over $\boldsymbol{\omega}$* if the data collector is capable of reconstructing every codeword $\mathbf{c} = (c_1, \ldots, c_n) \in \mathcal{C}$ from any set of prefixes of the representations of the c_i's according to $\boldsymbol{\omega}$, which contains at least $\alpha n - m$ symbols of $\mathbb{F}_q$ overall.

To define this notion formally, for positive integers n and m let $\mathcal{N}_m^n$ be the set of nonnegative integer vectors of length n whose sum of entries is m. An element $\mathbf{t} \in \mathcal{N}_m^n$ describes the number of missing q-ary symbols from each of the n nodes, and hence it is called an *erasure pattern*. For such $\mathbf{t} \in \mathcal{N}_m^n$, let

$$\begin{aligned} X_{\mathbf{t}}(\boldsymbol{\omega}) \triangleq & \left\langle \{(\omega_i, 0, \ldots, 0)\}_{i\in[t_1]} \right\rangle \oplus \\ & \left\langle \{(0, \omega_i, 0, \ldots, 0)\}_{i\in[t_2]} \right\rangle \oplus \\ & \cdots \\ & \left\langle \{(0, \ldots, 0, \omega_i)\}_{i\in[t_n]} \right\rangle, \end{aligned} \tag{1}$$

where each vector represented by $(\cdot)$ is an n-vector, $\langle\cdot\rangle$ denotes span over $\mathbb{F}_q$, and $\oplus$ is the sum of subspaces that intersect trivially. For example, for $n = 3$, $m = 4$, and $\mathbf{t} = (2, 1, 1) \in \mathcal{N}_4^3$ we have $X_{\mathbf{t}}(\boldsymbol{\omega}) = \langle (\omega_1, 0, 0), (\omega_2, 0, 0), (0, \omega_1, 0), (0, 0, \omega_1) \rangle$. Note that each element of $X_t(\boldsymbol{\omega})$ corresponds to an $\mathbb{F}_{q^\alpha}^n$ word that is mapped to the all-0 word by an m-hierarchical erasure.

Definition 1. *A code $\mathcal{C} \subseteq \mathbb{F}_{q^\alpha}^n$ is called an m-hierarchical erasure correcting code over $\boldsymbol{\omega}$ if $\mathcal{C} \cap X_{\mathbf{t}}(\boldsymbol{\omega}) = \{0\}$ for all $\mathbf{t} \in \mathcal{N}_m^n$. In addition, $\mathcal{C}$ is called m-correcting if there exists a basis $\boldsymbol{\omega}$ over which $\mathcal{C}$ is an m-hierarchical erasure correcting code.*

Similarly, for $T \subseteq \mathcal{N}_m^n$ a code $\mathcal{C} \subseteq \mathbb{F}_{q^\alpha}^n$ is T-correcting if there exists a basis $\boldsymbol{\omega}$ such that $\mathcal{C} \cap X_{\mathbf{t}}(\boldsymbol{\omega}) = \{0\}$ for all $\mathbf{t} \in T$. Notice that if $m = \ell\alpha$ for some positive integer ℓ, then an m-correcting code is also an ℓ-erasure correcting code in the ordinary sense, and hence its dimension is upper bounded by $n - \ell$. In this paper we mostly focus on α-correcting codes and T correcting codes for $T \subseteq \mathcal{N}_\alpha^n$.

IV. Code constructions

A. α-correcting codes of length two

In this section we provide a construction of an α-correcting code of length two for any q and any even α.

Theorem 1. *For any q and any even α, the code*

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

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, an ordered 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 $i \in [\alpha/2]$. Further, let $b \in \mathbb{F}_{q^\alpha}$ be a root of an irreducible quadratic polynomial $P(x) = x^2 + a_1 x + a_0$ over $\mathbb{F}_q$.

Lemma 2. *For any even α and any q, there exists a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$.*

Proof. Denote $\alpha = 2^t \ell$, where ℓ is odd and $t \geq 1$. 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 we have 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

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

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 b-symmetric, 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 all $t \in [\alpha]$.

[2] For example, computing $A \cdot v$ for $A \in \mathbb{F}_q^{\alpha\times\alpha}$ and $v \in \mathbb{F}_q^\alpha$ requires $\Theta(\alpha^2)$ operations over $\mathbb{F}_q$, whereas if A is representative of an element $a \in \mathbb{F}_{q^\alpha}$ then $a \cdot v$ can be computed in $\Theta(\alpha \log \alpha)$ field operations.

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

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

Lemma 2 and Lemma 3 imply the correctness of Theorem 1 as follows.

Proof. (of Theorem 1) Let $\boldsymbol{\omega}$ be a b-symmetric basis of $\mathbb{F}_{q^\alpha}$ over $\mathbb{F}_q$, as guaranteed by Lemma 2. Assume for contradiction 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 $\mathbf{c} \in X_{\mathbf{t}}(\boldsymbol{\omega})$. This readily implies that

$$c_1 \in \langle \omega_1, \ldots, \omega_{t_1} \rangle, \tag{2}$$

$$c_2 \in \langle \omega_1, \ldots, \omega_{t_2} \rangle, \text{ and} \tag{3}$$

$$c_1 + bc_2 = 0. \tag{4}$$

Furthermore, Lemma 3 and Eq. (3) imply that bc_2 is in $\langle \omega_\alpha, \omega_{\alpha-1}, \ldots, \omega_{\alpha-t_2+1} \rangle$. Alongside (2), this implies that (4) is a sum of elements from trivially intersecting subspaces that results in zero, and hence c_1 and c_2 are zero as well, a contradiction. $\square$

This $[2,1]_{q^\alpha}$ code is applicable for correcting partial erasures in flash-memory LDPC codes, by assigning the edge labels 1 and b to a check node of degree two [4]. Further, in the full version of this paper it will be extended to an $[n,1]_{q^\alpha}$ α-correcting code for any even $\alpha \geq n-1$ and any $q \geq n-1$.

B. Balanced erasure patterns

In this subsection, we restrict our attention to $\alpha = 2^\beta$ for some positive integer[3] β and to *balanced* erasure patterns. An erasure pattern $\mathbf{t} \in \mathcal{N}_\alpha^n$ is called balanced if there exists an integer $0 \leq i \leq \min\{\beta, \log n\}$, where the logarithm is to base 2, and a set $J \subseteq [n]$ with $|J| = 2^i$, such that for all $j \in [n]$

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

For example, if $n = 4$ then the erasure patterns $(\alpha/2, 0, \alpha/2, 0)$ 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 T_B.

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 $0 \leq i \leq \beta$. For a vector $\mathbf{h} = (h_1, \ldots, h_n) \in \mathbb{F}_{q^\alpha}^n$ and a code $\mathcal{C} \triangleq \ker_{\mathbb{F}_{q^\alpha}}(\mathbf{h})$, imposing these restrictions reduces the construction problem to one of linear independence of subsets of the h_i's over certain subfields of $\mathbb{F}_{q^\alpha}$, as we now show.

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

[3]In fact, any α of the form $\ell 2^\beta$ can be treated similarly, by applying the same techniques on $(q')^{\alpha'}$, where $q' \triangleq q^\ell$ and $\alpha' \triangleq 2^\beta$.

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 $0 \leq i \leq \min\{\beta, \log n\}$. If $\mathcal{C}$ is not T_B-correcting, then there exists a nonzero $\mathbf{c} = (c_1, c_2, \ldots, c_n)$ in $\mathcal{C}$ and an erasure pattern $\mathbf{t} \in T_B$ such that $\mathbf{c} \in \mathcal{C} \cap X_{\mathbf{t}}$. By the definition of T_B, it follows that there exists an integer i and a set $J \subseteq [n]$ of size 2^i such that $t_j = \alpha/2^i$ if $j \in J$, and zero otherwise. Hence, we have that

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

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 $[n, n-1]_{q^\alpha}$ T_B-correcting code, for any n and any α over a base field $\mathbb{F}_q$ with $q \geq n-1$. To this end, let $\{b_i\}_{i\in[\beta]} \subseteq \mathbb{F}_{q^\alpha}$ such that $\mathbb{F}_{q^{\alpha/2^{i-1}}} = \mathbb{F}_{q^{\alpha/2^i}}(b_i)$ 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 $0 \leq i \leq \beta$ and a $2^i \times n$ matrix M over $\mathbb{F}_{q^{\alpha/2^i}}$, let

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

where UH and LH denote the upper and lower halves of M, respectively. Further, for an integer $1 \leq i \leq \beta$ and an $\alpha \times n$ matrix M over $\mathbb{F}_q$ let $\mathcal{H}^{(i)}(M) \triangleq \mathcal{H}_{\beta-i+1}(\cdots(\mathcal{H}_{\beta-1}(\mathcal{H}_\beta(M))))$, and let $\mathcal{H}^{(0)}(M) = M$.

Finally, we say that a matrix V is a *generalized Vandermonde* (GV) matrix if $V = M \cdot \mathrm{diag}(\mathbf{d})$ for some Vandermonde matrix M 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 $s \geq r$ is also an MDS matrix, i.e., all its $r \times r$ submatrices are invertible.

Theorem 5. *For an integer α which is a power of two and an integer n, let q be a prime power such that $q \geq n-1$, and let $V \in \mathbb{F}_q^{\alpha\times n}$ be a Vandermonde matrix. Then, for $\mathbf{h} = (h_1, h_2, \ldots, h_n) \triangleq \mathcal{H}^{(\beta)}(V)$, the code $\mathcal{C} \triangleq \ker_{\mathbb{F}_{q^\alpha}}(\mathbf{h})$ is a T_B-correcting code.*

The proof of this theorem requires the following lemma.

Lemma 6. *For all $1 \leq i \leq \beta$, the matrix $\mathcal{H}^{(i)}(V)$ is a GV matrix over $\mathbb{F}_{q^{2^i}}$.*

Proof. We prove this claim by induction, in which the base case $i = 0$ is clear. 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 lower and upper halves, respectively. Since V_i is a GV matrix, there exists a Vandermonde matrix $M \in \mathbb{F}_{q^{2^i}}^{(\alpha/2^i)\times n}$ and a vector $\mathbf{d} \in (\mathbb{F}_{q^{2^i}}^*)^n$ such that $V_i = M\,\mathrm{diag}(\mathbf{d})$. Hence, it follows that $U_i = \mathrm{UH}(M)\,\mathrm{diag}(\mathbf{d})$ and $L_i = \mathrm{LH}(M)\,\mathrm{diag}(\mathbf{d})$, and therefore

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

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

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

where $\odot$ denotes the pointwise product of vectors (also called *Hadamard product*), and $\mathbb{1}$ is the all 1's vector. Since $\mathrm{UH}(M)$ is a Vandermonde matrix, to finish the proof we are only left to show the entries of $(\mathbb{1} + b_{\beta-i}\mathbf{x}) \odot \mathrm{diag}(\mathbf{d})$ 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. □

The recursive basis which is used to prove Theorem 5 is the one which is induced by the $\{b_i\}_{i\in[\beta]}$; that is, the basis is $\boldsymbol{\omega} \triangleq W_\beta$, where $W_0 \triangleq (1)$ and $W_{i+1} \triangleq W_i | (b_{\beta-i+1} \cdot W_i)$, and $|$ denotes concatenation.

Proof. (of Theorem 5) According to Lemma 4, it suffices to show that for any $1 \leq i \leq \min\{\log n, \beta\}$, 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 by Lemma 6, and hence also an MDS matrix. Thus, H_J is invertible, and the claim follows. □

V. Existence of α-correcting codes over large fields

In this section it is shown that for any n there exists an $[n, n-1]_{q^\alpha}$ α-correcting code, provided that q is large enough. In what follows, let $X(\boldsymbol{\omega}) \triangleq \cup_{\mathbf{t}\in\mathcal{N}_\alpha^n} X_{\mathbf{t}}(\boldsymbol{\omega})$, and let

$$\mathcal{S} \triangleq \{\ker_{\mathbb{F}_{q^\alpha}}(\mathbf{x}) | \mathbf{x} \in X(\boldsymbol{\omega})\}, \text{ where}$$
$$\ker_{\mathbb{F}_{q^\alpha}}(\mathbf{x}) \triangleq \{\mathbf{y} \in \mathbb{F}_{q^\alpha}^n | \mathbf{y}\cdot\mathbf{x}^\top = 0\}.$$

The proof of existence relies on the following simple lemmas, and some of the respective proofs are omitted due to space constraints.

Lemma 7. $|X(\boldsymbol{\omega})| \leq \binom{\alpha+n-1}{n-1} \cdot q^\alpha$.

Lemma 8. $|\mathcal{S}| \leq |X(\boldsymbol{\omega})|/(q-1)$.

Proof. Since $X(\boldsymbol{\omega})$ is a union of subspaces over $\mathbb{F}_q$, it follows that it is closed under multiplication by elements of $\mathbb{F}_q$. The claim follows since $\ker_{\mathbb{F}_{q^\alpha}}(\mathbf{v}) = \ker_{\mathbb{F}_{q^\alpha}}(\lambda\mathbf{v})$ for any $\mathbf{v} \in \mathbb{F}_{q^\alpha}^n$ and any nonzero $\lambda \in \mathbb{F}_{q^\alpha}$, and in particular, any nonzero $\lambda \in \mathbb{F}_q$. □

Lemma 9. *The size of a set of hyperplanes in $\mathbb{F}_{q^\alpha}^n$ (i.e., subspaces of dimension $n-1$ of $\mathbb{F}_{q^\alpha}^n$) whose union is the entire space is at least $q^\alpha + 1$.*

Theorem 10. *If $q \geq \binom{\alpha+n-1}{n-1} + 1$ then there exists an $(n-1)$-dimensional α-correcting code in $\mathbb{F}_{q^\alpha}^n$.*

Proof. According to Lemma 7, Lemma 8, and the bound on q, it follows that $|\mathcal{S}| \leq q^\alpha$. Therefore, the union of the hyperplanes in $\mathcal{S}$ is not the entire space $\mathbb{F}_{q^\alpha}^n$ by Lemma 9. Hence, there exists $\mathbf{h} \triangleq (h_1, \ldots, h_n) \in \mathbb{F}_{q^\alpha}^n$ such that $\mathbf{h}\cdot\mathbf{x}^\top \neq 0$ for every nonzero $\mathbf{x} \in X(\boldsymbol{\omega})$. In turn, this implies that $\ker_{\mathbb{F}_{q^\alpha}}(\mathbf{h}) \cap X(\boldsymbol{\omega}) = \{0\}$, and hence $\ker_{\mathbb{F}_{q^\alpha}}(\mathbf{h})$ in an $n-1$ dimensional α-correcting code. □

In the full version of this paper, a slightly more involved argument will prove the existence of non-MDS m-correcting codes, where the lower bound on q goes to 1 as α grows.

VI. Connections to known codes

A. GRS codes

The construction in Section IV-B is closely related to a classical coding theoretic notion called *alternant codes* [8, 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 $\mathcal{C}_{alt}$ 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 $0 \leq i \leq \beta$ let $\mathcal{C}_i$ be the right kernel of $\mathcal{H}^{(i)}(V)$ over $\mathbb{F}_{q^{2^i}}$. Notice that Lemma 6 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 $j \leq i$. Lemma 6 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.

B. Gabidulin codes

In this section we restrict our attention to $n \leq \alpha$ and to erasure patterns we call *bounded*, but allow more than α erasures overall. For positive integers r and m, an erasure pattern $\mathbf{t} \in \mathcal{N}_n^m$ is called r-bounded if $t_j \leq r$ for all $j \in [n]$, and let $T_r \subseteq \mathcal{N}_n^m$ be the set of r-bounded erasures. Note that bounded erasure patterns are more restrictive than balanced ones, since balanced patterns allow in particular a full symbol erasure. In what follows it is shown that Gabidulin codes are T_r-correcting, and the size of the resulting codes is in correspondence with r.

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 non-negative 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.

Theorem 11. *For nonnegative integers ℓ, n, and α such that $n \leq \alpha$ and $r < n$, the code*

$$\mathrm{Gab}[n, n-r] \triangleq \{(f(\omega_1), \ldots, f(\omega_n)) | \\ f \text{ is linearized and } \deg_q(f) < n-r\}.$$

is T_r-correcting.

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

$$f(\omega_j) \in \langle \omega_1, \ldots, \omega_{t_j} \rangle, \text{ for all } j \in [n]. \tag{5}$$

Since f is a linearized polynomial and since $\mathbf{t} \in T_r$, Eq. (5) implies that $f(\langle\omega_1, \ldots, \omega_n\rangle) \subseteq \langle\omega_1, \ldots, \omega_r\rangle$, which in turn implies that $\dim\ker(f) \geq n-r$. Thus, f has more roots than its degree, which is a contradiction. □

Note that $n \leq \alpha$ is necessary to have n linearly independent evaluation points for $\mathrm{Gab}[n, n-r]$. Further, the dimension of $\mathrm{Gab}[n, n-r]$ over $\mathbb{F}_{q^\alpha}$ is $n-r$, and hence, for this construction to be nontrivial we must have that $r < n$. Moreover, since every $\mathbf{t} \in T_r$ sums to m, the necessary condition $m \leq n(n-1)$ is inevitable. Finally, we emphasize that this construction applies to any q.

C. Universally decodable matrices

The problems addressed in this paper are intimately related to *universally decodable matrices* (UDMs) [6], [12], which are a useful tool in error correction of *slow-fading channels* [11].

Definition 2 ([6, Def. 1])**.** *The matrices $A_1, \ldots, A_n \in \mathbb{F}_q^{\alpha \times \alpha}$ are called Universally Decodable Matrices (UDMs) if for every $\mathbf{t} = (t_1, \ldots, t_n) \in \mathcal{N}_\alpha^n$ the following condition is satisfied: the $\alpha \times \alpha$ matrix composed of the first t_1 rows of A_1, the first t_2 rows of A_2, ..., the first t_α rows of A_α has full rank.*

In the following theorem let J_α be the $\alpha \times \alpha$ matrix which contains 1's in its anti-diagonal, and zero elsewhere.

Theorem 12 ([12, Prop. 14])**.** *Let n, m, and α be positive integers, let q be a prime power such that $q \geq n-1$, and let γ be a primitive element in $\mathbb{F}_q$. Then, the following are $\alpha \times m$ UDMs over $\mathbb{F}_q$*

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

UDMs can be used to obtain codes with hierarchical erasure correction capabilities that are not necessarily linear over $\mathbb{F}_{q^\alpha}$. It is fairly easy to show that if $\{A_i\}_{i \in [n]}$ are $\alpha \times m$ UDMs, then

$$\mathcal{C} \triangleq \{x \in (\mathbb{F}_q^\alpha)^n | (A_0^\top, A_1^\top, \ldots, A_{n-1}^\top) \cdot x = 0\},$$

is an m-correcting code that is not necessarily linear over $\mathbb{F}_{q^\alpha}$, and hence suffers from higher encoding complexity[4].

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

Theorem 13. *For $h_1, \ldots, h_n \in \mathbb{F}_{q^\alpha}$, a code $\mathcal{C} \subseteq \{\mathbf{c} \in \mathbb{F}_{q^\alpha}^n | (h_1, \ldots, h_n) \cdot \mathbf{c}^\top = 0\}$ is an α-correcting code 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 $\boldsymbol{\omega} \triangleq (\omega_1, \ldots, \omega_\alpha)^\top$.*

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.,

$$A_i \boldsymbol{\omega} = h_i \boldsymbol{\omega} \text{ for all } i \in [n]. \tag{6}$$

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

$$\begin{aligned} h_i c_i &\in \langle h_i \omega_1, \ldots, h_i \omega_{t_i} \rangle \\ &\overset{(6)}{=} \left\langle A_i^{(1)} \boldsymbol{\omega}, \ldots, A_i^{(t_i)} \boldsymbol{\omega} \right\rangle, \end{aligned}$$

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 $\mathbf{v}_i A_i^{(1:t_i)} = h_i c_i$, 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

$$\mathbf{v} \cdot \begin{pmatrix} A_1^{(1:t_1)} \\ A_2^{(1:t_2)} \\ \vdots \\ A_n^{(1:t_n)} \end{pmatrix} = \sum_{i \in [n]} \mathbf{v}_i A_i^{(t_i)} = \sum_{i \in [n]} h_i c_i = 0, \tag{7}$$

which is a contradiction to $\{A_i\}_{i=1}^n$ being UDMs.

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 for contradiction 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

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

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 $c_i \triangleq \mathbf{v}_i \cdot (\omega_1, \ldots, \omega_{t_i})$ for all $i \in [n]$, and let $\mathbf{c} \triangleq (c_1, \ldots, c_n)$. Hence, it is readily verified that $\mathbf{c}$ is a nonzero codeword in $\mathcal{C} \cap X_{\mathbf{t}}$, a contradiction to $\mathcal{C}$ being an α-correcting code. □

Finally, we note that Theorem 1 can alternatively be proved by a direct application of Theorem 13, and the details are left to the curious reader.

VII. Discussion

In this paper an algebraic framework was given for erasure correction of scalar codes in distributed systems under the straggler effect. In the general case, a construction of a code of length two was given, and the existence of longer codes was proved. In more restricted settings, two code constructions were given which apply to any code length. Further results will be given in the full version of this paper.

Many natural questions remain widely open. Aside of finding a general construction of a $k\alpha$-correcting $[n, n-k]_{q^\alpha}$ code for $k > 1$, an intriguing research direction is studying this question without the assumption that the non-erased bits are error free, and finding the respective error-correction radius. Another challenging direction is to find scalar codes that are capable of mitigating the straggler effect during node repair.

[4]Assuming that multiplication in $\mathbb{F}_{q^\alpha}$ requires $O(\alpha \log \alpha)$ field operations.

References

[1] B. Bartan and M. Wootters, "Repairing multiple failures for scalar MDS codes," *arXiv:1707.02241* [cs.IT], 2017.

[2] A. Chowdhury and A. Vardy, "Improved schemes for asymptotically optimal repair of MDS codes," *arXiv:1710.01867* [cs.IT], 2017.

[3] R. Cohen and Y. Cassuto, "Iterative decoding of LDPC codes over the q-ary partial erasure channel," *IEEE Trans. on Inf. Th.* vol. 62, no. 5, pp. 2658–2672, 2016.

[4] R. Cohen, N. Raviv, and Y. Cassuto, "LDPC codes over the q-ary multi-bit channel," *arXiv:1706.09146* [cs.IT], 2017.

[5] S. H. Dau, I. M. Duursma, H. M. Kiah, and O. Milenkovic, "Repairing Reed-Solomon codes with multiple erasures," *arXiv:1612.01361* [cs.IT], 2016.

[6] A. Ganesan and P. O. Vontobel, "On the existence of universally decodable matrices," *IEEE Trans. on Inf. Th.*, vol. 53, no. 7, pp. 2572–2575, 2007.

[7] V. Guruswami and M. Wootters, "Repairing Reed-Solomon codes," *IEEE Trans. on Inf. Th.*, vol. 63, no. 9, pp. 5684–5698, 2017.

[8] R. Roth, *Introduction to coding theory*, Cambridge University Press, 2006.

[9] I. Tamo, M. Ye, and A. Barg, "Optimal repair of Reed-Solomon codes: achieving the cut-set bound," *IEEE Ann. Symp. on Found. of Comp. Sci. (FOCS)*, pp. 216–227, 2017.

[10] ——, "Fractional decoding: Error correction from partial information," *IEEE Intl. Symp. on Inf. Th. (ISIT)*, pp. 998–1002, 2017.

[11] S. Tavildar and P. Viswanath, "Approximately universal codes over slow-fading channels," *IEEE Trans. on Inf. Th.*, vol. 52, no. 7, pp. 3233–3258, 2006.

[12] P. O. Vontobel and A. Ganesan, "On universally decodable matrices for spacetime coding," *Designs, Codes, and Cryptography*, vol. 41, no. 3, pp. 325–342, 2006.

[13] M. Ye and A. Barg, "Repairing Reed-Solomon codes: Universally achieving the cut-set bound for any number of erasures," *arXiv:1710.07216* [cs.IT], 2017.