Optimal repair of Reed-Solomon codes: Achieving the cut-set bound

Itzhak Tamo
Department of EE-Systems
Tel Aviv University
Tel Aviv, Israel
zactamo@gmail.com

Min Ye
Department of ECE/ISR
University of Maryland
College Park, MD, USA
yeemmi@gmail.com

Alexander Barg
Department of ECE/ISR
University of Maryland
College Park, MD, USA
abarg@umd.edu

***Abstract*—The repair problem for an (n, k) error-correcting code calls for recovery of an unavailable coordinate of the codeword by downloading as little information as possible from a subset of the remaining coordinates. Using the terminology motivated by coding in distributed storage, we attempt to repair a failed node by accessing information stored on d *helper nodes*, where $k \leqslant d \leqslant n-1$, and using as little *repair bandwidth* as possible to recover the lost information.**

By the so-called cut-set bound (Dimakis et al., 2010), the repair bandwidth of an $(n, k = n-r)$ MDS code using d helper nodes is at least $dl/(d+1-k)$, where l is the size of the node. A number of constructions of MDS array codes have been shown to meet this bound with equality. In a related but separate line of work, Guruswami and Wootters (2016) studied repair of Reed-Solomon (RS) codes, showing that it is possible to perform repair using a smaller bandwidth than under the trivial approach. At the same time, their work as well as follow-up papers stopped short of constructing RS codes (or any scalar MDS codes) that meet the cut-set bound with equality, which has been an open problem in coding theory.

In this work we present a solution to this problem, constructing RS codes of length n over the field of size $q^l, l = \exp((1+o(1))n \log n)$ that meet the cut-set bound. We also prove an almost matching lower bound on l, showing that super-exponential scaling is both necessary and sufficient for achieving the cut-set bound using linear repair schemes. More precisely, we prove that for scalar MDS codes (including the RS codes) to meet this bound, the sub-packetization l must satisfy $l \geqslant \exp((1+o(1))k \log k)$.

***Keywords*-Cut-set bound; Optimal sub-packetization; Repair bandwidth.**

I. Introduction

A. Minimum Storage Regenerating codes and optimal repair bandwidth

The amount of information produced has grown exponentially over the last decade, and large-scale storage systems are widely used to store the data. The problem that we consider is motivated by applications of codes in distributed storage wherein the data is written on a large number of physical storage nodes. Failure of an individual node renders a portion of the data inaccessible, and erasure-correcting codes are used to increase the reliability of the system. The repair task performed by the system relies on communication between individual nodes, and introduces new challenges in the code design. In particular, a new parameter that has a bearing on the overall efficiency of the system is the amount of data sent between the nodes in the process of repair.

To protect the information, we divide the original file into k information blocks and view each block as a single element of a finite field F or a vector over F. We encode the data by adding $r = n-k$ parity blocks (field symbols or vectors) and distributing the resulting n blocks across n storage nodes. In this paper we deal only with linear codes, so the parity blocks are formed as linear combinations of the information blocks over F. We use the notation (n, k) to refer to the length and dimension of a linear code. A well-known class of linear (n, k) Maximum Distance Separable (MDS) codes studied in this paper, has the favorable property that the original file can be recovered from the content stored on any k nodes, which provides the optimal tradeoff between failure tolerance and storage overhead.

Before proceeding further, we make a brief remark on the terminology used in the literature devoted to erasure correcting codes for distributed storage. The coordinates of the codeword are assumed to be stored on different nodes, and by extension are themselves referred to as nodes. In practice, single node failure is the most common scenario [1, Section 6.6], so we will be interested in the problem of designing codes that efficiently correct (repair) a single erasure (failed node). We assume that the data is encoded with a code $\mathcal{C}$ over a finite field F wherein each coordinate of the codeword is either an element of F or an l-dimensional vector over F, where $l > 1$. The latter construction, termed *array codes*, turns out to be better suited to the needs of the repair problem, as will be apparent in the later part of this section. To repair a failed node, the system needs to download the contents from some other nodes (*helper nodes*) of the codeword to the processor, and the total amount of the downloaded data is called the *repair bandwidth*. Coding solutions that support efficient repair are called *regenerating codes*, and they have been a focal point of current research in coding theory following their introduction in [2].

A traditional solution to recover a single node failure in an MDS-coded system is to download the content stored on any k nodes. The MDS property guarantees that we can recover the whole file, so we can also recover any single node failure. However, this method is far from efficient in

The research of I. Tamo is supported in part by ISF grant no. 1030/15 and the NSF-BSF grant no. 2015814. The research of M. Ye and A. Barg is supported by NSF grants CCF1422955 and CCF1618603. A. Barg is also with Inst. for Problems of Information Transmission, Russian Academy of Sciences, Moscow, Russia.

DOI 10.1109/FOCS.2017.28

the sense that the repair bandwidth that it requires is much larger than is needed for the repair of a single node. Indeed, by a rather counter-intuitive result of [2] it is possible to save on the repair bandwidth by contacting $d > k$ helper nodes, and the maximum savings are attained when d is the largest possible value, namely $d = n - 1$.

More specifically, suppose that an (n, k) MDS-coded system attempts to repair a failed node by connecting to d helper nodes. In this case, as shown in [2], the total amount of information that needs to be downloaded to complete the repair task is at least $dl/(d+1-k)$, where l is the size of each node. This lower bound on the repair bandwidth is called the *cut-set bound* since it is obtained from the cut-set bound in network information theory [3]. Given $k < d \leqslant n - 1$, an (n, k) MDS code achieving the cut-set bound for the repair of any single failed node from any d helper nodes is called an (n, k) *minimum storage regenerating* (MSR) code with *repair degree* d [2].

The definition of MSR codes, given above in an informal way, will be formalized for a particular subclass of codes known as *MDS array codes*. An (n, k) MDS array code $\mathcal{C}$ with *sub-packetization* l over a finite field F is formed of k information nodes and $r = n - k$ parity nodes, where every node is a column vector of length l over F (so $\dim_F(\mathcal{C}) = kl$). The MDS property requires that any k nodes of $\mathcal{C}$ suffice to recover the remaining r nodes of the codeword. Array codes are also called *vector codes*, while code families more common to coding theory (such as Reed-Solomon (RS) codes and others) are called *scalar codes* in the literature. Clearly, scalar codes correspond to the case $l = 1$ of the above definition.

Definition 1 (Repair bandwidth)**.** *Let $\mathcal{C}$ be an (n, k) MDS array code with sub-packetization l over a finite field F. We write a codeword of $\mathcal{C}$ as $c = (c_1, \ldots, c_n)$. For $i \in \{1, \ldots, n\}$ and $\mathcal{R} \subseteq [n]\backslash\{i\}$ of cardinality $|\mathcal{R}| \geqslant k$, define $N(\mathcal{C}, i, \mathcal{R})$ as the smallest number of symbols of F one needs to download from the set of helper nodes $\{c_j : j \in \mathcal{R}\}$ in order to repair the failed node c_i. The* repair bandwidth *of the code $\mathcal{C}$ with d helper nodes equals*

$$\max_{i \in [n]} \max_{\mathcal{R} \subseteq [n]\backslash\{i\}, |\mathcal{R}| = d} N(\mathcal{C}, i, \mathcal{R}).$$

We note that the symbols downloaded to repair the node c_i can be some functions of the contents of the helper nodes $\{c_j, j \in \mathcal{R}\}$.

Definition 2 (Cut-set bound [2])**.** *Let $\mathcal{C}$ be an (n, k) MDS array code with sub-packetization l and let $k \leqslant d \leqslant n - 1$. For any $i \in [n]$ and any subset $\mathcal{R} \subseteq [n]\backslash\{i\}$ of size d we have the following inequality:*

$$N(\mathcal{C}, i, \mathcal{R}) \geqslant \frac{dl}{d + 1 - k}. \tag{1}$$

An (n, k) MDS array code with sub-packetization l achieving the cut-set bound (1) *for the repair of any single failed node from any d helper nodes is called an (n, k, d, l)* MSR array code.

Several constructions of MSR codes are available in the literature: See [4]–[8] for the high-rate regime where $k > n/2$, and [9] for the low-rate regime where $k \leqslant n/2$. Recently the concept of repair bandwidth was extended in [10] to the problem of correcting errors; this paper also presented explicit code constructions that support error correction under the minimum possible amount of information downloaded during the decoding process.

Due to the limited storage capacity of each node, we would like the sub-packetization l to be as small as possible. However, it is shown in [11] that for an $(n, k, d = n - 1, l)$ MSR array code, $l \geqslant \exp(\sqrt{k/(2r-1)})$ (i.e., l is exponential in n for fixed r and growing n).

B. Repair schemes for scalar linear MDS codes

While there has been much research into constructions and properties of MSR codes specifically designed for the repair task, it is also of interest to study the repair bandwidth of general families of MDS codes, for instance, RS codes. In [12], Shanmugam et al. proposed a framework for studying the repair bandwidth of a scalar linear (n, k) MDS code $\mathcal{C}$ over some finite field E (called the symbol field below). The idea of [12] is to "vectorize" the code construction by considering $\mathcal{C}$ as an array code over some subfield F of E. This approach provides a bridge between RS codes and MDS array codes, wherein the extension degree $l := [E : F]$ can be viewed as the value of sub-packetization. The code $\mathcal{C}$ is viewed as an (n, k) MDS array code with sub-packetization l, and the repair bandwidth is defined exactly in the same way as above. The cut-set bound (1) and the definition of MSR codes also apply to this setup.

In this paper we study repair of RS codes, focusing on linear repair schemes, i.e., we assume that the repair operations are linear over the field F. Guruswami and Wootters [13] gave a characterization for linear repair schemes of scalar linear MDS codes based on the framework in [12]. We will use this characterization to prove one of our main results, namely, a lower bound on the sub-packetization, so we recall it below. Let us start with the definition of the dual code.

Definition 3 (Dual code)**.** *The dual code of a linear code $\mathcal{C} \subseteq E^n$ is the linear subspace of E^n defined by*

$$\mathcal{C}^\perp = \Big\{x = (x_1, \ldots, x_n) \in E^n \Big| \sum_{i=1}^n x_i c_i = 0 \\ \forall c = (c_1, \ldots, c_n) \in \mathcal{C}\Big\}.$$

In the next theorem E is the degree-l extension of the field F. Viewing E as an l-dimensional vector space over F, we use the notation $\dim_F(a_1, a_2, \ldots, a_t)$ to refer to the dimension of the subspace spanned by the set $\{a_1, a_2, \ldots, a_t\} \subset E$ over F.

We will need a result from [13] which we state in the form that is suited to our needs.

Theorem 1 ([13])**.** *Let $\mathcal{C} \subseteq E^n$ be a scalar linear MDS code of length n. Let F be a subfield of E such that $[E :$

$F] = l$. For a given $i \in \{1, \dots, n\}$ the following statements are equivalent.

(1) There is a linear repair scheme of the node c_i over F such that the repair bandwidth $N(\mathcal{C}, i, [n]\backslash\{i\}) \leqslant b$.

(2) There is a subset of codewords $\mathcal{P}_i \subseteq \mathcal{C}^\perp$ with size $|\mathcal{P}_i| = l$ such that

$$\dim_F(\{x_i : x \in \mathcal{P}_i\}) = l,$$

and

$$b \geqslant \sum_{j \in [n]\backslash\{i\}} \dim_F(\{x_j : x \in \mathcal{P}_i\}).$$

In addition to this general linear repair scheme for scalar linear MDS codes, the authors of [13] also presented a specific repair scheme for a family of RS codes and further proved that (in some cases) the repair bandwidth of RS codes using this scheme is the smallest possible among all linear repair schemes and all scalar linear MDS codes with the same parameters. At the same time, the approach of [13] has some limitations. Namely, their repair scheme applies only for small sub-packetization $l = \log_{n/r} n$, and the optimality claim only holds for this specific sub-packetization value. At the same time, in order to achieve the cut-set bound, l needs to be exponentially large in n for a fixed value of r [11], so the repair bandwidth of this scheme is rather far from the bound. Subsequently, two of the present authors [14] used the general linear repair scheme in [13] to construct an explicit family of RS codes with asymptotically optimal repair bandwidth: the ratio between the actual repair bandwidth of the codes and the cut-set bound approaches 1 as the code length n goes to infinity.

In [13], there is one more restriction on the parameters of the RS codes, namely they achieve the smallest possible repair bandwidth only if the number of parities is of the form $r = q^s, (l-s)|l$. In [15], Dau and Milenkovic generalized the scheme in [13] and extended their results to all values of $s = 1, \dots, l-1$. The repair bandwidth attained in [15] is $(n-1)(l-s)$ symbols of F for $r \geqslant q^s$, and is the smallest possible whenever r is a power of q. In [16], Dau et al. extended the results of [13] to repair of multiple erasures.

To summarize the earlier work, constructions of RS codes (or any scalar MDS codes) that meet the cut-set bound have as yet been unknown, so the existence question of such codes has been an open problem. In this paper, we resolve this problem in the affirmative, presenting such a construction. We also prove a lower bound on the sub-packetization of scalar linear MDS codes that attain the cut-set bound with a linear repair scheme, showing that there is a penalty for the scalar case compared to MDS array codes.

C. Our Results

(1) **Explicit constructions of RS codes achieving the cut-set bound:** Given any n, k and $d, k \leqslant d \leqslant n-1$, we construct an (n, k) RS code over the field $E = \mathbb{F}_{q^l}$ that achieves the cut-set bound (1) when repairing *any* single failed node from *any* d helper nodes. As above, we view RS codes over E as vector codes over the subfield $F = \mathbb{F}_q$. The main novelty in our construction is the choice of the evaluation points for the code in such a way that their over F are distinct primes. As a result, the symbol field is an extension field of F with degree no smaller than the product of these distinct primes. For the actual repair we rely on the linear scheme proposed in [13] (this is essentially the only possible linear repair approach).
The value of sub-packetization l of our construction equals s times the product of the first n distinct primes in an arithmetic progression,

$$l = s \prod_{\substack{i=1 \\ p_i \equiv 1 \bmod s}}^{n} p_i,$$

where $s := d + 1 - k$. This product is a well-studied function in number theory, related to a classical arithmetic function $\psi(n, s, a)$ (which is essentially the sum of logarithms of the primes). The prime number theorem in arithmetic progressions (for instance, [17, p.121]) yields asymptotic estimates for l. In particular, for fixed s and large n, we have $l = e^{(1+o(1))n \log n}$.
In contrast, for the case $d = n-1$ (i.e., $s = r = n-k$), there exist MSR array codes that attain sub-packetization $l = r^{\lceil n/(r+1) \rceil}$ [18], which is the smallest known value among MSR codes[1]. So although this distinct prime structure allows us to achieve the cut-set bound, it makes us pay a penalty on the sub-packetization.

(2) **A lower bound on the sub-packetization of scalar MDS codes achieving the cut-set bound:** Surprisingly, we also show that the distinct prime structure discussed above is necessary for any scalar linear MDS code (not just the RS codes) to achieve the cut-set bound under linear repair. Namely, given d such that $k+1 \leqslant d \leqslant n-1$, we prove that for any (n, k) scalar linear MSR code with repair degree d, the sub-packetization l is bounded below by $l \geqslant \prod_{i=1}^{k-1} p_i$, where p_i is the i-th smallest prime. By the Prime Number Theorem [17], we obtain the lower asymptotic bound on l of the form $l \geqslant e^{(1+o(1))k \log k}$.

(3) **Main result:** In summary, we obtain the following results for the smallest possible sub-packetization of scalar linear MDS codes, including the RS codes, whose repair bandwidth achieves the cut-set bound.

Theorem 2. *Let $\mathcal{C}$ be an $(n, k = n-r)$ scalar linear MDS code over the field $E = \mathbb{F}_{q^l}$, and let d be an integer satisfying $k+1 \leqslant d \leqslant n-1$. Suppose that for any single failed node of $\mathcal{C}$ and any d helper nodes there is a linear repair scheme over $\mathbb{F}_q$ that uses the bandwidth $dl/(d+1-k)$ symbols of $\mathbb{F}_q$, i.e., it achieves the cut-set bound* (1)*. For a fixed $s = d+1-k$ and*

[1] The construction of [18] achieves the cut-set bound only for repair of systematic nodes, and gives $l = r^{\lceil k/(r+1) \rceil}$. Using the approach of [4], it is possible to modify the construction of [18] and to obtain an MSR code with $l = r^{\lceil n/(r+1) \rceil}$.

Table I: Tradeoff between repair bandwidth and sub-packetization

Code construction	Repair bandwidth	sub-packetization	achieving cut-set bound
Array codes			
$(n, k = n - r, n - 1, l)$ MSR array codes for $2k \leqslant (n+1)$, [9]	$\frac{(n-1)l}{r}$	$l = r$	Yes
$(n, k, n-1, l)$ MSR array codes (a modification of [18])	$\frac{(n-1)l}{r}$	$l = r^{\lceil n/(r+1) \rceil}$	Yes
$(n, k, n-1, l)$ MSR array codes [5]	$\frac{(n-1)l}{r}$	$l = r^{\lceil n/r \rceil}$	Yes
(n, k) MDS array codes with design parameter $t \geqslant 1$ [19]	$(1 + \frac{1}{t})\frac{(n-1)l}{r}$	$l = r^t$	No
Scalar codes			
(n, k) RS code [14]	$< \frac{(n+1)l}{r}$	$l = r^n$	No
(n, k) RS code [13]	$n - 1$	$l = \log_{n/r} n$	No
(n, k) RS code [15]	$(n-1)l(1 - \log_n r)$	$\log_q n$	No
(n, k) RS code (this paper)	$\frac{(n-1)l}{r}$	$l \approx n^n$	Yes

$n, k \to \infty$ *the following bounds on the smallest possible sub-packetization hold true:*

$$e^{(1+o(1))k \log k} \leqslant l \leqslant e^{(1+o(1))n \log n}. \qquad (2)$$

For large s, *we have* $l \leqslant s \prod_{i:p_i \equiv 1 \bmod s}^{n} p_i$, *where the product goes over the first* n *distinct primes in the arithmetic progression.*

Remark 1. *The upper bound on* l *can be made more explicit even for large* s, *and the answer depends on whether we accept the Generalized Riemann Hypothesis (if yes, we can still claim the bound* $l \leqslant \exp((1 + o(1))n \log n)$*).*

(4) **Discussion: Array codes and scalar codes** The lower bound in (2) is much larger than the sub-packetization of many known MSR array code constructions. To make the comparison between the repair parameters of scalar codes and array codes more clear, we summarize the tradeoff between the repair bandwidth and the sub-packetization of some known MDS code constructions in Table I. We only list papers considering the repair of a single node from all the remaining $n-1$ helper nodes. Moreover, in the table we limit ourselves to explicit code constructions, and do not list multiple existence results that appeared in recent years.

As discussed earlier, the constructions of [13], [15] have optimal repair bandwidth among all the RS codes with the same sub-packetization value as in these papers[2]. At the same time, these values are too small for the constructions of [13], [15] to achieve the cut-set bound. From the first three rows of the table one can clearly see that the achievable sub-packetization values for MSR array codes are much smaller than the lower bound for scalar linear MSR codes derived in this paper. This is to be expected since for array codes we only require the code to be linear over the "repair field," i.e., F, and not the symbol field E as in the case of scalar codes.

D. Organization of the paper

In Sec. II, we present a simple construction of RS codes that achieve the cut-set bound for some of the nodes. This construction is inferior to the more involved construction of Sec. III, but simple to follow, and already contains some of the main ideas of the later part, so we include it as a warm-up for the later results. In Sec. III, we present our main construction of RS codes that achieve the cut-set bound for the repair of any single node, proving the upper estimate in (2). In Sec. IV, we prove the lower bound on the sub-packetization of scalar linear MSR codes, finishing the proof of (2).

II. A simple construction

In this section we present a simple construction of RS codes that achieve the cut-set bound for the repair of certain nodes. We note that any (n, k) MDS code trivially allows repair that achieves the cut-set bound for $d = k$. We say that a node in an MDS code has a *nontrivial optimal repair scheme* if for a given $d > k$ it is possible to repair this node from any d helper nodes with repair bandwidth achieving the cut-set bound. The code family presented in this section is different from standard MSR codes in the sense that although

[2] Expressing the sub-packetization of the construction in [15] via n and r is difficult. The precise form of the result in [15] is as follows: for every $s < l$ and $r \geqslant q^s$, the authors construct repair schemes of RS codes of length $n = q^l$ with repair bandwidth $(n-1)(l-s)$. Moreover, if $r = q^s$, then the schemes proposed in [15] achieve the smallest possible repair bandwidth for codes with these parameters.

the repair bandwidth of our construction achieves the cut-set bound, the number of helper nodes depends on the node being repaired.

Let us first recall the definition of (generalized) Reed-Solomon codes.

Definition 4. *A* Generalized Reed-Solomon code $\mathrm{GRS}_F(n,k,\Omega,v) \subseteq F^n$ *of dimension* k *over* F *with evaluation points* $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\} \subseteq F$ *is the set of vectors*

$$\{(v_1 f(\omega_1), \dots, v_n f(\omega_n)) \in F^n : f \in F[x], \deg f \leqslant k-1\}$$

where $v = (v_1, \dots, v_n) \in (F^*)^n$ *are some nonzero elements. If* $v = (1, \dots, 1)$*, then the GRS code is called a Reed-Solomon code and is denoted as* $\mathrm{RS}_F(n,k,\Omega)$*.*

It is well known [20, p.304] that

$$(\mathrm{RS}_F(n,k,\Omega))^\perp = \mathrm{GRS}_F(n,n-k,\Omega,v) \qquad (3)$$

where $v_i = \prod_{j \neq i} (\omega_i - \omega_j)^{-1}, i = 1, \dots, n$*. (The dual of an RS code is a GRS code.)*

Denote by $\pi(t)$ the number of primes less than or equal to t. Let F be a finite field and let E be the extension of F of degree t. The trace function $\mathrm{tr}_{E/F} : E \to F$ is defined by

$$\mathrm{tr}_{E/F}(x) := x + x^{|F|} + x^{|F|^2} + \cdots + x^{|F|^{t-1}}.$$

In the next theorem we construct a special subfamily of RS codes. Our construction enables nontrivial repair of $\pi(r)$ nodes, which without loss of generality we take to be nodes $1, 2, \dots, \pi(r)$. Let $d_i, i = 1, 2, \dots, \pi(r)$ be the number of helper nodes used to repair the i-th node. We will take $d_i = p_i + k - 1$, where p_i is the i-th smallest prime number. The repair scheme presented below supports repair of node i by connecting to any d_i helper nodes and downloading a $\frac{1}{p_i}$-th proportion of information stored at each of these nodes. Since $p_i = d_i - k + 1$, this justifies the claim of achieving the cut-set bound for repair of a single node.

Theorem 3. *Let* $n \geqslant k$ *be two positive integers, and let* $r = n - k$*. There exists an* (n,k) *RS code over a field* E *such that* $\pi(r)$ *of its coordinates admit nontrivial optimal repair schemes.*

Proof: Let $m := \pi(r)$ and let $q \geqslant n - m$ be a prime power. Let E be the $\big(\prod_{i=1}^m p_i\big)$-th degree extension of the finite field $\mathbb{F}_q$.

Let $\alpha_i, i = 1, \dots, m$ be an element of order p_i over $\mathbb{F}_q$, so that $\mathbb{F}_{q^{p_i}} = \mathbb{F}_q(\alpha_i)$, where $\mathbb{F}_q(\alpha_i)$ denotes the field obtained by adjoining α_i to $\mathbb{F}_q$. It is clear that $E = \mathbb{F}_q(\alpha_1, \dots, \alpha_m)$. Define m subfields F_i of E by setting

$$F_i = \mathbb{F}_q(\alpha_j : j \neq i),$$

so that $E = F_i(\alpha_i)$ and $[E : F_i] = p_i$, $i = 1, \dots, m$. Let $\alpha_{m+1}, \dots, \alpha_n \in \mathbb{F}_q$ be arbitrary $n - m$ distinct elements of the field, and let $\Omega = \{\alpha_1, \alpha_2, \dots, \alpha_n\}$.

Let $\mathcal{C} = \mathrm{RS}_E(n,k,\Omega)$ be the RS code of dimension k with evaluation points Ω and let $\mathcal{C}^\perp$ be its dual code. We claim that for $i = 1, 2, \dots, m$, the i-th coordinate (node) of $\mathcal{C}$ can be optimally repaired from any d_i helper nodes, where

$$d_i = p_i + k - 1.$$

Let $i \in \{1, 2, \dots, m\}$ and let us show how to repair the ith node. Choose a subset of helper nodes $\mathcal{R}_i \subseteq [n] \backslash \{i\}, |\mathcal{R}_i| = d_i$, and note that since $p_i \leqslant r$, we have $d_i \leqslant n - 1$. Let $h(x)$ be the annihilator polynomial of the set $\{\alpha_j : j \in [n] \backslash (\mathcal{R}_i \cup \{i\})\}$, i.e.,

$$h(x) = \prod_{j \in [n] \backslash (\mathcal{R}_i \cup \{i\})} (x - \alpha_j). \qquad (4)$$

Since $\deg(h(x)) = n - k - p_i$, we have $\deg(x^s h(x)) < r$ for all $s = 0, 1, \dots, p_i - 1$. As a result, for all $s = 0, \dots, p_i - 1$, the vector

$$(v_1 \alpha_1^s h(\alpha_1), \dots, v_n \alpha_n^s h(\alpha_n)) \in \mathcal{C}^\perp, \qquad (5)$$

cf. (3). Let $c = (c_1, \dots, c_n) \in \mathcal{C}$ be a codeword. By (5) we have

$$\sum_{j=1}^n v_j h(\alpha_j) \alpha_j^s c_j = 0, \quad s = 0, \dots, p_i - 1.$$

Let $\mathrm{tr}_i := \mathrm{tr}_{E/F_i}$ denote the trace from E to F_i. We have

$$\sum_{j=1}^n \mathrm{tr}_i(v_j h(\alpha_j) \alpha_j^s c_j) = 0, \quad s = 0, \dots, p_i - 1.$$

Equivalently, we can write for each $s = 0, \dots, p_i - 1$

$$\begin{aligned} \mathrm{tr}_i(v_i h(\alpha_i) \alpha_i^s c_i) &= - \sum_{j \neq i} \mathrm{tr}_i(v_j h(\alpha_j) \alpha_j^s c_j) \\ &= - \sum_{j \in \mathcal{R}_i} \mathrm{tr}_i(v_j h(\alpha_j) \alpha_j^s c_j) \\ &= - \sum_{j \in \mathcal{R}_i} \alpha_j^s \, \mathrm{tr}_i(v_j h(\alpha_j) c_j), \end{aligned} \qquad (6)$$

where the second equality follows from (4) and the third follows because $\alpha_j \in F_i$ for all $j \neq i$ and tr_i is an F_i-linear map.

The information used to recover the value c_i (to repair the ith node) is comprised of the following d_i elements of F_i :

$$\mathrm{tr}_i(v_j h(\alpha_j) c_j), \ \ j \in \mathcal{R}_i.$$

Let us show that these elements indeed suffice. First, by (6), given these elements, we can calculate the values of $\mathrm{tr}_i(v_i h(\alpha_i) \alpha_i^s c_i)$ for all $s = 0, \dots, p_i - 1$. The mapping

$$\begin{aligned} &E \to F_i^{p_i} \\ &\gamma \mapsto \big(\mathrm{tr}_i\left(v_i h(\alpha_i)\gamma\right), \mathrm{tr}_i\left(v_i h(\alpha_i)\alpha_i\gamma\right), \dots, \\ &\qquad\qquad \mathrm{tr}_i\left(v_i h(\alpha_i)\alpha_i^{p_i-1}\gamma\right)\big) \end{aligned}$$

is in fact a bijection, which can be realized as follows. Since the set $\{1, \alpha_i, \dots, \alpha_i^{p_i-1}\}$ forms a basis of E over F_i and $v_i h(\alpha_i) \neq 0$, the set $\{v_i h(\alpha_i), v_i h(\alpha_i)\alpha_i, \dots, v_i h(\alpha_i)\alpha_i^{p_i-1}\}$ also forms a basis. Let $\{\theta_0, \theta_1, \dots, \theta_{p_i-1}\}$ be the dual basis of $\{v_i h(\alpha_i), v_i h(\alpha_i)\alpha_i, \dots, v_i h(\alpha_i)\alpha_i^{p_i-1}\}$, i.e.,

$$\mathrm{tr}_i(v_i h(\alpha_i)\alpha_i^s \theta_j) = \begin{cases} 0 & \text{if } s \neq j, \\ 1 & \text{if } s = j \end{cases} \quad \forall s, j \in \{0, 1, \dots, p_i - 1\}.$$

The value c_i can now be found as follows:

$$c_i = \sum_{s=0}^{p_i-1} \mathrm{tr}_i(v_i h(\alpha_i)\alpha_i^s c_i)\theta_s$$

(this is the essence of the repair scheme proposed in [13]).

The presented arguments constitute a linear repair scheme of the node $c_i, i = 1, \dots m$ over F_i. The information downloaded from each of the helper nodes consists of one element of F_i, or, in other words, the $(1/p_i)$th proportion of the contents of each node. This shows that node i admits nontrivial optimal repair. The proof is thereby complete. ■

Example 1. Take $q = 5$, $k = 3, r = 5$. We have $\pi(r) = 3$ and $p_1 = 2, p_2 = 3, p_3 = 5$. Let us construct an $(8, 3)$ RS code over the field $E = \mathbb{F}_{5^{30}}$, where the first 3 nodes admit nontrivial optimal repair schemes. Let α be a primitive element of E. Choose the set $\Omega = \{\alpha_1, \dots, \alpha_8\}$ as follows:

$$\alpha_1 = \alpha^{\frac{5^{30}-1}{5^2-1}}, \alpha_2 = \alpha^{\frac{5^{30}-1}{5^3-1}}, \alpha_3 = \alpha^{\frac{5^{30}-1}{5^5-1}},$$
$$\alpha_4 = 0, \alpha_5 = 1, \alpha_6 = 2, \alpha_7 = 3, \alpha_8 = 4.$$

The number of helper nodes for the first 3 nodes is $(d_1, d_2, d_3) = (4, 5, 7)$. It is easy to verify that for any subset $A \subseteq \{1, 2, 3\}$

$$\mathbb{F}_5(\alpha_i : i \in A) = \mathbb{F}_{m_A}, \text{ where } m_A = 5^{(\prod_{i \in A} p_i)}.$$

The code $\mathcal{C}$ constructed in the above proof is given by $\mathcal{C} = \mathrm{RS}_E(8, 3, \Omega)$. Let us address the task of repairing c_3 from all the remaining 7 helper nodes with repair bandwidth achieving the cut-set bound. Let $\mathcal{C}^\perp = \mathrm{GRS}_E(8, 5, \Omega, v)$, where $v = (v_1, \dots, v_8) \in (E^*)^8$. We download the value $\mathrm{tr}_{E/\mathbb{F}_{5^6}}(v_j c_j)$ from each helper node $c_j, j \neq 3$. Since $[E : \mathbb{F}_{5^6}] = p_3$, this amounts to downloading exactly a $1/p_3 = (1/5)$-th fraction of the information stored at each helper node, which achieves the cut-set bound. The value of c_3 can be found from the downloaded information using the following 5 equations:

$$\begin{aligned} \mathrm{tr}_{E/\mathbb{F}_{5^6}}(\alpha_3^s v_3 c_3) &= -\sum_{j \neq 3} \mathrm{tr}_{E/\mathbb{F}_{5^6}}(\alpha_j^s v_j c_j) \\ &= -\sum_{j \neq 3} \alpha_j^s \, \mathrm{tr}_{E/\mathbb{F}_{5^6}}(v_j c_j), \quad s = 0, \dots, 4. \end{aligned}$$

Indeed, the downloaded symbols suffice to recover the vector $(\mathrm{tr}_{E/\mathbb{F}_{5^6}}(\alpha_3^s v_3 c_3), s = 0, \dots, 4)$, and therefore also suffice to repair the symbol c_3.

III. A family of RS codes achieving the cut-set bound

In this section we develop the ideas discussed above and construct RS codes achieving the cut-set bound with nontrivial optimal repair of all nodes. More precisely, given any positive integers $k < d \leqslant n - 1$, we explicitly construct an (n, k) RS code $\mathcal{C}$ achieving the cut-set bound for the repair of *any* single node from *any* d helper nodes. In other words, $\mathcal{C}$ is an (n, k) MSR code with repair degree d.

Let $\mathbb{F}_p$ be a finite field of prime order (for simplicity we can take $p = 2$). Denote $s := d - k + 1$ and let $p_1, \dots, p_n$ be n distinct primes such that

$$p_i \equiv 1 \bmod s \quad \text{for all } i = 1, 2, \dots, n. \tag{7}$$

According to Dirichlet's theorem, there are infinitely many such primes. For $i = 1, \dots, n$, let α_i be an element of degree p_i over $\mathbb{F}_p$, i.e., $[\mathbb{F}_p(\alpha_i) : \mathbb{F}_p] = p_i$, and define

$$\mathbb{F} := \mathbb{F}_p(\alpha_1, \dots, \alpha_n). \tag{8}$$

Note that for any subset of indices $A \subseteq [n]$, the field $\mathbb{F}_p(\{\alpha_i : i \in A\})$ is an extension of $\mathbb{F}_p$ of degree $\prod_{i \in A} p_i$, and in particular, $\mathbb{F}$ has degree $\prod_{i=1}^n p_i$ over $\mathbb{F}_p$. Next, we define n distinct subfields F_i of the field $\mathbb{F}$ and one extension field $\mathbb{K}$ of $\mathbb{F}$.

1) For $i = 1, \dots, n$, define $F_i = \mathbb{F}_p(\{\alpha_j : j \neq i\})$. Note that $\mathbb{F} = F_i(\alpha_i)$ and $[\mathbb{F} : F_i] = p_i$.
2) The field $\mathbb{K}$ is defined to be the degree-s extension of the field $\mathbb{F}$, i.e. there exists an element $\beta \in \mathbb{K}$ of degree s over $\mathbb{F}$ such that $\mathbb{K} = \mathbb{F}(\beta)$. We also have $[\mathbb{K} : F_i] = sp_i$ for all i.

We are ready to construct a family of RS codes that can be optimally repaired for each node. The set $\alpha_1, \dots, \alpha_n$ serves as the set of evaluation points of the code.

The following theorem is the main result of this section.

Theorem 4. *Let k, n, d be any positive integers such that $k < d < n$. Let $\Omega = \{\alpha_1, \dots, \alpha_n\}$, where $\alpha_i, i = 1, \dots, n$ is an element of degree p_i over $\mathbb{F}_p$ and p_i is the ith smallest prime that satisfies* (7). *The code $\mathcal{C} := \mathrm{RS}_{\mathbb{K}}(n, k, \Omega)$ achieves the cut-set bound for the repair of any single node from any d helper nodes. In other words, $\mathcal{C}$ is an (n, k) MSR code with repair degree d.*

Remark 2. *The code constructions in this paper rely on the condition of the form $\alpha_i \notin \mathbb{F}_q(\alpha_j, j \neq i), i = 1, \dots, n$ (in this section we also require that the extension degree $[\mathbb{F} : F_i] \equiv 1$ mod $s, i = 1, \dots, n$). The most efficient way to accomplish this in terms of the value of sub-packetization l is to take the extension degrees to be the smallest (distinct) primes, and this is the underlying idea behind the code constructions presented in this paper.*

Proof: Our repair scheme of the i-th node is performed over the field F_i. More specifically, for every $i \in [n]$, we explicitly construct a vector space S_i over the field F_i such that

$$\dim_{F_i} S_i = p_i, \quad S_i + S_i\alpha_i + \dots + S_i\alpha_i^{s-1} = \mathbb{K}, \tag{9}$$

where $S_i\alpha := \{\gamma\alpha : \gamma \in S_i\}$, and the operation $+$ is the Minkowski sum of sets, $T_1 + T_2 := \{\gamma_1 + \gamma_2 : \gamma_1 \in T_1, \gamma_2 \in T_2\}$. Note that the sum in (9) is in fact a direct sum since the dimension of each summand is p_i, and $[\mathbb{K} : F_i] = sp_i$. We will describe a construction of S_i and prove that S_i satisfies (9) in Lemma 1 later in this section. For now let us assume that we have such vector spaces $S_i, i = 1, 2, \ldots, n$ and continue the proof of the theorem.

Suppose that we want to repair the i-th node from a subset $\mathcal{R} \subseteq [n] \backslash \{i\}$ of $|\mathcal{R}| = d$ helper nodes. Let $h(x)$ be the annihilator polynomial of the set $\{\alpha_j : j \in [n] \backslash (\mathcal{R} \cup \{i\})\}$, i.e.,

$$h(x) = \prod_{j \in [n] \backslash (\mathcal{R} \cup \{i\})} (x - \alpha_j). \tag{10}$$

By (3) the dual code of $\mathcal{C}$ is $\mathcal{C}^{\perp} = \mathrm{GRS}_{\mathbb{K}}(n, n-k, \Omega, v)$ where the coefficients $v = (v_1, \ldots, v_n) \in (\mathbb{K}^*)^n$ are nonzero. Clearly, $\deg(x^t h(x)) \leqslant s-1+n-(d+1) < n-k$ for all $t = 0, 1, \ldots, s-1$, so for any such t we have

$$(v_1 \alpha_1^t h(\alpha_1), \ldots, v_n \alpha_n^t h(\alpha_n)) \in \mathcal{C}^{\perp}. \tag{11}$$

These s dual codewords will be used to recover the i-th coordinate. Let $c = (c_1, \ldots, c_n) \in \mathcal{C}$ be a codeword, and let us construct a repair scheme for the coordinate (node) c_i using the values $\{c_j : j \in \mathcal{R}\}$. Rewrite (11) as follows:

$$\sum_{j=1}^{n} v_j \alpha_j^t h(\alpha_j) c_j = 0 \text{ for all } t = 0, \ldots, s-1. \tag{12}$$

Let $e_1, \ldots, e_{p_i}$ be an arbitrary basis of the subspace S_i over the field F_i. From (12) we obtain the following system of sp_i equations:

$$\sum_{j=1}^{n} e_m v_j \alpha_j^t h(\alpha_j) c_j = 0, \quad t = 0, \ldots, s-1; m = 1, \ldots, p_i.$$

Let $\mathrm{tr}_i := \mathrm{tr}_{\mathbb{K}/F_i}$ be the trace map to the subfield F_i. From the last set of equations we have, for all $t = 0, \ldots, s-1$ and all $m = 1, \ldots, p_i$,

$$\sum_{j=1}^{n} \mathrm{tr}_i(e_m v_j \alpha_j^t h(\alpha_j) c_j) = 0. \tag{13}$$

Arguing as in (6), let us write (13) in the following form:

$$\begin{aligned} \mathrm{tr}_i(e_m \alpha_i^t v_i h(\alpha_i) c_i) &= -\sum_{j \neq i} \mathrm{tr}_i(e_m v_j \alpha_j^t h(\alpha_j) c_j) \\ &= -\sum_{j \in \mathcal{R}} \mathrm{tr}_i(e_m v_j \alpha_j^t h(\alpha_j) c_j) \\ &= -\sum_{j \in \mathcal{R}} \alpha_j^t h(\alpha_j) \, \mathrm{tr}_i(e_m v_j c_j) \end{aligned} \tag{14}$$

for all $t = 0, \ldots, s-1$ and $m = 1, \ldots, p_i$, where the second equality follows from (10) and the third follows from the fact that the trace mapping tr_i is F_i-linear, and that $\alpha_j \in F_i$ for all $j \neq i$.

As before, to recover c_i, we download the following p_i symbols of F_i from each helper node $c_j, j \in \mathcal{R}$:

$$\mathrm{tr}_i(e_m v_j c_j) \text{ for } m = 1, \ldots, p_i. \tag{15}$$

These field elements suffice to recover the node c_i. Indeed, according to (14), we can calculate the values of $\mathrm{tr}_i(e_m \alpha_i^t v_i h(\alpha_i) c_i)$ for all $t = 0, \ldots, s-1$ and all $m = 1, \ldots, p_i$ from the set of elements in (15). By definition, $e_1, \ldots, e_{p_i}$ is a basis of the subspace S_i over the field F_i. According to (9), $\mathbb{K} = S_i + S_i \alpha_i + \cdots + S_i \alpha_i^{s-1}$. Therefore, the set $\{e_m \alpha_i^t : t = 0, \ldots, s-1; m = 1, \ldots, p_i\}$ forms a basis of $\mathbb{K}$ over F_i and so does the set $\{e_m \alpha_i^t v_i h(\alpha_i) : t = 0, \ldots, s-1; m = 1, \ldots, p_i\}$ (recall that $v_i \cdot h(\alpha_i) \neq 0$). Hence the mapping

$$\begin{aligned} &\mathbb{K} \to F_i^{sp_i} \\ &\gamma \mapsto (\mathrm{tr}_i(e_m \alpha_i^t v_i h(\alpha_i) \gamma), m = 1, \ldots, p_i; t = 0, \ldots, s-1) \end{aligned}$$

is a bijection. This means that c_i is uniquely determined by the set of values $\{\mathrm{tr}_i(e_m \alpha_i^t v_i h(\alpha_i) c_i), m = 1, \ldots, p_i; t = 0, \ldots, s-1\}$, validating our repair scheme.

It is also clear that the construction meets the cut-set bound. Indeed, $c_j \in \mathbb{K}$ for all j and $[\mathbb{K} : F_i] = sp_i$, so the amount of information required from each helper node (15) is exactly $(1/s)$th fraction of its contents.

This completes the proof of Theorem 4. ■

In the proof above we assumed the existence of vector spaces $S_i, i = 1, 2, \ldots, n$ that satisfy (9). In the next lemma we construct such a space and establish its properties.

For a vector space V over a field F and a set of vectors $A = (a_1, \ldots, a_l) \subset V$, let $\mathrm{Span}_F(A) = \{\sum_{i=1}^{l} \gamma_i a_i, \gamma_i \in F\}$ be the span of A over F.

Lemma 1. *Let β be a generating element of $\mathbb{K}$ over $\mathbb{F} = \mathbb{F}_p(\alpha_1, \ldots, \alpha_n)$. Given $i \in [n]$, define the following vector spaces over F_i:*

$$\begin{aligned} S_i^{(1)} &= \mathrm{Span}_{F_i}\big(\beta^u \alpha_i^{u+qs}, u = 0, 1, \ldots, s-1; \\ &\qquad\qquad q = 0, 1, \ldots, \tfrac{p_i - 1}{s} - 1\big) \\ S_i^{(2)} &= \mathrm{Span}_{F_i}\Big(\sum_{t=0}^{s-1} \beta^t \alpha_i^{p_i - 1}\Big) \\ S_i &= S_i^{(1)} + S_i^{(2)}. \end{aligned}$$

Then

$$\dim_{F_i} S_i = p_i, \quad S_i + S_i \alpha_i + \cdots + S_i \alpha_i^{s-1} = \mathbb{K}.$$

Proof: Let $K := S_i + S_i \alpha_i + \cdots + S_i \alpha_i^{s-1}$. If $K = \mathbb{K}$, then $\dim_{F_i} S_i = p_i$ easily follows. Indeed, by definition $\dim_{F_i} S_i \leqslant p_i$. On the other hand, $[\mathbb{K} : F_i] = sp_i$ and $K = \mathbb{K}$ together imply that $\dim_{F_i} S_i \geqslant p_i$.

Let us prove that $K = \mathbb{K}$. Clearly K is a vector space over F_i, and $K \subseteq \mathbb{K}$. Let us show the reverse inclusion, namely that $\mathbb{K} \subseteq K$. To prove this, recall that $\mathbb{K}$ is a vector space of dimension s over $\mathbb{F}$, and the set $1, \beta, \ldots, \beta^{s-1}$ forms a basis, i.e., $\mathbb{K} = \oplus_{u=0}^{s-1} \beta^u \mathbb{F}$. Thus, the lemma will be proved if we show that $\beta^u \mathbb{F} \subseteq K$ for all $u = 0, 1, \ldots, s-1$. To prove this inclusion we will use induction on u.

For the induction base, let $u = 0$. In this case, we have $\alpha_i^{qs} \in S_i^{(1)}$ for all $0 \leqslant q < \frac{p_i - 1}{s}$. Therefore $\alpha_i^{qs+j} \in S_i^{(1)} \alpha_i^j$ for all $0 \leqslant q < \frac{p_i - 1}{s}$. As a result, $\alpha_i^{qs+j} \in K$ for all $0 \leqslant q < \frac{p_i - 1}{s}$ and all $0 \leqslant j \leqslant s-1$. In other words,

$$\alpha_i^t \in K, \ t = 0, 1, \ldots, p_i - 2. \tag{16}$$

Next we show that also $\alpha_i^{p_i-1} \in K$. For every $t = 1, \dots, s-1$ we have $0 \leqslant \lfloor \frac{p_i-1-t}{s} \rfloor < \frac{p_i-1}{s}$. As a result,

$$\beta^t \alpha_i^{t+\lfloor \frac{p_i-1-t}{s} \rfloor s} \in S_i^{(1)},\ t = 1, \dots, s-1.$$

We obtain, for each $t = 1, \dots, s-1$,

$$\begin{aligned} \beta^t \alpha_i^{p_i-1} &= \beta^t \alpha_i^{t+\lfloor \frac{p_i-1-t}{s} \rfloor s} \alpha_i^{p_i-1-t-\lfloor \frac{p_i-1-t}{s} \rfloor s} \\ &\in S_i \alpha_i^{p_i-1-t-\lfloor \frac{p_i-1-t}{s} \rfloor s} \subseteq K. \end{aligned}$$

At the same time,

$$\sum_{t=0}^{s-1} \beta^t \alpha_i^{p_i-1} \in S_i^{(2)} \subseteq K.$$

The last two statements together imply that

$$\alpha_i^{p_i-1} = \sum_{t=0}^{s-1} \beta^t \alpha_i^{p_i-1} - \sum_{t=1}^{s-1} \beta^t \alpha_i^{p_i-1} \in K.$$

Combining this with (16), we conclude that $\alpha_i^t \in K$ for all $t = 0, 1, \dots, p_i - 1$. Recall that $1, \alpha_i, \dots, \alpha_i^{p_i-1}$ is a basis of $\mathbb{F}$ over F_i, and that K is a vector space over F_i, so $\mathbb{F} \subseteq K$. This establishes the induction base.

Now let us fix $u \geqslant 1$ and let us assume that $\beta^{u'}\mathbb{F} \subseteq K$ for all $u' < u$. To prove the induction step, we need to show that $\beta^u \mathbb{F} \subseteq K$. Mimicking the argument that led to (16), we can easily show that

$$\beta^u \alpha_i^{u+t} \in K,\ t = 0, 1, \dots, p_i - 2. \tag{17}$$

Let us show that (17) is also true for $t = p_i - 1$, i.e., that $\beta^u \alpha_i^{u+p_i-1} \in K$. For every $1 \leqslant t \leqslant s - 1 - u$, we have $0 \leqslant \lfloor \frac{p_i-1-t}{s} \rfloor < \frac{p_i-1}{s}$. As a result,

$$\beta^{u+t} \alpha_i^{u+t+\lfloor \frac{p_i-1-t}{s} \rfloor s} \in S_i^{(1)},\ t = 1, \dots, s-1-u.$$

Therefore, for all such t

$$\begin{aligned} \beta^{u+t} \alpha_i^{u+p_i-1} &= \beta^{u+t} \alpha_i^{u+t+\lfloor \frac{p_i-1-t}{s} \rfloor s} \alpha_i^{p_i-1-t-\lfloor \frac{p_i-1-t}{s} \rfloor s} \\ &\in S_i \alpha_i^{p_i-1-t-\lfloor \frac{p_i-1-t}{s} \rfloor s} \subseteq K. \end{aligned} \tag{18}$$

By the induction hypothesis, $\beta^{u'}\mathbb{F} \subseteq K$ for all $u' = 0, 1, \dots, u-1$. As a result,

$$\beta^{u'} \alpha_i^{u+p_i-1} \in K,\ u' = 0, 1, \dots, u-1. \tag{19}$$

At the same time,

$$\sum_{t=0}^{s-1} \beta^t \alpha_i^{u+p_i-1} = \Big(\sum_{t=0}^{s-1} \beta^t \alpha_i^{p_i-1} \Big) \alpha_i^u \in S_i^{(2)} \alpha_i^u \subseteq K. \tag{20}$$

Combining (18), (19) and (20), we obtain

$$\begin{aligned} \beta^u \alpha_i^{u+p_i-1} = \sum_{t=0}^{s-1} \beta^t \alpha_i^{u+p_i-1} &- \sum_{u'=0}^{u-1} \beta^{u'} \alpha_i^{u+p_i-1} \\ &- \sum_{t=1}^{s-1-u} \beta^{u+t} \alpha_i^{u+p_i-1} \in K. \end{aligned}$$

Now on account of (17) we can conclude that $\beta^u \alpha_i^{u+t} \in K$ for all $t = 0, 1, \dots, p_i - 1$. Therefore, $\beta^u \mathbb{F} \subseteq K$. This establishes the induction step and completes the proof of the lemma. ■

The value of sub-packetization of the constructed codes is given in the following obvious proposition.

Proposition 1. *The sub-packetization of our construction is $l = [\mathbb{K} : \mathbb{F}_p] = s \prod_{i=1}^n p_i$, where p_i's are the smallest n distinct primes satisfying* (7).

The proof follows immediately from the fact that the repair of the i-th coordinate is performed over the field F_i, so the repair field of our construction is $\cap_{i=1}^n F_i = \mathbb{F}_p$. To estimate the asymptotics of l for $n \to \infty$, recall that our discussion of Dirichlet's prime number theorem in Sec. I-C above implies that, for fixed s, $l = e^{(1+o(1))n \log n}$. This proves the upper bound in (2).

IV. A lower bound on the sub-packetization of scalar linear MSR codes

In this section we prove a lower bound on the sub-packetization value l of (n,k) scalar linear MSR codes, which implies that $l \geqslant e^{(1+o(1))k \log k}$. In contrast, for MSR array codes, a much smaller sub-packetization value $l = r^{\lceil n/(r+1) \rceil}$ is achievable [18]. This shows that limiting oneself to scalar linear codes necessarily leads to a much larger sub-packetization, and constructing such codes in real storage systems is even less feasible than their array code counterparts. The main result of this section is the following theorem:

Theorem 5. *Let $F = \mathbb{F}_q$ and $E = \mathbb{F}_{q^l}$ for a prime power q. Let d be an integer between $k+1$ and $n-1$. Let $\mathcal{C} \subseteq E^n$ be an (n,k) scalar linear MDS code with a linear repair scheme over F. Suppose that the repair bandwidth of the scheme achieves the cut-set bound with equality for the repair of any single node from any d helper nodes. Then the sub-packetization l is at least*

$$l \geqslant \prod_{i=1}^{k-1} p_i$$

where p_i is the i-th smallest prime.

As discussed above in Sec. I-C, this theorem implies the asymptotic lower bound $l \geqslant e^{(1+o(1))k \log k}$.

In the proof of Theorem 5, we will need the following auxiliary lemmas.

Lemma 2. (Subfield criterion [21, Theorem 2.6]) *Each subfield of the field $\mathbb{F}_{p^n}$ is of order p^m, where $m|n$. For every positive divisor m of n there exists a unique subfield of $\mathbb{F}_{p^n}$ that contains p^m elements.*

Lemma 3. *Let E be an extension field of $\mathbb{F}_q$ and let $\alpha_1, \dots, \alpha_n \in E$. Then*

$$[\mathbb{F}_q(\alpha_1, \dots, \alpha_n) : \mathbb{F}_q] = \operatorname{lcm}(d_1, \dots, d_n),$$

where $d_i = [\mathbb{F}_q(\alpha_i) : \mathbb{F}_q]$.

Proof: Obvious.

Lemma 4. *Let* $a_1, a_2, \dots, a_n$ *and* $b_1, b_2, \dots, b_n$ *be two sets of vectors of the same dimension over a field* F*, and let* A *and* B *denote their spans over* F*. Let* $c_i = a_i + b_i, i = 1, \dots, n$ *then*

$$\dim_F(c_1, \dots, c_n) \leqslant \dim A + \dim B. \tag{21}$$

The lemma follows immediately from the fact that, for any two subspaces A and B of a linear space,

$$\dim(A + B) + \dim(A \cap B) = \dim A + \dim B.$$

In the next lemma $\mathcal{S}_F(\cdot)$ refers to the row space of the matrix argument over the field F.

Lemma 5. *Let* E *be an extension of a finite field* F *of degree* l*. Let* $A = (a_{i,j})$ *be an* $m \times n$ *matrix over* E*. Then*

$$\dim(\mathcal{S}_F(A)) \leqslant \sum_{j=1}^{n} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}). \tag{22}$$

Moreover, if (22) *holds with equality, then for every* $\mathcal{J} \subseteq [n]$,

$$\dim(\mathcal{S}_F(A_{\mathcal{J}})) = \sum_{j \in \mathcal{J}} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}) \tag{23}$$

where $A_{\mathcal{J}}$ *is the restriction of* A *to the columns with indices in* $\mathcal{J}$.

Proof: Inequality (22) is an immediate consequence of Lemma 4. Indeed, suppose that $n = 2$ and view the ith row of A as the sum of two 2-dimensional vectors over E, namely $(a_{i,1}|0)$ and $(0|a_{i,2}), i = 1, \dots, m$; then (22) is the same as (21). The extension to $n > 2$ follows by straightforward induction.

Now let us prove the second part of the claim. Suppose that

$$\dim(\mathcal{S}_F(A)) = \sum_{j=1}^{n} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}).$$

Then for every $\mathcal{J} \subseteq [n]$,

$$\begin{aligned}&\sum_{j \in \mathcal{J}} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}) \\ &\quad + \sum_{j \in \mathcal{J}^c} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}) \\ &= \dim(\mathcal{S}_F(A)) \leqslant \dim(\mathcal{S}_F(A_{\mathcal{J}})) + \dim(\mathcal{S}_F(A_{\mathcal{J}^c})).\end{aligned}$$

But according to (22),

$$\dim(\mathcal{S}_F(A_{\mathcal{J}})) \leqslant \sum_{j \in \mathcal{J}} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}),$$

$$\dim(\mathcal{S}_F(A_{\mathcal{J}^c})) \leqslant \sum_{j \in \mathcal{J}^c} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}).$$

Therefore

$$\dim(\mathcal{S}_F(A_{\mathcal{J}})) = \sum_{j \in \mathcal{J}} \dim_F(a_{1,j}, a_{2,j}, \dots, a_{m,j}).$$

This completes the proof of the lemma. ∎

Now we are ready to prove Theorem 5.

Proof of Theorem 5: Let $\mathcal{C}$ be an (n, k) MSR code with repair degree d. By puncturing the code $\mathcal{C}$ to any $d + 1$ coordinates, we obtain a $(d + 1, k)$ MSR code with repair degree d. Therefore without loss of generality below we assume that $d = n - 1$.

Let $H = [M|I_r]$ be the parity-check matrix of the code $\mathcal{C}$ over E, written in systematic form, where M is an $r \times k$ matrix and I_r is the $r \times r$ identity matrix. Let h_{ij} be the entry of H in position (i, j). Since $\mathcal{C}$ is an MDS code, every square submatrix of M is invertible. In particular, every entry of M is nonzero, so without loss of generality we may assume that $h_{1,j} = 1, j = 1, 2, \dots, k$. Since $d \geqslant k + 1$, we also have $n \geqslant k + 2$, and therefore H contains at least two rows.

The theorem will follow from the following claim.

Claim 1. *For* $j = 1, \dots, k - 1$ *define* $\alpha_j := \frac{h_{2,j}}{h_{2,k}}$*. Then for every* $j = 1, \dots, k - 1$,

$$\alpha_j \notin \mathbb{F}_q\big(\{\alpha_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\}\}\big). \tag{24}$$

In other words, α_j *is not generated by the remaining* α_i*'s over* $\mathbb{F}_q$.

We first show that this claim indeed implies the theorem. Let $d_i = [\mathbb{F}_q(\alpha_i) : \mathbb{F}_q]$ be the degree of the field extension generated by α_i. We prove by contradiction that for all $j = 1, 2, \dots, k-1$, d_j does not divide $\mathrm{lcm}(d_i : i \in \{1, 2, \dots, k-1\} \backslash \{j\})$. Suppose the contrary, i.e., that there is a j such that $d_j | \mathrm{lcm}(d_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\})$. According to Lemma 3,

$$\begin{aligned}&[\mathbb{F}_q\big(\{\alpha_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\}\}\big) : \mathbb{F}_q] \\ &\qquad\qquad = \mathrm{lcm}(d_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\}).\end{aligned}$$

Then by Lemma 2, there is a subfield

$$F_j \subseteq \mathbb{F}_q\big(\{\alpha_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\}\}\big) \tag{25}$$

such that $[F_j : \mathbb{F}_q] = d_j$. Notice that $E = \mathbb{F}_{q^l}$ contains all $\alpha_u, u = 1, 2, \dots, k - 1$. So both F_j and $\mathbb{F}_q(\alpha_j)$ are subfields of E, and they have the same order q^{d_j}. Consequently, $\mathbb{F}_q(\alpha_j) = F_j$. Then from (25) we conclude that $\alpha_j \in \mathbb{F}_q\big(\{\alpha_i : i \in \{1, 2, \dots, k-1\} \backslash \{j\}\}\big)$, which contradicts (24). Thus, our assumption is wrong, and $d_j \ | \ \mathrm{lcm}(d_i : i \in \{1, 2, \dots, k - 1\} \backslash \{j\})$. As an immediate corollary,

$$\begin{aligned}l = [E : \mathbb{F}_q] &\geqslant [\mathbb{F}_q(\{\alpha_i : i = 1, \dots, k - 1\}) : \mathbb{F}_q] \\ &= \mathrm{lcm}(d_1, \dots, d_{k-1}) \geqslant \prod_{i=1}^{k-1} p_i.\end{aligned}$$

Thus we have shown that this claim indeed implies the theorem. Now let us prove the claim.

Proof of the Claim: Consider the repair of the j-th node of the code $\mathcal{C}$ for some $j \in \{1, 2, \dots, k-1\}$. Since $\mathcal{C}$ can be viewed as an $(n, k, n-1, l)$ MSR code with a linear repair scheme over $\mathbb{F}_q$, node c_j can be repaired by downloading $(n-1)l/r$ symbols of $\mathbb{F}_q$ from all the remaining nodes $\{c_i : i \in [n]\backslash\{j\}\}$, where $r = n - k$. Therefore by Theorem 1, there exist l codewords

$$(c_{t,1}, c_{t,2}, \dots, c_{t,n}) \in \mathcal{C}^{\perp}, t = 1, 2, \dots, l$$

such that

$$\dim_{\mathbb{F}_q}(c_{1,j}, c_{2,j}, \dots, c_{l,j}) = l, \text{ and} \tag{26}$$

$$\sum_{i \neq j} \dim_{\mathbb{F}_q}(c_{1,i}, c_{2,i}, \dots, c_{l,i}) = \frac{(n-1)l}{r}. \tag{27}$$

Since H is a generator matrix of $\mathcal{C}^{\perp}$, for each $t = 1, 2, \dots, l$ there is a column vector $b_t \in E^r$ such that $(c_{t,1}, c_{t,2}, \dots, c_{t,n}) = b_t^T H$. We define an $l \times r$ matrix B over the field E as $B = [b_1 b_2 \dots b_l]^T$. We claim that the $\mathbb{F}_q$-rank of the row space of B is l. Indeed, assume the contrary, then there exists a nonzero vector $w \in \mathbb{F}_q^l$ such that $wB = 0$. Therefore,

$$wBH = w \begin{bmatrix} c_{1,1} & c_{1,2} & \cdots & c_{1,n} \\ c_{2,1} & c_{2,2} & \cdots & c_{2,n} \\ \vdots & \vdots & \vdots & \vdots \\ c_{l,1} & c_{l,2} & \cdots & c_{l,n} \end{bmatrix} = 0.$$

This implies that $w(c_{1,j}, c_{2,j}, \dots, c_{l,j})^T = 0$, contradicting (26). Thus we conclude that B has l linearly independent rows over $\mathbb{F}_q$.

Now we want to show that there exists an $l \times l$ invertible matrix A over $\mathbb{F}_q$ such that the matrix AB is an $r \times r$ block-diagonal matrix $\text{Diag}(a_1, \dots, a_r)$, where each block a_i is formed of a column vector of length $\frac{l}{r}$. In other words, by performing elementary row operations over $\mathbb{F}_q$, B can be transformed into an $r \times r$ block-diagonal matrix $\text{Diag}(a_1, \dots, a_r)$. Indeed, for $i \in [n]$, let h_i be the i-th column of the matrix H, and define

$$t_i = \dim_{\mathbb{F}_q}(Bh_i) = \dim_{\mathbb{F}_q}(c_{1,i}, c_{2,i}, \dots, c_{l,i}).$$

By (27), we have

$$\sum_{i \neq j}^{n} t_i = \frac{(n-1)l}{r}. \tag{28}$$

Since H generates an (n, r) MDS code, for any subset of indices $\mathcal{J} \subseteq [n]$ of size $|\mathcal{J}| = r$, the matrix $H_{\mathcal{J}}$ is of full rank. Therefore, the $l \times r$ matrix $BH_{\mathcal{J}}$ satisfies the conditions

$$l = \dim(\mathcal{S}_{\mathbb{F}_q}(B)) = \dim(\mathcal{S}_{\mathbb{F}_q}(BH_{\mathcal{J}})) \leqslant \sum_{i \in \mathcal{J}} \dim_{\mathbb{F}_q}(Bh_i), \tag{29}$$

where the last inequality follows from Lemma 5. Summing both sides of (29) over all subsets $\mathcal{J} \subseteq [n]\backslash\{j\}$ of size $|\mathcal{J}| = r$, we obtain that

$$\begin{aligned} l\binom{n-1}{r} &\leqslant \sum_{\substack{\mathcal{J} \subseteq [n]\backslash\{j\} \\ |\mathcal{J}| = r}} \sum_{i \in \mathcal{J}} \dim_{\mathbb{F}_q}(Bh_i) \\ &= \binom{n-2}{r-1} \sum_{i \neq j} t_i \\ &\overset{(28)}{=} \binom{n-2}{r-1} \frac{(n-1)l}{r} \\ &= l\binom{n-1}{r}. \end{aligned} \tag{30}$$

This implies that the inequality above is in fact an equality, and therefore, on account of (29) for every subset $\mathcal{J} \subseteq [n]\backslash\{j\}, |\mathcal{J}| = r$ we have

$$l = \sum_{i \in \mathcal{J}} \dim_{\mathbb{F}_q}(Bh_i) = \sum_{i \in \mathcal{J}} t_i. \tag{31}$$

From (31) we obtain that for all $i \in [n]\backslash\{j\}$

$$\dim_{\mathbb{F}_q}(Bh_i) = t_i = l/r. \tag{32}$$

Moreover, since (29) holds with equality, we can use the second part of Lemma 5 to claim that, for $\mathcal{J} \subseteq [n]\backslash\{j\}$ of size $|\mathcal{J}| \leqslant r$,

$$\dim(\mathcal{S}_{\mathbb{F}_q}(BH_{\mathcal{J}})) = \sum_{i \in \mathcal{J}} \dim_{\mathbb{F}_q}(Bh_i) = \frac{|\mathcal{J}|l}{r}. \tag{33}$$

Let us take $\mathcal{J}$ to be a subset of $\{k+1, k+2, \dots, n\}$. Since the last r columns of H form an identity matrix, (33) becomes

$$\dim(\mathcal{S}_{\mathbb{F}_q}(B_{\mathcal{J}})) = \frac{|\mathcal{J}|l}{r} \text{ for all } \mathcal{J} \subseteq [r] \text{ with size } |\mathcal{J}| \leqslant r. \tag{34}$$

Now we are ready to prove that by performing elementary row operations over $\mathbb{F}_q$, B can be transformed into an $r \times r$ block diagonal matrix $\text{Diag}(a_1, \dots, a_r)$, where each block a_i is a single column vector of length $\frac{l}{r}$. We proceed by induction. More specifically, we prove that for $i = 1, 2, \dots, r$, we can use elementary row operations over $\mathbb{F}_q$ to transform the first i columns of B into the following form:

$$\begin{bmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & a_i \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix},$$

where each $\mathbf{0}$ in the last row of the above matrix is a column vector of length $l(1 - \frac{i}{r})$.

Let $i = 1$. According to (34), each column of B has dimension l/r over $\mathbb{F}_q$. Thus the induction base holds trivially. Now assume that there is an $l \times l$ invertible matrix A over $\mathbb{F}_q$ such that

$$AB_{[i-1]} = \begin{bmatrix} a_1 & 0 & \dots & 0 \\ 0 & a_2 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & a_{i-1} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \end{bmatrix},$$

where each $\mathbf{0}$ in the last row of this matrix is a column vector of length $l(1-\frac{i-1}{r})$. Let us write the i-th column of AB as $(v_1, v_2, \dots, v_l)^T$. Since each column of B has dimension l/r over $\mathbb{F}_q$, $(v_1, v_2, \dots, v_l)^T$ also has dimension l/r over $\mathbb{F}_q$. Since the last $l(1-\frac{i-1}{r})$ rows of the matrix $AB_{[i-1]}$ are all zero, we can easily deduce that

$$\begin{aligned}\dim(\mathcal{S}_{\mathbb{F}_q}(AB_{[i]})) \leqslant & \frac{i-1}{r} l \\ & + \dim_{\mathbb{F}_q}(v_{(i-1)l/r+1}, v_{(i-1)l/r+2}, \dots, v_l).\end{aligned}$$

By (34), $\dim(\mathcal{S}_{\mathbb{F}_q}(AB_{[i]})) = \dim(\mathcal{S}_{\mathbb{F}_q}(B_{[i]})) = \frac{il}{r}$. As a result,

$$\begin{aligned}\dim_{\mathbb{F}_q}(v_{(i-1)l/r+1}, v_{(i-1)l/r+2}, \dots, v_l) \geqslant l/r \\ = \dim_{\mathbb{F}_q}(v_1, v_2, \dots, v_l).\end{aligned}$$

In other words, $(v_{(i-1)l/r+1}, v_{(i-1)l/r+2}, \dots, v_l)$ contains a basis of the set $(v_1, v_2, \dots, v_l)$ over $\mathbb{F}_q$. This implies that we can use elementary row operations on the matrix AB to eliminate all the nonzero entries v_m for $m \leqslant (i-1)l/r$, and thus obtain the desired block-diagonal structure for the first i columns. This establishes the induction step.

We conclude that there exists an $l \times l$ invertible matrix A over $\mathbb{F}_q$ such that $AB = \mathrm{Diag}(a_1, \dots, a_r)$, where each block a_i is a single column vector of length $\frac{l}{r}$. For $u \in [r]$, let A_u be the vector space spanned by the entries of a_u over $\mathbb{F}_q$. According to (32), for all $i \in [n]\backslash\{j\}$

$$\dim_{\mathbb{F}_q}(ABh_i) = \dim_{\mathbb{F}_q}(Bh_i) = l/r.$$

Since for every $i = 1, 2, \dots, n$

$$\begin{aligned}\dim_{\mathbb{F}_q}(ABh_i) &= \dim_{\mathbb{F}_q}(\mathrm{Diag}(a_1, \dots, a_r)h_i) \\ &= \dim_{\mathbb{F}_q}(A_1 h_{1,i} + \dots + A_r h_{r,i}),\end{aligned}$$

for all $i \in [n]\backslash\{j\}$ we have

$$\dim_{\mathbb{F}_q}(A_1 h_{1,i} + \dots + A_r h_{r,i}) = l/r.$$

Since each column of B has dimension l/r over $\mathbb{F}_q$, A_u also has dimension l/r over $\mathbb{F}_q$ for every $u \in [r]$. Recall that $h_{u,i} \neq 0$ for all $u \in [r]$ and all $i \in [k]$. Thus

$$\dim_{\mathbb{F}_q}(A_u h_{u,i}) = l/r = \dim_{\mathbb{F}_q}(A_1 h_{1,i} + \dots + A_r h_{r,i})$$

for all $u = 1, \dots, r$ and $i \in [k]\backslash\{j\}$. Therefore,

$$A_1 h_{1,i} = A_2 h_{2,i} = \dots = A_r h_{r,i} \text{ and all } i \in [k]\backslash\{j\}.$$

Since $h_{1,i} = 1$ for all $i = 1, 2, \dots, k$, we have

$$A_2 h_{2,i} = A_1 \text{ for all } i \in [k]\backslash\{j\}. \tag{35}$$

Equivalently,

$$A_2 \alpha_i = A_2 \text{ for all } i \in \{1, 2, \dots, k-1\}\backslash\{j\}.$$

By definition A_2 is a vector space over $\mathbb{F}_q$, so

$$A_2 \gamma = A_2 \text{ for all } \gamma \in \mathbb{F}_q(\{\alpha_i : i \in \{1, 2, \dots, k-1\}\backslash\{j\}\}). \tag{36}$$

On the other hand,

$$\begin{aligned}&\dim_{\mathbb{F}_q}(A_1 h_{1,j} + \dots + A_r h_{r,j}) \\ &\quad = \dim_{\mathbb{F}_q}(\mathrm{Diag}(a_1, \dots, a_r)h_j) = \dim_{\mathbb{F}_q}(ABh_j) \\ &\quad = \dim_{\mathbb{F}_q}(Bh_j) = \dim_{\mathbb{F}_q}\{c_{1,j}, c_{2,j}, \dots, c_{l,j}\} = l,\end{aligned} \tag{37}$$

while

$$\dim_{\mathbb{F}_q}(A_u h_{u,j}) = l/r, \quad u = 1, 2, \dots, r. \tag{38}$$

Equations (37) and (38) together imply that the vector spaces $A_1 h_{1,j}, A_2 h_{2,j}, \dots, A_r h_{r,j}$ are pairwise disjoint. In particular, $A_1 \cap A_2 h_{2,j} = \{0\}$. On account of (35), we therefore have $A_2 h_{2,k} \cap A_2 h_{2,j} = \{0\}$. This implies that $A_2 \alpha_j \neq A_2$. By (36), we conclude that $\alpha_j \notin \mathbb{F}_q(\{\alpha_i : i \in \{1, 2, \dots, k-1\}\backslash\{j\}\})$. This completes the proof of the claim.

References

[1] K. V. Rashmi, N. B. Shah, D. Gu, H. Kuang, D. Borthakur, and K. Ramchandran, "A Hitchhiker's guide to fast and efficient data reconstruction in erasure-coded data centers," in *ACM SIGCOMM Computer Communication Review*, vol. 44, no. 4. ACM, 2014, pp. 331–342.

[2] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran, "Network coding for distributed storage systems," *IEEE Trans. Inform. Theory*, vol. 56, no. 9, pp. 4539–4551, 2010.

[3] A. El Gamal, "On information flow in relay networks," in *Proc. IEEE National Telecom Conf.*, vol. 2, New Orleans, LA, 1981, pp. D4.1.1–D4.1.4.

[4] M. Ye and A. Barg, "Explicit constructions of high-rate MDS array codes with optimal repair bandwidth," *IEEE Transactions on Information Theory*, vol. 63, no. 4, pp. 2001–2014, 2017.

[5] ——, "Explicit constructions of optimal-access MDS codes with nearly optimal sub-packetization," 2016, arXiv:1605.08630.

[6] S. Goparaju, A. Fazeli, and A. Vardy, "Minimum storage regenerating codes for all parameters," *IEEE Transactions on Information Theory*, vol. 63, 2017.

[7] N. Raviv, N. Silberstein, and T. Etzion, "Constructions of high-rate minimum storage regenerating codes over small fields," *IEEE Transactions on Information Theory*, vol. 63, no. 4, pp. 2015–2038, 2017.

[8] I. Tamo, Z. Wang, and J. Bruck, "Zigzag codes: MDS array codes with optimal rebuilding," *IEEE Transactions on Information Theory*, vol. 59, no. 3, pp. 1597–1616, 2013.

[9] K. V. Rashmi, N. B. Shah, and P. V. Kumar, "Optimal exact-regenerating codes for distributed storage at the MSR and MBR points via a product-matrix construction," *IEEE Trans. Inform. Theory*, vol. 57, no. 8, pp. 5227–5239, 2011.

[10] I. Tamo, M. Ye, and A. Barg, "Fractional decoding: Error correction from partial information," 2017, arXiv:1701.06969.

[11] S. Goparaju, I. Tamo, and R. Calderbank, "An improved sub-packetization bound for minimum storage regenerating codes," *IEEE Transactions on Information Theory*, vol. 60, no. 5, pp. 2770–2779, 2014.

[12] K. Shanmugam, D. S. Papailiopoulos, A. G. Dimakis, and G. Caire, "A repair framework for scalar MDS codes," *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 5, pp. 998–1007, 2014.

[13] V. Guruswami and M. Wootters, "Repairing Reed-Solomon codes," *IEEE Transactions on Information Theory*, vol. 63, 2017.

[14] M. Ye and A. Barg, "Explicit constructions of MDS array codes and RS codes with optimal repair bandwidth," in *2016 IEEE International Symposium on Information Theory (ISIT)*. IEEE, 2016, pp. 1202–1206.

[15] H. Dau and O. Milenkovic, "Optimal repair schemes for some families of full-length Reed-Solomon codes," 2017, arXiv:1701.04120.

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

[17] H. Iwaniec and E. Kowalski, *Analytic number theory*. American Mathematical Society Providence, RI, 2004, vol. 53.

[18] Z. Wang, I. Tamo, and J. Bruck, "Explicit minimum storage regenerating codes," *IEEE Transactions on Information Theory*, vol. 62, no. 8, pp. 4466–4480, 2016.

[19] V. Guruswami and A. S. Rawat, "MDS code constructions with small sub-packetization and near-optimal repair bandwidth," in *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2017, pp. 2109–2122.

[20] F. J. MacWilliams and N. J. A. Sloane, *The Theory of Error-Correcting Codes*. Elsevier, 1977.

[21] R. Lidl and H. Niederreiter, *Introduction to Finite Fields and Their Applications*. Cambridge University Press, 1994.