On Decoding Rank-Metric Codes over Large Fields

Ron M. Roth, *Fellow, IEEE*

***Abstract*—A decoding algorithm is presented for rank-metric array codes that are based on diagonal interleaving of maximum-distance separable (MDS) codes. With respect to this metric, such array codes are known to be optimal when the underlying field is algebraically closed. It is also shown that for any list decoding radius that is smaller than the minimum rank distance, the list size can be bounded from above by an expression that is independent of the field.**

***Index Terms*—Array codes, decoding, list decoding, rank metric.**

I. Introduction

Rank-metric codes over finite fields were introduced almost 40 years ago [3], [6], [14], yet the interest in them was revived in recent years due to the application of (variants of) such codes to certain network coding schemes [10], [17].

The (optimal) finite-field maximum-rank distance (MRD) construction of [3], [6] has been generalized also to certain infinite fields [1], [12], [15, §6]. However, since the construction assumes the existence of algebraic field extensions with prescribed extension degrees over the underlying field, such generalizations are not applicable to fields such as algebraically closed fields—in particular to the complex field, which is more relevant to low-rank metric recovery problems [2]. In fact, except for trivial cases, the rank-metric Singleton bound cannot be attained over algebraically closed fields (see Section II below).

In [14, §5] a simple construction of rank-metric codes was presented, which is optimal in case the field is algebraically closed. Yet no decoding algorithm was presented for these codes and, so, one purpose of this work is to present such an algorithm. We recall the construction in Section II, and present a linear-algebraic decoding algorithm in Section III. Then, in Section IV, we show that for any list decoding radius that is smaller than the minimum rank distance, the list size of a list decoding algorithm for these codes can be bounded from above by an expression that is independent of the field. This, in turn, provides evidence that, with respect to list decoding, there is a range of parameters for which the rank-metric construction studied here has advantages also when the underlying field is a (sufficiently large) finite field.

II. Construction

Let F be a field and let n and μ be positive integers such that $\mu \leq n$. For integers $a < b$, we denote by $[a,b\rangle$ the set $\{a, a+1, a+2, \ldots, b-1\}$, and we let $[b\rangle$ be a shorthand notation for $[0,b\rangle$. We index the rows and columns in an array (matrix) A over F starting at 0, with $A_{i,j}$ denoting the entry at row i and column j.

For every integer $m \in [1-n,n\rangle$, *diagonal* m in an $n \times n$ matrix refers to the set of $n-|m|$ row–column index pairs

$$\mathcal{D}_m = \mathcal{D}_m(n) = \Big\{(i,j) \in [n\rangle \times [n\rangle \ : \ j-i = m\Big\} .$$

In particular, $\mathcal{D}_0$ indexes the main diagonal and $\mathcal{D}_{1-n}$ (respectively, $\mathcal{D}_{n-1}$) indexes the length-1 diagonal consisting of the lower-left (respectively, upper-right) entry (see Figure 1: the spade suit marks the set $\mathcal{D}_{-2}$ and the diamond suit marks the set $\mathcal{D}_1$). We assume the obvious ordering on the elements of

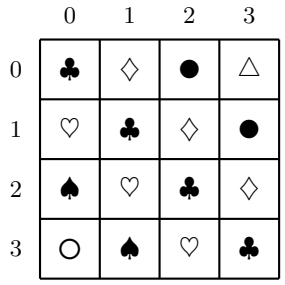

Fig. 1. Diagonals in an $n \times n$ matrix, for $n = 4$.

$\mathcal{D}_m$, where $(i, i+m) < (i', i'+m)$ if and only if $i < i'$. For a matrix $A \in F^{n\times n}$, the notation $(A)_{\mathcal{D}_m}$ stands for the row vector in $F^{n-|m|}$ that is formed by the entries indexed by $\mathcal{D}_m$ (in the order assumed on the elements of $\mathcal{D}_m$).

In [14, §5], the following array code $\mathsf{C} = \mathsf{C}_F(n,\mu)$ over (a sufficiently large field) F was presented:

$$\mathsf{C} = \Big\{\Gamma \in F^{n\times n} \ : \ (\Gamma)_{\mathcal{D}_m} \in \mathcal{C}_m \ , \quad \text{for } m \in [1-n,n\rangle\Big\} ,$$

where the set $\mathcal{C}_m = \mathcal{C}_m(n-|m|,\mu)$ is a prescribed linear maximum-distance separable (MDS) code over F of length $n-|m|$ and minimum (Hamming) distance μ (and dimension $n-|m|-\mu+1$; when $n-|m| < \mu$, the code $\mathcal{C}_m$ contains just the all-zero codeword). As a concrete example, each constituent code $\mathcal{C}_m$ can be taken as a (possibly extended) generalized Reed–Solomon (GRS) code over F, assuming that $|F| \geq n-1$ when $1 < \mu < n$.

The code C is linear over F with dimension

$$\begin{aligned} k &= \sum_{m=\mu-n}^{n-\mu} (n-|m|-\mu+1) \\ &= (n-\mu+1) + 2\sum_{m=1}^{n-\mu}(n-m-\mu+1) \end{aligned}$$

Ron M. Rothis with the Computer Science Department, Technion, Haifa 3200003, Israel. Email: ronny@cs.technion.ac.il.

This work was supported in part by Grants Nos. 1092/12 and 1396/16 from the Israel Science Foundation. This work was presented in part at the IEEE Int'l Symposium on Information Theory, Aachen, Germany (June 2017).

$$= (n-\mu+1) + 2\sum_{\ell=1}^{n-\mu} \ell$$
$$= (n-\mu+1)^2 .$$

Moreover, in every nonzero $\Gamma \in \mathsf{C}$, there is necessarily a smallest (or largest) index m such that $(\Gamma)_{\mathcal{D}_m}$ is a nonzero vector, which, in turn, is a codeword of $\mathcal{C}_m$ and therefore must contain at least μ nonzero entries; thus, the rank of Γ must be at least μ (and there are code arrays $\Gamma \in \mathsf{C}$ of rank exactly μ). Hence, in the notation of [14], the code C is a linear μ-$[n \times n, k{=}(n{-}\mu{+}1)^2]$ array code over F (with minimum rank μ).

Conversely, it is known that the dimension k of any linear μ-$[n \times n, k]$ array code $\boldsymbol{C}$ over an algebraically closed field F must satisfy the upper bound

$$k \le (n-\mu+1)^2 . \tag{1}$$

This bound is obtained through the characterization of the set $\mathcal{M} = \mathcal{M}_F(n, \mu{-}1)$ of $n \times n$ matrices of rank $< \mu$ over such a field F as an irreducible algebraic variety of dimension $\dim \mathcal{M} = n^2 - (n-\mu+1)^2$ [8, Prop. 12.2]; the bound (1) then follows from

$$0 = \dim(\{0\}) = \dim(\mathcal{M} \cap \boldsymbol{C}) \ge \dim \mathcal{M} + \dim \boldsymbol{C} - n^2$$

(see [9, p. 48, Prop. 7.1]). Note that unless $\mu = 1$, the bound (1) is strictly stronger than the rank-metric Singleton bound $k \le n(n{-}\mu{+}1)$, which applies to any field and which is attained by the MRD construction for finite fields (as well as for certain other fields [1], [15]).[1]

III. Decoding

Let $\Gamma \in \mathsf{C} = \mathsf{C}_F(n, \mu)$ represent a transmitted code array and let the received array be $Y = \Gamma + E$, where E is in $F^{n \times n}$ and has rank less than $\mu/2$. In this section, we present an efficient algorithm for recovering Γ out of Y. The decoding problem can be seen as a variant of a matrix completion problem of the low-rank matrix E, where the contents of E along the diagonals $\mathcal{D}_m$ is known to the decoder for $|m| > n - \mu$ (since along those diagonals Γ is zero), and is concealed by additive "noise"—namely, codewords of $\mathcal{C}_m(n{-}|m|, \mu)$—along the remaining diagonals.

A. Overview

We start by presenting the principles of the decoding (albeit not in a manner that yields the most efficient implementation of it). Our decoder recovers iteratively the contents of E (or, rather, of Γ) along diagonal $\mathcal{D}_m$, as m ranges from $1{-}n$ to $n{-}1$.

For $m \in [1{-}n, n\rangle$, let

$$\mathcal{T}_m = \mathcal{T}_m(n) = \cup_{m' \in [1-n, m\rangle} \mathcal{D}_{m'} \tag{2}$$
$$= \Big\{(i,j) \in [n\rangle \times [n\rangle \ : \ j - i < m\Big\} ,$$

namely, $\mathcal{T}_m$ indexes all entries in an $n \times n$ array that lie below diagonal $\mathcal{D}_m$ (the shape of $\mathcal{T}_m$ is a triangle for $m \le 1$ and a pentagon for $m > 1$).

Suppose that the entries of E have already been recovered at all positions $(i,j) \in \mathcal{T}_m$, for some $m \in [1{-}n, n\rangle$ (recall that this holds vacuously for $m = \mu - n$, since $(\Gamma)_{\mathcal{D}_{m'}}$ is necessarily the all-zero codeword when $m' < \mu{-}n$). Consider the $n \times n$ array $\tilde{Y}$ whose entries are given by:

$$\tilde{Y}_{i,j} = \begin{cases} E_{i,j} & \text{if } (i,j) \in \mathcal{T}_m \\ Y_{i,j} & \text{otherwise} \end{cases}$$

(see an example in Figure 2; note that $\tilde{Y}$ would be the received array if $(\Gamma)_{\mathcal{D}_{m'}}$ were the all-zero codeword for all $m' < m$). Now, perform a (partial) Gaussian elimination on $\tilde{Y}$,

♣	♢	●	△	♠	○	■
♡	♣	♢	●	△	♠	○
$e_{2,0}$	♡	♣	♢	●	△	♠
$e_{3,0}$	$e_{3,1}$	♡	♣	♢	●	△
$e_{4,0}$	$e_{4,1}$	$e_{4,2}$	♡	♣	♢	●
$e_{5,0}$	$e_{5,1}$	$e_{5,2}$	$e_{5,3}$	♡	♣	♢
$e_{6,0}$	$e_{6,1}$	$e_{6,2}$	$e_{6,3}$	$e_{6,4}$	♡	♣

Fig. 2. Array $\tilde{Y}$, for $n = 7$ and $m = -1$. The heart suit marks $\mathcal{D}_m$.

by subtracting scalar multiples of rows (respectively, columns) that intersect $\mathcal{T}_m$, from rows of a *lower* index (respectively, from columns of a *higher* index), thereby eliminating as many nonzero entries as possible among the entries that are indexed by $\mathcal{T}_m$. It is easy to see that during this process, entries in $\tilde{Y}$ that are indexed by $\mathcal{T}_m \cup \mathcal{D}_m$ are modified by terms that depend only on the already-recovered entries E (and not on any of the remaining entries of $\tilde{Y}$).

Denote by $\tilde{Z}$ the result of the Gaussian elimination, and let $\mathcal{P}$ be the set of index pairs $(i,j) \in \mathcal{T}_m$ such that $\tilde{Z}_{i,j} \neq 0$ (see an example in Figure 3). Next, we mark as erasures the entries along $\mathcal{D}_m$ that share a row index or a column index with elements of $\mathcal{P}$. As we show in Section III-C below, the number of erasures plus twice the number of *additional* errors in $(\tilde{Z})_{\mathcal{D}_m}$ must be less than μ, thereby allowing us to recover $(E)_{\mathcal{D}_m}$ by applying a combined error–erasure decoder for $\mathcal{C}_m(n{-}|m|, \mu)$ to $(\tilde{Z})_{\mathcal{D}_m}$.

In the next subsections, we present and analyze a more efficient implementation of this decoding idea. In particular, through bookkeeping of the elementary operations applied when $m' < m$, we do not need to start the Gaussian elimination all over again when decoding $(E)_{\mathcal{D}_m}$.

B. The decoding algorithm

Figure 4 presents a decoder for C. As mentioned in Section III-A, the algorithm recovers iteratively the contents of Γ along the diagonals $\mathcal{D}_m$, for $m \leftarrow 1{-}n, 2{-}n, \ldots, n{-}1$ (see the main loop in Figure 4). The input to the algorithm is the $n \times n$ received array Y, and the output is a reconstructed code

[1] The results of this work can be extended easily through code shortening to $\ell \times n$ rectangular arrays: the dimension of C then becomes $(\ell{-}\mu{+}1)(n{-}\mu+1)$, which is also the value that will replace the right-hand side of (1). For simplicity, however, we have opted to assume that $\ell = n$.

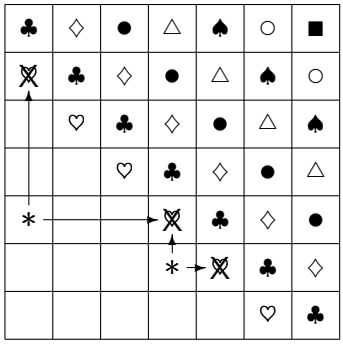

Fig. 3. Array $\tilde{Z}$, for $n = 7$ and $m = -1$. The stars mark the positions indexed by $\mathcal{P}$, and the crossed entries are the (at most $2|\mathcal{P}|$) erasure positions along diagonal $\mathcal{D}_m$.

array $\widehat{\Gamma}$. The algorithm uses several data structures: a subset $\mathcal{P} \subseteq [n\rangle \times [n\rangle$ (as in Section III-A) of size less than $\mu/2$, another subset $\mathcal{Q} \subseteq [n\rangle \times [n\rangle$ (marking erasure positions) of size less than μ, and several $n \times n$ matrices over F, namely, L, R, U, and V. These matrices are sparse: L and R contain $O(\mu n)$ nonzero entries and U and V contain less than $\mu/2$ nonzero entries. A fifth matrix, Z, is represented as such only for convenience: except for less than $\mu/2$ of its entries (which are indexed by the set $\mathcal{P}$), each iteration of the algorithm utilizes only one diagonal of Z and, so, different iterations may reuse the same area ($(Z)_{\mathcal{D}_m}$ will coincide with $(\tilde{Z})_{\mathcal{D}_m}$ described in Section III-A).

Input: received array $Y \in F^{n\times n}$;
Output: code array $\widehat{\Gamma} \in \mathsf{C}_F(n,\mu)$.

$\widehat{\Gamma} \leftarrow 0$;
$L, R \leftarrow I_n$; /* Initialize to the identity matrix */
$\mathcal{P} \leftarrow \emptyset$;

For $m \leftarrow 1{-}n, 2{-}n, \ldots, n{-}1$ do:

1) $(Z)_{\mathcal{D}_m} \leftarrow (L \cdot (Y - \widehat{\Gamma}) \cdot R)_{\mathcal{D}_m}$;
2) $\mathcal{Q} \leftarrow \{(i,j) \in \mathcal{D}_m \ : \ \exists r$ s.t. $(i,r) \in \mathcal{P}$ or $(r,j) \in \mathcal{P}\}$;
3) Decode $(Z)_{\mathcal{D}_m}$ into $(\widehat{\Gamma})_{\mathcal{D}_m} \in \mathcal{C}_m$ to recover $|\mathcal{Q}|$ erasures indexed by $\mathcal{Q}$ and less than $(\mu{-}|\mathcal{Q}|)/2$ additional errors;
4) $(Z)_{\mathcal{D}_m} \leftarrow (Z)_{\mathcal{D}_m} - (\widehat{\Gamma})_{\mathcal{D}_m}$;
5) $U, V \leftarrow 0$;
6) For $(i, i{+}m) \in \mathcal{D}_m$ s.t. $(j, i{+}m) \in \mathcal{P}$ for some j do:
   a) $U_{i,j} \leftarrow Z_{i,i+m}/Z_{j,i+m}$;
   b) $Z_{i,i+m} \leftarrow 0$;
7) For $(j{-}m, j) \in \mathcal{D}_m$ s.t. $(j{-}m, i) \in \mathcal{P}$ for some i do:
   a) $V_{i,j} \leftarrow Z_{j-m,j}/Z_{j-m,i}$;
   b) $Z_{j-m,j} \leftarrow 0$;
8) Update:
   a) $L \leftarrow (I_n - U) \cdot L$;
   b) $R \leftarrow R \cdot (I_n - V)$;
   c) $\mathcal{P} \leftarrow \mathcal{P} \cup \{(i,j) \in \mathcal{D}_m \ : \ Z_{i,j} \neq 0\}$;
   d) If $|\mathcal{P}| \geq \mu/2$ declare "decoding failure" and stop.

Fig. 4. Decoding algorithm for $\mathsf{C}_F(n,\mu)$.

Each iteration m starts by setting an initial contents to the entries of Z that are indexed by diagonal $\mathcal{D}_m$ (step 1); that contents is a function of the received array Y, the already-decoded diagonals of Γ, and matrices L and R computed in previous rounds. A combined error–erasure decoders for $\mathcal{C}_m = \mathcal{C}_m(n{-}|m|,\mu)$ is then applied to $(Z)_{\mathcal{D}_m}$ to recover $\varrho \in [\mu\rangle$ erasures and less than $(\mu - \varrho)/2$ additional errors (steps 2–4); the ϱ erasures are indexed by the set $\mathcal{Q}$ which consists of all index pairs $(i,j) \in \mathcal{D}_m$ that share a row or a column with the support, $\mathcal{P}$, of the previously-computed entries of Z (along diagonals $\mathcal{D}_{m'}$ for $m' < m$). In steps 5–7, a partial Gaussian elimination is performed on the entries of Z that are indexed by $(i,j) \in \mathcal{D}_m$, with $I_n - U$ representing row operations and $I_n - V$ column operations. The overall effect of the row and column operations that are performed throughout the course of the main loop is accumulated into the matrices L and R. Those matrices, along with the support $\mathcal{P}$ of Z, are updated in step 8.

C. Validity

Next, we analyze the algorithm. The value of a variable in the algorithm (e.g., the matrix Z) right after step 2 in iteration m of the main loop will be denoted by adding the superscript m (e.g., $Z^{(m)}$). Also, define the matrix $E^{(m)}$ by

$$E^{(m)} = L^{(m)} \cdot E \cdot R^{(m)} \ , \tag{3}$$

and recall the definition of $\mathcal{T}_m$ in (2).

We have the next proposition, part (vi) of which establishes the validity of the algorithm in Figure 4.

Proposition 1. *Let $Y = \Gamma + E$, where $\Gamma \in \mathsf{C}$ and $\mathrm{rank}(E) < \mu/2$. The following properties hold for iteration m of the algorithm in Figure 4, for all $m \in [1{-}n, n\rangle$.*

i) $\mathcal{P}^{(m)} = \left\{(i,j) \in \mathcal{T}_m \ : \ Z^{(m)}_{i,j} \neq 0\right\}$, *and no two index pairs in $\mathcal{P}^{(m)}$ share the same row or the same column.*
ii) $Z^{(m)}_{i,j} = E^{(m)}_{i,j}$, *for $(i,j) \in \mathcal{T}_m$.*
iii) $L^{(m)}$ *and* $R^{(m)}$ *are upper-triangular, with the main diagonal entries being all* 1*, and the nonzero entries in $I_n - L^{(m)}$ (respectively, in $I_n - R^{(m)}$) are all in columns j (respectively, rows i) such that $(j,i) \in \mathcal{P}^{(m)}$ for some i (respectively, j).*
iv) $|\mathcal{P}^{(m)}| \leq \mathrm{rank}(E^{(m)}) = \mathrm{rank}(E) < \mu/2$.
v) $(Z^{(m)})_{\mathcal{D}_m} = (\Gamma)_{\mathcal{D}_m} + (E^{(m)})_{\mathcal{D}_m}$.
vi) *The decoding in step 3 is successful, namely,*

$$(\widehat{\Gamma}^{(m+1)})_{\mathcal{D}_m} = (\Gamma)_{\mathcal{D}_m} \ .$$

Proof. We start by showing that parts (i), (ii), and (v) imply the other parts (for the *same* iteration m). We will then use induction to establish those former parts.

(i) ⇒ (iii). By part (i) it follows that $\mathcal{P}^{(m)} \subseteq \mathcal{T}_m$, and by steps 6 and 7 in Figure 4 we then get that $U^{(m)}$ and $V^{(m)}$ are strictly upper-triangular. This and steps 8a–8c imply the result.

(i)–(iii) ⇒ (iv). By part (iii) and the definition of $E^{(m)}$ it follows that $\mathrm{rank}(E^{(m)}) = \mathrm{rank}(E)$. Also, parts (i) and (ii) imply that $|\mathcal{P}^{(m)}| \leq \mathrm{rank}(E^{(m)})$.

(iv)–(v) ⇒ (vi). The set $\mathcal{Q}^{(m)}$, which is computed in step 2 during iteration m, contains all index pairs in $\mathcal{D}^{(m)}$ that share a row or a column with some index pair in $\mathcal{P}^{(m)}$. Clearly, $|\mathcal{Q}^{(m)}| \leq 2|\mathcal{P}^{(m)}|$ (see Figure 3). Now, among the entries of $(E^{(m)})_{\mathcal{D}_m}$ that are indexed by $\mathcal{D}_m \setminus \mathcal{Q}^{(m)}$, only less than

$$\mu/2 - |\mathcal{P}^{(m)}| \leq (\mu - |\mathcal{Q}^{(m)}|)/2$$

can be nonzero, or else the rank of $E^{(m)}$ would be at least $\mu/2$, thereby contradicting part (iv). From part (v), we conclude that the decoding in step 3 is guaranteed to be successful.

We now prove parts (i), (ii), and (v) for iteration m, assuming by induction that all parts hold for some iteration $m-1$ (the induction base, $m = 1-n$, is obvious).

Induction step for part (i). Steps 4, 6b, and 7b in iteration $m-1$ do not affect the contents of Z at positions $(i,j) \in \mathcal{T}_{m-1}$, yet eliminate nonzero entries of Z at positions $(i,j) \in \mathcal{D}_{m-1}$ that share rows and columns with $\mathcal{P}^{(m-1)}$. The result follows from step 8c and the induction hypothesis on part (i).

(As we pointed out earlier, steps 5–7 in the iterations $1-n, 2-n, \ldots, m-1$ perform a Gaussian elimination on the rows and columns of the portion of Z that is indexed by $(i,j) \in \mathcal{D}_{m-1} \cup \mathcal{T}_{m-1} = \mathcal{T}_m$, with $I_n - U$ representing the row operations and $I_n - V$ the column operations in each iteration. As a result of this elimination, the nonzero entries in that portion of Z, which are indexed by $\mathcal{P}^{(m)}$, must be in distinct rows and columns.)

Induction step for part (ii). Assuming the induction hypothesis for iteration $m-1$, we get, for every $(i,j) \in \mathcal{T}_{m-1}$:

$$\begin{aligned} E^{(m)}_{i,j} &= (L^{(m)} \cdot E \cdot R^{(m)})_{i,j} \\ &= ((I_n - U^{(m)}) \\ &\qquad \cdot L^{(m-1)} \cdot E \cdot R^{(m-1)} \cdot (I_n - V^{(m)}))_{i,j} \\ &\overset{(3)}{=} ((I_n - U^{(m)}) \cdot E^{(m-1)} \cdot (I_n - V^{(m)}))_{i,j} \\ &= ((I_n - U^{(m)}) \cdot Z^{(m-1)} \cdot (I_n - V^{(m)}))_{i,j} \\ &= Z^{(m-1)}_{i,j} , \end{aligned}$$

where the penultimate equality follows from the induction hypothesis on part (ii) and the definition of the matrices U and V; the properties of these matrices also imply the last equality. Since iteration $m-1$ affects $Z_{i,j}$ only at positions $(i,j) \in \mathcal{D}_{m-1}$, it follows that $Z^{(m)}_{i,j} = Z^{(m-1)}_{i,j}$ for every $(i,j) \in \mathcal{T}_{m-1}$, namely, for this range of index pairs,

$$Z^{(m)}_{i,j} = E^{(m)}_{i,j} .$$

It remains to show that the latter equality holds also for $(i,j) \in \mathcal{D}_{m-1}$ $(= \mathcal{T}_m \setminus \mathcal{T}_{m-1})$. Let $\tilde{Z}$ be the contents of Z right after step 4 in iteration $m-1$. It follows from the induction hypothesis on parts (ii), (v), and (vi) that $\tilde{Z}_{i,j} = E^{(m-1)}_{i,j}$ for $(i,j) \in \mathcal{T}_m$. From steps 6 and 7 we then get, for $(i,j) \in \mathcal{D}_{m-1}$,

$$\begin{aligned} Z^{(m)}_{i,j} &= ((I_n - U^{(m)}) \cdot \tilde{Z} \cdot (I - V^{(m)}))_{i,j} \\ &= (I_n - U^{(m)}) \cdot E^{(m-1)} \cdot (I_n - V^{(m)})_{i,j} \\ &\overset{(3)}{=} ((I_n - U^{(m)}) \\ &\qquad \cdot L^{(m-1)} \cdot E \cdot R^{(m-1)} \cdot (I_n - V^{(m)}))_{i,j} \\ &= (L^{(m)} \cdot E \cdot R^{(m)})_{i,j} \\ &= E^{(m)}_{i,j} . \end{aligned}$$

Induction step for part (v). From step 1 we have,

$$\begin{aligned} &(Z^{(m)})_{\mathcal{D}_m} \\ &\quad = (L^{(m)} \cdot (Y - \widehat{\Gamma}^{(m)}) \cdot R^{(m)})_{\mathcal{D}_m} \\ &\quad = (L^{(m)} \cdot (E + \Gamma - \widehat{\Gamma}^{(m)}) \cdot R^{(m)})_{\mathcal{D}_m} \\ &\quad \overset{(3)}{=} (L^{(m)} \cdot (\Gamma - \widehat{\Gamma}^{(m)}) \cdot R^{(m)})_{\mathcal{D}_m} + (E^{(m)})_{\mathcal{D}_m} . \end{aligned} \quad (4)$$

At this point, we can already assume that part (iii) holds for iteration m, since this part was proved based only on part (i), for which we have already established the induction step; thus, $L^{(m)}$ and $R^{(m)}$ are upper-triangular. In addition, by the induction hypothesis on part (vi) for iterations up to $m-1$, we have $\widehat{\Gamma}^{(m)}_{i,j} = \Gamma_{i,j}$ for all $(i,j) \in \mathcal{T}_m$. Therefore,

$$(L^{(m)} \cdot (\Gamma - \widehat{\Gamma}^{(m)}) \cdot R^{(m)})_{\mathcal{D}_m} = (\Gamma - \widehat{\Gamma}^{(m)})_{\mathcal{D}_m}$$

and, so, from (4),

$$\begin{aligned} (Z^{(m)})_{\mathcal{D}_m} &= (\Gamma - \widehat{\Gamma}^{(m)})_{\mathcal{D}_m} + (E^{(m)})_{\mathcal{D}_m} \\ &= (\Gamma)_{\mathcal{D}_m} + (E^{(m)})_{\mathcal{D}_m} , \end{aligned}$$

where the last equality follows from the fact that $\widehat{\Gamma}^{(m)}_{i,j} = 0$ for $(i,j) \notin \mathcal{T}_m$. □

Remark 1. When $\mathrm{rank}(E) < \mu/2$, the algorithm in Figure 4 produces the correct code array already at iteration $m = n-\mu$, since each subsequent iteration $m \in [n-\mu+1, n\rangle$ should result in an all-zero codeword, $(\widehat{\Gamma})_{\mathcal{D}_m}$, in step 3. Those last $\mu-1$ iterations, however, could be used for (partial) error detection through step 8d, in case $\mathrm{rank}(E) \geq \mu/2$. □

D. Complexity

We turn now to analyzing the time complexity of the algorithm in Figure 4, expressed as a number of arithmetic operations and element insertions into sets. It is easy to see that with the exception of steps 1, 3, and 8a–8b, all steps can be carried out in time complexity of $O(|\mathcal{P}|) = O(\mu)$ per iteration. (Step 4, as written, seemingly requires $O(n)$ operations per iteration, yet in fact it requires only $O(\mu)$ operations assuming that the decoder of $\mathcal{C}_m$ in step 3 also outputs the error locations and the error–erasure values.)

By part (iii) in Proposition 1 it follows that for every $(i,j) \in \mathcal{D}_m$, the update of $Z_{i,j}$ in step 1 requires $O(\mu)$ operations. Hence, this step requires $O(\mu n)$ operations per iteration.

Step 3 applies a decoder for $\mathcal{C}_m$; when this code is taken as a GRS code, then a straightforward implementation of a syndrome-based GRS decoder requires $O(\mu n)$ operations.

Finally, steps 8a–8b involve $|\mathcal{P}|$ elementary operations on the rows of L and on the columns of R, amounting to $O(\mu n)$ operations per iteration.

In summary, the time complexity of the decoder in Figure 4 is $O(\mu n^2)$. The space complexity has already been discussed at the beginning of Section III-B: with the exception of the input and output arrays (each containing n^2 elements of F), the algorithm uses two subsets of $[n\rangle \times [n\rangle$ of size $O(\mu)$, four sparse matrices (namely, L, R, U, and V) containing $O(\mu n)$ nonzero elements of F, and a matrix Z in which less than $n+\mu/2$ entries need to be kept from one step of the algorithm to the next.

In comparison, implementations of the known decoding algorithms for the MRD construction of [6] require a number of arithmetic operations (of F) which scales at least as $\mu n^2 \log n \log\log n$, whenever $\mu = O(n/\log\log n)$; see [7], [16], [19].

E. Algebraic formulation of the decoding problem

We consider now the special case where each constituent code $\mathcal{C}_m$ is a GRS code with the $(\mu-1) \times (n-|m|)$ parity-check matrix

$$H_m = \left(\alpha^{\ell(b+i)}\right)_{\ell\in[\mu-1\rangle,\, i\in[n-|m|\rangle} = \begin{pmatrix} 1 & 1 & \dots & 1 \\ \alpha^b & \alpha^{b+1} & \dots & \alpha^{n-|m|-1+b} \\ \alpha^{2b} & \alpha^{2(b+1)} & \dots & \alpha^{2(n-|m|-1+b)} \\ \vdots & \vdots & \vdots & \vdots \\ \alpha^{(\mu-2)b} & \alpha^{(\mu-2)(b+1)} & \dots & \alpha^{(\mu-2)(n-|m|-1+b)} \end{pmatrix},$$

where α is a prescribed nonzero element in F of multiplicative order $\mathcal{O}(\alpha) \geq n$ (e.g., $\mathcal{O}(\alpha) = \infty$), and $b = \max\{0, -m\}$.

Our analysis of this case will be under the more general setting of list decoding, where the error array is assumed to have rank at most τ, for a prescribed *decoding radius* $\tau \in \mathbb{Z}^+$ (we will elaborate more on list decoding in Section IV).

Let $Y = \Gamma + E$, where $\Gamma \in \mathsf{C}$ and $\mathrm{rank}(E) = \rho$. Then E can be written as

$$E = \sum_{r\in[\rho\rangle} \boldsymbol{u}_r^\top \boldsymbol{v}_r \;, \tag{5}$$

where each of the ordered sets $\{\boldsymbol{u}_r\}_{r\in[\rho\rangle}$ and $\{\boldsymbol{v}_r\}_{r\in[\rho\rangle}$ contains ρ row vectors in F^n that are linearly independent over F. For each vector $\boldsymbol{u}_r = (u_{r,i})_{i\in[n\rangle}$, we associate the polynomial

$$u_r(x) = \sum_{i\in[n\rangle} u_{r,i} x^i$$

in the set, $F_n[x]$, of polynomials of degree less than n over F. Also, for each $\boldsymbol{v}_r = (v_{r,j})_{j\in[n\rangle}$, we associate the polynomial

$$v_r^*(x) = \sum_{j\in[n\rangle} v_{r,j}^* x^j = \sum_{j\in[n\rangle} v_{r,n-1-j} x^j$$

(notice the reversed order of coefficients).

For $m \in [1-n, n\rangle$ and $\ell \in [\mu-1\rangle$, we define the syndrome component

$$\begin{aligned} S_{\ell,n-1-m} &= \left(H_m(Y)^\top_{\mathcal{D}_m}\right)_\ell \\ &= \left(H_m(E)^\top_{\mathcal{D}_m}\right)_\ell \;. \end{aligned}$$

The syndrome (polynomial) vector is defined by

$$\boldsymbol{S}(x) = (S_0(x)\; S_1(x)\; \dots\; S_{\mu-2}(x)) \;,$$

where, for each $\ell \in [\mu-1\rangle$,

$$\begin{aligned} S_\ell(x) &= \sum_{m\in[1-n,n\rangle} S_{\ell,n-1-m} x^{n-1-m} \\ &= \sum_{j\in[2n-1\rangle} S_{\ell,j} x^j \;\; (\in F_{2n-1}[x]) \;. \end{aligned}$$

We have:

$$\begin{aligned} S_\ell(x) &= \sum_{m\in[1-n,n\rangle} \left(H_m(E)^\top_{\mathcal{D}_m}\right)_\ell x^{n-1-m} \\ &= \sum_{m\in[1-n,n\rangle} x^{n-1-m} \sum_{i\,:\,(i,i+m)\in\mathcal{D}_m} \alpha^{i\ell} E_{i,i+m} \\ &\overset{(5)}{=} \sum_{m\in[1-n,n\rangle} x^{n-1-m} \sum_{i\,:\,(i,i+m)\in\mathcal{D}_m} \alpha^{i\ell} \sum_{r\in[\rho\rangle} u_{r,i} v_{r,i+m} \\ &= \sum_{r\in[\rho\rangle} \sum_{m\in[1-n,n\rangle} x^{n-1-m} \sum_{i\,:\,(i,i+m)\in\mathcal{D}_m} \alpha^{i\ell} u_{r,i} v^*_{r,n-1-m-i} \\ &= \sum_{r\in[\rho\rangle} u_r(\alpha^\ell x)\, v_r^*(x) \;. \end{aligned}$$

Recalling that the syndrome vector $\boldsymbol{S}(x)$ can be computed from the received array Y, the (list-)decoding problem up to a given decoding radius τ can be formulated as finding ordered pairs of vectors in $(F_n[x])^\rho$ for $\rho \in [\tau+1\rangle$,

$$\left(\boldsymbol{U}(x) = (u_r(x))_{r\in[\rho\rangle}, \boldsymbol{V}(x) = (v_r(x))_{r\in[\rho\rangle}\right) \;,$$

with the components of $\boldsymbol{U}(x)$ (respectively, $\boldsymbol{V}(x)$) being linearly independent polynomials over F, such that

$$S_\ell(x) = \sum_{r\in[\rho\rangle} u_r(\alpha^\ell x)\, v_r^*(x) \;, \quad \ell \in [\mu-1\rangle \;. \tag{6}$$

For each such pair $(\boldsymbol{U}(x), \boldsymbol{V}(x))$, the respective ordered sets, $\{\boldsymbol{u}_r\}_{r\in[\rho\rangle}$ and $\{\boldsymbol{v}_r\}_{r\in[\rho\rangle}$, define an error array E via (5). Moreover, the association $(\boldsymbol{U}(x), \boldsymbol{V}(x)) \mapsto E$ is one-to-one if $\boldsymbol{U}(x)$ is in the following canonical form:

- $u_r(x)$ is monic, for every $r \in [\rho\rangle$;
- $\deg u_{r-1}(x) < \deg u_r(x)$, for every $r \in [1, \rho\rangle$; and—
- $u_{r,i} = 0$, for every $r \in [1, \rho\rangle$ and $i \in \{\deg u_s(x)\}_{s\in[r\rangle}$.

Clearly, when $\tau < \mu/2$, Eq. (6) has at most one canonical solution $(\boldsymbol{U}(x), \boldsymbol{V}(x))$.

Example 1. We consider the case $\mu = 3$ and $\tau = (\mu-1)/2 = 1$ (correcting an error array of rank at most 1): here (6) reduces to

$$\begin{aligned} S_0(x) &= u(x)\, v^*(x) \\ S_1(x) &= u(\alpha x)\, v^*(x) \;, \end{aligned}$$

where, for simplicity, we have omitted the subscript 0 from $u_0(x)$ and $v_0(x)$. Clearly, $S_0(x) = S_1(x) = 0$ implies $E = 0$. Otherwise, $u(x)$ satisfies the equality

$$S_0(x)\, u(\alpha x) = S_1(x)\, u(x) \;, \tag{7}$$

which is, in fact, a set of linear equations in the coefficients of $u(x) \in F_n[x]$. Specifically, letting $d = \deg S_0(x)$, from (7) we have $d = \deg S_1(x)$ and

$$\sum_{j=i}^{\deg u} u_j \left(\alpha^j \cdot S_{0,d+i-j} - S_{1,d+i-j}\right) = 0 \;, \tag{8}$$

$$i = \deg u, \deg u - 1, \dots, 0$$

(note that $S_0(x) = u(x)v^*(x)$ implies that $\deg u \leq d$). It follows from (8) that $\deg u$ is the unique $t \in [n\rangle$ for which

$$\alpha^t = \frac{S_{1,d}}{S_{0,d}} \;.$$

Assuming that $u(x)$ is monic (i.e., $u_t = 1$), one can solve (8) iteratively—and uniquely—for u_i, $i = t-1, t-2, \dots, 0$ (the

coefficient of u_i in (8) equals $\alpha^i \cdot S_{0,d} - S_{1,d}$, which is necessarily nonzero for $i \in [t\rangle$).

Thus, we decode $E = \boldsymbol{u}^\top \boldsymbol{v}$ by first solving (7) for a monic $u(x) \in F_n[x]$, and then let $v^*(x) = S_0(x)/u(x)$. □

More generally, the decoding algorithm in Figure 4 can be reformulated to have as input the syndrome vector $\boldsymbol{S}(x)$ (instead of the array Y) and as output the error array E; a decomposition of E as in (5) (which can be carried out through a simple Gaussian elimination) then yields a solution $\big(\boldsymbol{U}(x)=(u_r(x))_{r\in[\rho\rangle}, \boldsymbol{V}(x)=(v_r(x))_{r\in[\rho\rangle}\big)$ of (6). Going into more detail, in each iteration m of the main loop in Figure 4, step 1 inserts into $(Z)_{\mathcal{D}_m}$ the vector

$$(L \cdot (Y - \widehat{\Gamma}) \cdot R)_{\mathcal{D}_m} = (\Gamma + E + \widehat{E}^{(m)})_{\mathcal{D}_m} \ ,$$

with the entries of $(\widehat{E}^{(m)})_{\mathcal{D}_m}$ standing for linear combinations of (already decoded) entries $E_{i,j} = Y_{i,j} - \widehat{\Gamma}_{i,j}$, for $(i,j) \in \mathcal{T}_m$. Hence, the syndrome of $(Z)_{\mathcal{D}_m}$ with respect to the parity-check matrix H_m can be expressed as:

$$\begin{aligned} H_m(Z)^\top_{\mathcal{D}_m} &= H_m\big(\Gamma + E + \widehat{E}^{(m)}\big)^\top_{\mathcal{D}_m} \\ &= H_m(E)^\top_{\mathcal{D}_m} + H_m\big(\widehat{E}^{(m)}\big)^\top_{\mathcal{D}_m} \\ &= \big(S_{\ell,n-1-m}\big)_{\ell\in[\mu-1\rangle} + H_m\big(\widehat{E}^{(m)}\big)^\top_{\mathcal{D}_m} \end{aligned}$$

(where $(S_{\ell,n-1-m})_{\ell\in[\mu-1\rangle}$ is the vector of coefficients of x^{n-1-m} in $\boldsymbol{S}(x)$). The syndrome $H_m(Z)^\top_{\mathcal{D}_m}$ can then be used in steps 3 and 4 to recover $(E + \widehat{E}^{(m)})_{\mathcal{D}_m}$, from which one can extract $(E)_{\mathcal{D}_m}$. Thus, throughout the iterations, $\boldsymbol{S}(x)$ (rather than the whole received array Y) suffices for obtaining the syndromes $H_m(Z)^\top_{\mathcal{D}_m}$, which, in turn, suffice in order to recover E (and, therefore, $(\boldsymbol{U}(x), \boldsymbol{V}(x))$).

IV. List decoding

We consider now the list decoding capabilities of the array code $\mathsf{C} = \mathsf{C}_F(n,\mu)$: given a decoding radius $\tau \in \mathbb{Z}^+$, we seek an upper bound on the list size $\mathsf{L}_\mathsf{C}(\tau)$, being the largest intersection of any coset of C (within $F^{n\times n}$) with the set of $n\times n$ matrices of rank at most τ over F. Clearly, $\mathsf{L}_\mathsf{C}(\tau)$ grows unboundedly with $|F|$ when $\tau \geq \mu$, since C contains at least $|F|-1$ arrays of rank μ (see also [4, Thm. 1]). Our main result (Theorem 4 below) implies that the converse is also true: when $\tau < \mu$, the list size is bounded from above by an expression that depends on n and μ, but not on F.

Before getting to the main result, we consider, in Propositions 2 and 3, the special cases $\mu \in \{2, n\}$ in more detail.

Proposition 2. *When each constituent code $\mathcal{C}_m$ in $\mathsf{C} = \mathsf{C}_F(n,2)$ is taken to be a single-parity code over F,*

$$\mathsf{L}_\mathsf{C}(1) \leq \binom{2n-2}{n-1} \ , \tag{9}$$

with equality holding if $|F| \geq 2n-3$.

Proof. For $\mu = 2$ and $\tau = 1$, Eq. (6) becomes

$$S_0(x) = u(x)\, v^*(x) \tag{10}$$

(where we have omitted the subscript from $u_0(x)$ and $v_0(x)$). In particular, $\boldsymbol{S}(x) = (S_0(x))$ forms the (whole) syndrome that is associated with a given coset of C. Constraining $u(x)$ to be monic, each pair $(u(x), v(x))$ uniquely defines an error array $E = \boldsymbol{u}^\top \boldsymbol{v}$ in that coset. Recalling that $S_0(x)$ is a nonzero polynomial in $F_{2n-1}[x]$, Eq. (10) is satisfied by at most $\binom{2n-2}{n-1}$ pairs $(u(x), v(x)) \in (F_n[x])^2$ where $u(x)$ is monic. Moreover, this bound is tight when $S_0(x)$ factors into distinct degree-1 terms over F and $\deg S_0(x) \in \{2n-3, 2n-2\}$. □

In contrast, the finite-field MRD construction of [3], [6] requires, for $\mu = 2$ and $\tau = 1$, a list size of $(|F|^n - 1)/(|F| - 1)$, thereby growing unboundedly with $|F|$ for every fixed $n > 1$ [18]. To avoid such an unbounded growth for $\tau = 1$, we must therefore select $\mu \geq 3$ and incur a redundancy of $2n$ (instead of n when $\mu = 2$). The list size of $\mathsf{C}_F(n,2)$ for $\tau = 1$, on the other hand, is bounded from above by an expression that does not depend on F, and the redundancy is $2n-1$.

Remark 2. When $|F| < 2n-3$ and $\mu = 2$, the exact value of $\mathsf{L}_\mathsf{C}(1)$ becomes generally more involved to analyze, as the attaining syndrome $S_0(x)$ may contain factors with multiplicities as well as irreducible factors of degree greater than 1. Table I lists, for $F = \mathrm{GF}(2)$ and several values of n, syndromes $S_0(x)$ that correspond to the largest possible list size, along with the value of that list size (for $n = 4, 5, 11, 12, 21, 22$, the maximizing syndromes are in fact unique). □

Proposition 3. *For $\mathsf{C} = \mathsf{C}_F(n,n)$ and every $\tau \in [1, n\rangle$,*

$$\mathsf{L}_\mathsf{C}(\tau) = \min\left\{ \left\lfloor \frac{n}{n-\tau} \right\rfloor, |F| \right\} \ . \tag{11}$$

Proof. We prove the proposition more generally for every one-dimensional array code over F of the form $\{\lambda \cdot A \ : \ \lambda \in F\}$, where A is nonsingular in $F^{n\times n}$. Given any $Y \in F^{n\times n}$, the matrix $E = Y - \lambda \cdot A$ in the coset of Y will have rank at most τ for the eigenvalues $\lambda \in F$ of $Y \cdot A^{-1}$ whose geometric (and therefore algebraic) multiplicities are at least $n-\tau$. Since $Y \cdot A^{-1}$ can have at most $\min\{\lfloor n/(n-\tau)\rfloor, |F|\}$ such eigenvalues, we get the right-hand side of (11) as an upper bound on $\mathsf{L}_\mathsf{C}(\tau)$. This bound is attained when $Y \cdot A^{-1}$ is a diagonal matrix whose main diagonal contains $\min\{\lfloor n/(n-\tau)\rfloor, |F|\}$ distinct elements of F, each appearing at least $n-\tau$ times. □

Recall that the finite-field MRD construction for even $\mu = n$ (which has dimension n) requires a list size of $|F|^{n/2} + 1$ already for $\tau = n/2$ [13].

Theorem 4. *For $\mathsf{C} = \mathsf{C}_F(n,\mu)$,*

$$\mathsf{L}_\mathsf{C}(\mu-1) \leq \prod_{i=0}^{n-\mu} \left(\binom{n+i}{\mu-1} \Big/ \binom{\mu-1+i}{\mu-1} \right) \ . \tag{12}$$

Proof. We assume that F is algebraically closed; otherwise, replace F with its algebraic closure K and C with the span over K of a basis of C. Denoting by $\mathcal{M} = \mathcal{M}_F(n,\mu)$ the set of $n\times n$ matrices of rank $< \mu$ over F, we show that for every $\Gamma \in F^{n\times n}$, the cardinality of the intersection of the coset $\mathsf{C}+\Gamma$ with $\mathcal{M}$ does not exceed the right-hand side of (12). Clearly, this holds when $\Gamma \in \mathsf{C}$, so we assume hereafter in the proof that $\Gamma \notin \mathsf{C}$ (and, in particular, that $\mu > 1$).

Fixing $\Gamma \in F^{n\times n} \setminus \mathsf{C}$, we denote by $\mathsf{G} = \mathsf{G}(\Gamma)$ the one-dimensional subspace $\{\lambda \cdot \Gamma \ : \ \lambda \in F\}$ of $F^{n\times n}$, and we let H be a hyperplane in $F^{n\times n}$ such that $\mathsf{C} \subseteq \mathsf{H}$ yet $\mathsf{G} \not\subseteq \mathsf{H}$ (e.g., H is the dual space of the one-dimensional space spanned by a

TABLE I
VALUES OF $\mathsf{L}_{\mathsf{C}}(1)$ WHEN $\mathcal{C}_m$ IS A SINGLE-PARITY CODE OVER $F = \mathrm{GF}(2)$ AND ATTAINING SYNDROMES FOR SEVERAL VALUES OF n.

n	$S_0(x)$			$\deg S_0(x)$	$\mathsf{L}_{\mathsf{C}}(1)$
3	x^2+x	$=$	$x(x+1)$	2	4
4	x^4+x^2	$=$	$x^2(x+1)^2$	4	7
5	$(x^4+x)(x^2+x)$	$=$	$x^2(x+1)^2(x^2+x+1)$	6	12
6	$(x^4+x)(x^2+x)x$	$=$	$x^3(x+1)^2(x^2+x+1)$	7	18
7	$(x^4+x^2)(x^4+x)$	$=$	$x^3(x+1)^3(x^2+x+1)$	8	26
8	$(x^4+x^2)(x^4+x)(x^2+x)$	$=$	$x^4(x+1)^4(x^2+x+1)$	10	36
9	$(x^4+x^2+x)(x^4+x^2)(x^4+x)$	$=$	$x^4(x+1)^3(x^2+x+1)(x^3+x+1)$	12	52
10	$(x^8+x)(x^4+x)(x^2+x)$	$=$	$x^3(x+1)^3(x^2+x+1)(x^3+x+1)(x^3+x^2+1)$	14	78
11	$(x^8+x)(x^4+x^2)(x^4+x)$	$=$	$x^4(x+1)^4(x^2+x+1)(x^3+x+1)(x^3+x^2+1)$	16	114
12	$(x^8+x^2)(x^8+x)(x^2+x)$	$=$	$x^4(x+1)^4(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$	18	160
13	$(x^8+x^2)(x^8+x)(x^2+x)x$	$=$	$x^5(x+1)^4(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$	19	216
14	$(x^8+x^2)(x^8+x)(x^4+x^2)$	$=$	$x^5(x+1)^5(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$	20	284
$\vdots$					
19	$(x^{16}+x)(x^8+x)(x^4+x)(x^2+x)x$	$=$	$x^5(x+1)^4(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$ $\cdot(x^4+x+1)(x^4+x^3+1)(x^4+x^3+x^2+x+1)$	31	1,288
20	$(x^{16}+x)(x^8+x)(x^4+x^2)(x^4+x)$	$=$	$x^5(x+1)^5(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$ $\cdot(x^4+x+1)(x^4+x^3+1)(x^4+x^3+x^2+x+1)$	32	1,744
21	$(x^{16}+x)(x^8+x)(x^4+x^2)(x^4+x)(x^2+x)$	$=$	$x^6(x+1)^6(x^2+x+1)^2(x^3+x+1)(x^3+x^2+1)$ $\cdot(x^4+x+1)(x^4+x^3+1)(x^4+x^3+x^2+x+1)$	34	2,294
22	$(x^{16}+x)(x^8+x^2)(x^8+x)(x^4+x^2)$	$=$	$x^6(x+1)^6(x^2+x+1)^3(x^3+x+1)(x^3+x^2+1)$ $\cdot(x^4+x+1)(x^4+x^3+1)(x^4+x^3+x^2+x+1)$	36	2,954

row vector that is added to a parity-check matrix of the direct sum $\mathsf{C} \oplus \mathsf{G}$ to form a parity-check matrix of C). We have

$$\mathsf{C} = (\mathsf{C} \oplus \mathsf{G}) \cap \mathsf{H} \, .$$

The sets C, $\mathsf{C} \oplus \mathsf{G}$, H, and $\mathcal{M}$ are irreducible algebraic varieties in the affine space $\mathbb{A}^{n^2}(F)$, each forming the common set of roots of a respective set of homogeneous multivariate polynomials over F. Hence, their projective counterparts, C', $(\mathsf{C} \oplus \mathsf{G})'$, H', and $\mathcal{M}'$ (obtained by removing the all-zero matrix and collapsing, for a nonzero matrix A in the respective variety, the set $\{\lambda \cdot A \, : \, \lambda \in F^*\}$ into one element) are well-defined projective irreducible varieties in the projective space $\mathbb{P}^{n^2-1}(F)$. Since $\mathcal{M}' \cap (\mathsf{C} \oplus \mathsf{G})' \cap \mathsf{H}' = \mathcal{M}' \cap \mathsf{C}'$ is empty, it follows from [9, p. 48, Thm. 7.2] that

$$\dim(\mathcal{M}' \cap (\mathsf{C} \oplus \mathsf{G})') + \dim \mathsf{H}' - (n^2 - 1) < 0 \, ,$$

which, with $\dim \mathsf{H}' = n^2 - 2$, implies that $\mathcal{M}' \cap (\mathsf{C} \oplus \mathsf{G})'$ is finite.

The degree of $\mathcal{M}'$, as a variety, is known to equal the right-hand size of (12) [5, §12.4.1], [8, pp. 243–244].[2] Since $(\mathsf{C} \oplus \mathsf{G})'$, being a linear projective space, has degree 1, it follows from Bézout's theorem that the intersection $\mathcal{M}' \cap (\mathsf{C} \oplus \mathsf{G})'$ cannot be larger that the right-hand side of (12) [8, Thms. 18.3 and 18.4]. The equality

$$|\mathcal{M} \cap (\mathsf{C} + \Gamma)| = |\mathcal{M}' \cap (\mathsf{C} \oplus \mathsf{G})'|$$

concludes the proof. □

The bound in Theorem 4 can grow exponentially with n^2 (e.g., when $\mu = n/2$); still, as shown for $\mu = 2$, the required redundancy for a given decoding radius and list size can be smaller for C than for other known schemes (albeit for very large fields). Specifically, given n and decoding radius τ, by selecting $\mu = \tau + 1$ the code C has redundancy $2\tau n - \tau^2$ and list size as in Theorem 4. In contrast, it is known that for many values of $\tau \leq n/5$ and for sufficiently large finite fields F, such a list size can be achieved with the finite-field MRD construction, only when $\mu \geq 2\tau + 1$ [13], thereby incurring a redundancy of $2\tau n$ (see also [4], [11], and references therein). It is worth noting that the parameters $(k{=}(n{-}\mu{+}1)^2, \tau{=}\mu{-}1)$ correspond to a point $(k/n^2, \tau/n)$ which lies right on a threshold curve of rate versus relative decoding radius τ/n: below (any fixed positive margin of) the curve most random rank-metric codes over $F = \mathrm{GF}(q)$ are list decodable with a fixed list size, while above (any fixed positive margin of) the curve the list size grows exponentially with $n^2 \log q$ [4].

The bound (12) is generally not tight for fields that are not algebraically closed. E.g., for $n = 10$ and $\mu = 2$, the bound coincides with (9) and equals $48,620$, yet for the construction of Section III-E over $F = \mathrm{GF}(2)$ the actual list size turns out to be 78 (see Table I). Computing upper bounds on $\mathsf{L}_{\mathsf{C}}(\tau)$ for $\mu/2 \leq \tau < \mu - 1$ is left for future work, and so is finding a respective list decoding algorithm (which is efficient in the value of $\mathsf{L}_{\mathsf{C}}(\tau)$).

[2] This generalizes to $\ell \times n$ matrices by changing any one of the two instances of n in (12) to ℓ (the resulting expression is symmetric in ℓ and n).

REFERENCES

[1] D. Augot, P. Loidreau, G. Robert, "Rank metric and Gabidulin codes in characteristic zero," *Proc. IEEE Int'l Symp. Inf. Theory (ISIT)*, Istanbul, Turkey (2013), 509–513.

[2] E.J. Candès, B. Recht, "Exact matrix completion via convex optimization," *Found. Comput. Math.,* 9 (2009), 717-772.

[3] P. Delsarte, "Bilinear forms over a finite field, with applications to coding theory," *J. Comb. Theory A,* 25 (1978), 226–241.

[4] Y. Ding, "On list-decodability of random rank metric codes and subspace codes," *IEEE Trans. Inf. Theory*, 61 (2015), 51–59.

[5] D. Eisenbud, J. Harris, *3264 and All That, A Second Course in Algebraic Geometry,* Cambridge University Press, Cambridge, UK, 2016.

[6] E.M. Gabidulin, "Theory of codes with maximum rank distance," *Probl. Inform. Transm.,* 21 (1985), 1–12.

[7] M. Gadouleau, Z. Yan, "Complexity of decoding Gabidulin codes," *Proc. 42nd Annual Conf. Inf. Sciences and Systems (CISS 2008)*, Princeton, New Jersey (Mar. 2008), 1081–1085.

[8] J. Harris, *Algebraic Geometry, A First Course,* Springer, New York, 1992.

[9] E. Hartshorne, *Algebraic Geometry,* Springer, New York, 1977.
[10] R. Kötter, F.R. Kschischang, "Coding for errors and erasures in random network coding," *IEEE Trans. Inf. Theory*, 54 (2008), 3579–3591.
[11] S. Liu, C. Xing, C. Yuan, "List decodability of random subcodes of Gabidulin codes," *IEEE Trans. Inf. Theory*, 63 (2017), 159–163.
[12] S. Müelich, S. Puchinger, D. Mödinger, M. Bossert, "An alternative decoding method for Gabidulin codes in characteristic zero," *Proc. 2016 IEEE Int'l Symp. Inf. Theory (ISIT 2016)*, Barcelona, Spain (2016), 2549–2553.
[13] N. Raviv, A. Wachter-Zeh, "Some Gabidulin codes cannot be list decoded efficiently at any radius," *IEEE Trans. Inf. Theory*, 62 (2016), 1605–1615.
[14] R.M. Roth, "Maximum-rank array codes and their application to criss-cross error correction," *IEEE Trans. Inf. Theory*, 37 (1991), 328–336.
[15] R.M. Roth, "Tensor codes for the rank metric," *IEEE Trans. Inf. Theory*, 42 (1996), 2146–2157.
[16] D. Silva, F.R. Kschischang, "Fast encoding and decoding of Gabildulin codes," *Proc. 2009 IEEE Int'l Symp. Inf. Theory (ISIT 2009)*, Seoul, Korea (June–July 2009), 2858–2862.
[17] D. Silva, F.R. Kschischang, R. Kötter, "A rank-metric approach to error control in random network coding," *IEEE Trans. Inf. Theory*, 54 (2008), 3951–3967.
[18] A. Wachter-Zeh, "Bounds on list decoding of rank-metric codes," *IEEE Trans. Inf. Theory*, 59 (2013), 7268–7277.
[19] A. Wachter-Zeh, V. Afanassiev, V. Sidorenko, "Fast decoding of Gabildulin codes," *Des. Codes Cryptogr.*, 66 (2013), 57–73.

Ron M. Roth (M'88–SM'97–F'03) received the B.Sc. degree in computer engineering, the M.Sc. in electrical engineering, and the D.Sc. in computer science from Technion—Israel Institute of Technology, Haifa, Israel, in 1980, 1984, and 1988, respectively. Since 1988 he has been with the Computer Science Department at Technion, where he now holds the General Yaakov Dori Chair in Engineering. During the academic years 1989–91 he was a Visiting Scientist at IBM Research Division, Almaden Research Center, San Jose, California, and during 1996–97, 2004–05, and 2011–2012 he was on sabbatical leave at Hewlett–Packard Laboratories, Palo Alto, California. He is the author of the book *Introduction to Coding Theory*, published by Cambridge University Press in 2006. Dr. Roth was an associate editor for coding theory in IEEE TRANSACTIONS ON INFORMATION THEORY from 1998 till 2001, and he is now serving as an associate editor in *SIAM Journal on Discrete Mathematics*. His research interests include coding theory, information theory, and their application to the theory of complexity.