On the I/O Costs in Repairing Short-Length Reed-Solomon Codes

Weiqi Li*, Hoang Dau†, Zhiying Wang*, Hamid Jafarkhani*, and Emanuele Viterbo‡
*University of California, Irvine, †RMIT University, ‡Monash University
Emails: *{weiqil4, zhiying, hamidj}@uci.edu, †sonhoang.dau@rmit.edu.au, ‡emanuele.viterbo@monash.edu

***Abstract*—Minimizing the *repair bandwidth*, i.e., the amount of information from the helper nodes needed for recovering the content of one failed node in an erasure-coded distributed storage system, has been the focus of many works in the literature. We investigate another important performance metric, namely the *I/O cost*, which specifies the amount of information that needs to be *read* by the helper nodes during the repair process of one failed node. We analyze the I/O costs of a few known repair schemes for Reed-Solomon codes of various lengths, in contrast to the previous works in this direction, which only studied the I/O costs in repairing full-length Reed-Solomon codes.**

I. Introduction

Reed-Solomon (RS) codes, although invented back in the sixties [1], still play a crucial role in major distributed storage systems (DSSs), including Google's Colossus, Quantcast File System, Facebook's f4, Yahoo Object Store, Baidu's Atlas, and Backblaze's Vaults (see [2, Tab. I]). The ubiquity of RS codes stems from their numerous advantages, such as optimal storage overhead, widest range of code parameters, and simple implementations. In the recently released version 3.0.0 of the Hadoop Distributed File System [3], the default erasure coding policy is based on a RS code.

There has been a considerable effort by the research community to optimize the *repair bandwidth* of RS codes [4]–[11]. Here, the repair bandwidth of a repair scheme refers to the amount of data to be transmitted from the helper nodes to the replacement node during the *repair process* that recovers the lost content of one failed node in a distributed storage system deploying the erasure code. Several extensions to the case of multiple erasures were also studied [2], [11]–[14]. The optimal repair bandwidth of RS codes is generally unknown, except for some full-length RS codes [5], [6], [8] and for RS codes with exponential large sub-packetizations [10], [11]. Constructions of RS codes and repair schemes that achieve a trade-off between the repair bandwidth and the sub-packetization were also investigated [15]–[17].

The *I/O cost*, which measures the total amount of data being *read* from the physical disks located at the helper nodes during the repair process of a failed node, is another important performance metric in repairing erasure codes. In the very first work along this line of research [18], the authors showed that known repair schemes of some families of full-length Reed-Solomon codes incur a *trivial* I/O cost, i.e., the entire file needs to be read from the system. Conversely, they also pointed out that for a particular family of RS codes with length $n = 2^\ell$ and two parities over $\mathbb{F}_{2^\ell}$, such a high I/O cost is necessary for achieving the optimal repair bandwidth. It was then a natural question to ask whether a trivial I/O cost is always required, or more generally, what is the lowest I/O cost possible. This question was answered for a family of RS codes of length $n = 2^\ell$ and two parities over $\mathbb{F}_{2^\ell}$ in the subsequent work [19], in which the authors constructed repair schemes that incur an *optimal* I/O cost of $(n-1)\ell - 2^{\ell-1}$ bits.

In this work, we first consider two-parity codes and $q = 2$. We generalize the construction for full-lengh RS codes [19] to RS codes of length $n = 2^m$, $m \leq \ell$, and obtain the I/O cost of $(n-1)\ell - 2^{m-1}$ bits. When $m \mid \ell$, we also provide a "lifting" transformation that transforms a repair scheme for a RS code of length q^m over $\mathbb{F}_{q^m}$ with I/O cost I into a repair scheme for a RS code of the same length q^m over the larger field $\mathbb{F}_{q^\ell}$ with I/O cost $\frac{\ell}{m}I$. Thus we obtain an improved I/O cost of $(n-1)\ell - \frac{\ell}{m}2^{m-1}$ bits. There is still a gap with our newly established lower bound of $(n-1)\ell-(\ell-m+1)2^{m-1}$ bits and it remains an open problem to construct repair schemes achieving optimal I/Os for these RS codes. Second, we consider short-length codes with arbitrary number of parities and field size. We explicitly determine the I/O costs of two existing repair schemes for short-length RS codes established in [10], [16], [17].

II. Preliminaries

Let $[n]$ denote the set $\{1, 2, \ldots, n\}$. Let $F = \mathbb{F}_q$ be the finite field of q elements, for some prime power q. Let $E = \mathbb{F}_{q^\ell}$ be an extension field of F, where $\ell \geq 1$, and let $E^* = E \backslash \{0\}$. We refer to the elements of E as *symbols* and the elements of F as *sub-symbols*. The field E may also be viewed as a vector space of dimension ℓ over F, i.e. $E \cong F^\ell$, and hence each symbol in E may be represented as a vector of length ℓ over F. More specifically, suppose $\mathcal{B} = \{\boldsymbol{\beta}_i\}_{i=1}^\ell$ is a basis of E over F, then any element $\boldsymbol{\alpha} \in E$ can be written as $\boldsymbol{\alpha} = \sum_{i=1}^\ell \alpha_i \boldsymbol{\beta}_i$. The unique vector $\phi_\mathcal{B}(\boldsymbol{\alpha}) = (\alpha_1, \ldots, \alpha_\ell) \in F^\ell$ is called the *vector representation* of $\boldsymbol{\alpha}$ w.r.t. the basis $\mathcal{B}$.

The (field) trace of a symbol $\boldsymbol{\alpha} \in E$ over F is defined to be $\mathsf{Tr}_{E/F}(\boldsymbol{\alpha}) = \sum_{i=0}^{\ell-1} \boldsymbol{\alpha}^{|F|^i}$ (the subscript E/F is often omitted). The support of a vector $\boldsymbol{u} = (u_1, \ldots, u_\ell)$, denoted $\mathsf{supp}(\boldsymbol{u})$, is the set $\{j \colon u_j \neq 0\}$. The (Hamming) weight of $\boldsymbol{u}$, denoted $\mathsf{wt}(\boldsymbol{u})$, is $|\mathsf{supp}(\boldsymbol{u})|$. The support of a set of vectors U is $\mathsf{supp}(U) \triangleq \cup_{\boldsymbol{u} \in U} \mathsf{supp}(\boldsymbol{u})$. A *linear* $[n, k]$ *code* $\mathcal{C}$ over E is an E-subspace of E^n of dimension k. Each element of a code is referred to as a *codeword*. The *dual* of a code $\mathcal{C}$, denoted $\mathcal{C}^\perp$, is the orthogonal complement of $\mathcal{C}$ in E^n and has dimension $r = n - k$. Elements of $\mathcal{C}^\perp$ are called *dual codewords*.

Definition 1. Let $E[x]$ denote the ring of polynomials over a finite field E. The Reed-Solomon code $\mathrm{RS}_E(\mathcal{A}, k) \subseteq E^n$ of dimension k with evaluation points $\mathcal{A} = \{\boldsymbol{\alpha}_j\}_{j=1}^n \subseteq E$ is defined as

$$\mathrm{RS}_E(\mathcal{A}, k) = \Big\{ \big(f(\boldsymbol{\alpha}_1), \ldots, f(\boldsymbol{\alpha}_n)\big) \colon f \in E[x], \deg(f) < k \Big\}.$$

The RS code is *full length* if $n = |E|$. When $\mathcal{A}$ forms an F-subspace of E, the dual of the RS code $\mathrm{RS}_E(\mathcal{A}, k)$ is another RS code $\mathrm{RS}_E(\mathcal{A}, n-k)$ (see Section IV-A).

Trace repair framework. First, note that each symbol in E can be recovered from ℓ independent traces. More precisely, given a basis $\{\boldsymbol{\beta}_i\}_{i=1}^\ell$ of E over F, any $\boldsymbol{\alpha} \in E$ can be uniquely determined given the values of $\mathsf{Tr}(\boldsymbol{\alpha}\boldsymbol{\beta}_i)$ for $i \in [\ell]$, i.e. $\boldsymbol{\alpha} = \sum_{i=1}^\ell \mathsf{Tr}(\boldsymbol{\alpha}\boldsymbol{\beta}_i)\boldsymbol{\beta}'_i$, where $\{\boldsymbol{\beta}'_i\}_{i=1}^\ell$ is the *dual (trace-orthogonal) basis* of $\{\boldsymbol{\beta}_i\}_{i=1}^\ell$ (see, e.g., [20, Ch. 2, Def. 2.30]).

Let $\mathcal{C}$ be an $[n, k]$ linear code over E and $\mathcal{C}^\perp$ be its dual. If $\boldsymbol{c} = (\boldsymbol{c}_1, \ldots, \boldsymbol{c}_n) \in \mathcal{C}$ and $\boldsymbol{g} = (\boldsymbol{g}_1, \ldots, \boldsymbol{g}_n) \in \mathcal{C}^\perp$ then $\boldsymbol{c} \cdot \boldsymbol{g}$

$= \sum_{j=1}^{n} \boldsymbol{c}_j \boldsymbol{g}_j = 0$. Note that when $\mathcal{A}$ is an F-subspace of E and $\mathcal{C} = \mathrm{RS}_E(\mathcal{A}, k)$, $\boldsymbol{g}$ is a codeword of $\mathrm{RS}_E(\mathcal{A}, n-k)$ and hence corresponds to a polynomial $g(x)$ of degree less than $n-k$. Suppose $\boldsymbol{c}_{j^*}$ is erased and needs to be recovered. In the trace repair framework, depending on j^*, we choose a set of ℓ dual codewords $\boldsymbol{g}^{(1)}, \ldots, \boldsymbol{g}^{(\ell)}$ such that $\dim_F\left(\{\boldsymbol{g}_{j^*}^{(i)}\}_{i=1}^{\ell}\right) = \ell$. Since trace is linear, we obtain the following ℓ equations

$$\mathsf{Tr}\big(\boldsymbol{g}_{j^*}^{(i)} \boldsymbol{c}_{j^*}\big) = -\sum_{j \neq j^*} \mathsf{Tr}\big(\boldsymbol{g}_j^{(i)} \boldsymbol{c}_j\big), \quad i \in [\ell]. \tag{1}$$

In order to recover $\boldsymbol{c}_{j^*}$, one needs to retrieve sufficient information from $\{\boldsymbol{c}_j\}_{j \neq j^*}$ to compute the right-hand sides of (1). We define, for every $j \in [n]$,

$$\mathcal{S}_{j \to j^*} \triangleq \mathsf{span}_F\Big(\Big\{ \boldsymbol{g}_j^{(1)}, \ldots, \boldsymbol{g}_j^{(\ell)} \Big\} \Big). \tag{2}$$

Then for each $j \neq j^*$, to determine $\mathsf{Tr}(\boldsymbol{g}_j^{(i)} \boldsymbol{c}_j)$ for all $i \in [\ell]$, it suffices to retrieve $\dim_F(\mathcal{S}_{j \to j^*})$ sub-symbols (in F) only. Indeed, suppose $\{\boldsymbol{g}_j^{(i_t)}\}_{t=1}^{s}$ is an F-basis of $\mathcal{S}_{j \to j^*}$, then by retrieving just s traces $\mathsf{Tr}(\boldsymbol{g}_j^{(i_1)} \boldsymbol{c}_j), \ldots, \mathsf{Tr}(\boldsymbol{g}_j^{(i_s)} \boldsymbol{c}_j)$ of $\boldsymbol{c}_j$, all other traces $\mathsf{Tr}(\boldsymbol{g}_j^{(i)} \boldsymbol{c}_j)$ can be computed as F-linear combinations of those s traces without any knowledge of $\boldsymbol{c}$. Finally, since $\{\boldsymbol{g}_{j^*}^{(i)}\}_{i=1}^{\ell}$ is F-linearly independent, $\boldsymbol{c}_{j^*}$ can be recovered from its ℓ corresponding traces on the left-hand side of (1). We refer to such a scheme as a *repair scheme based on* $\{\boldsymbol{g}^{(i)}\}_{i=1}^{\ell}$. It was known that this type of repair schemes includes every possible linear repair scheme for RS codes [5].

Lemma 1 (Guruswami-Wootters [5])**.** *Suppose $E = \mathbb{F}_{q^\ell}$, $F = \mathbb{F}_q$, $\mathcal{C}$ is an $[n,k]$ linear code over E and $\mathcal{C}^\perp$ is its dual. The repair scheme for $\boldsymbol{c}_{j^*}$ based on ℓ dual codewords $\boldsymbol{g}^{(1)}, \ldots, \boldsymbol{g}^{(\ell)}$, where $\dim_F\left(\mathcal{S}_{j^* \to j^*}\right) = \ell$, incurs a repair bandwidth of $\sum_{j \neq j^*} \dim_F(\mathcal{S}_{j \to j^*})$ sub-symbols in F, where $\mathcal{S}_{j \to j^*}$ is defined as in* (2)*.*

I/O Cost of a Repair Scheme Let $\mathcal{B} = \{\boldsymbol{\beta}_i\}_{i=1}^{\ell}$ be an F-basis of E. Then each element $\boldsymbol{\alpha} = \sum_{i=1}^{\ell} \alpha_i \boldsymbol{\beta}_i \in E$ can be represented by a vector $\phi(\boldsymbol{\alpha}) = (\alpha_1, \ldots, \alpha_\ell) \in F^\ell$ as defined earlier. We assume throughout this work that every node uses a fixed (but probably different) basis to represent and store the finite field elements. Another underlying assumption is that each sub-symbol α_i of $\boldsymbol{\alpha}$ can be read from the storage disk *separately* without accessing other sub-symbols. We first define the I/O cost of a function and then proceed to describe the I/O cost of a repair scheme.

Definition 2 ([18])**.** The (read) I/O cost of a function $\mathfrak{f}(\cdot)$ w.r.t. a basis $\mathcal{B}$ is the minimum number of sub-symbols of $\boldsymbol{\alpha} \in E$ needed for the computation of $\mathfrak{f}(\boldsymbol{\alpha})$. The I/O cost of a set of functions $\mathcal{F}$ is the minimum number of sub-symbols of $\boldsymbol{\alpha}$ needed for the computation of $\{\mathfrak{f}(\boldsymbol{\alpha})\colon \mathfrak{f} \in \mathcal{F}\}$.

Lemma 2 ([18])**.** *The following statements hold.*

(a) The I/O cost of the trace functional $\mathsf{Tr}^\gamma(\cdot)$, defined by $\mathsf{Tr}^\gamma(\boldsymbol{\alpha}) \triangleq \mathsf{Tr}(\boldsymbol{\gamma}\boldsymbol{\alpha})$, w.r.t. $\mathcal{B}$ is $\mathsf{wt}\big(\boldsymbol{w}^{\gamma,\mathcal{B}}\big)$, where

$$\boldsymbol{w}^{\gamma,\mathcal{B}} \triangleq \big(\mathsf{Tr}(\boldsymbol{\gamma}\boldsymbol{\beta}_1), \ldots, \mathsf{Tr}(\boldsymbol{\gamma}\boldsymbol{\beta}_\ell)\big) \in F^\ell. \tag{3}$$

(b) The I/O cost of the set of trace functionals $\{\mathsf{Tr}^\gamma(\cdot)\colon \boldsymbol{\gamma} \in \Gamma\}$ w.r.t. $\mathcal{B}$ is $|\cup_{\boldsymbol{\gamma} \in \Gamma} \mathsf{supp}(\boldsymbol{w}^{\gamma,\mathcal{B}})|$.

The I/O cost of the repair scheme for $\boldsymbol{c}_{j^*}$ based on a set of dual codewords $\{\boldsymbol{g}^{(i)}\}_{i=1}^{\ell}$ is the minimum number of sub-symbols of $\boldsymbol{c}_j$'s, $j \neq j^*$, needed in the computation of the right-hand side of (1). Note that each node may use a fixed but different basis $\mathcal{B}_j$, $j \in [n]$, to represent the finite field elements. The formal definition is given below.

Definition 3. The I/O cost of the repair scheme for $\boldsymbol{c}_{j^*}$ based on a set of dual codewords $\{\boldsymbol{g}^{(i)}\}_{i=1}^{\ell}$ w.r.t. a set of bases $\{\mathcal{B}_j\}_{j \neq j^*}$ is the sum of the I/O costs of the sets of trace functionals $\mathcal{F}_j = \big\{\mathsf{Tr}^{\boldsymbol{g}_j^{(i)}}(\cdot)\big\}_{i=1}^{\ell}$ w.r.t. $\mathcal{B}_j$, $j \in [n] \setminus \{j^*\}$.

Lemma 3 follows directly from Lemma 2 and Definition 3.

Lemma 3. *The I/O cost of the repair scheme for $\boldsymbol{c}_{j^*}$ based on a set of dual codewords $\{\boldsymbol{g}^{(i)}\}_{i=1}^{\ell}$ w.r.t. a set of bases $\{\mathcal{B}_j\}_{j \neq j^*}$ is $\sum_{j \in [n] \setminus \{j^*\}} \mathsf{nz}(\boldsymbol{W}_j)$, where $\mathsf{nz}(\boldsymbol{W}_j)$ specifies the number of nonzero columns in the $\ell \times \ell$ I/O matrix $\boldsymbol{W}_j$ defined as*

$$\boldsymbol{W}_j \triangleq \begin{pmatrix} \boldsymbol{w}^{\boldsymbol{g}_j^{(1)},\mathcal{B}_j} \\ \boldsymbol{w}^{\boldsymbol{g}_j^{(2)},\mathcal{B}_j} \\ \vdots \\ \boldsymbol{w}^{\boldsymbol{g}_j^{(\ell)},\mathcal{B}_j} \end{pmatrix} = \begin{pmatrix} \mathsf{Tr}(\boldsymbol{g}_j^{(1)}\boldsymbol{\beta}_{1,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(1)}\boldsymbol{\beta}_{i,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(1)}\boldsymbol{\beta}_{\ell,j}) \\ \mathsf{Tr}(\boldsymbol{g}_j^{(2)}\boldsymbol{\beta}_{1,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(2)}\boldsymbol{\beta}_{i,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(2)}\boldsymbol{\beta}_{\ell,j}) \\ \vdots \quad\quad \ddots \quad\quad \vdots \quad\quad \ddots \quad\quad \vdots \\ \mathsf{Tr}(\boldsymbol{g}_j^{(\ell)}\boldsymbol{\beta}_{1,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(\ell)}\boldsymbol{\beta}_{i,j}) \cdots \mathsf{Tr}(\boldsymbol{g}_j^{(\ell)}\boldsymbol{\beta}_{\ell,j}) \end{pmatrix},$$

where $\mathcal{B}_j = \{\boldsymbol{\beta}_{i,j}\}_{i=1}^{\ell}$. In this repair scheme the node storing $\boldsymbol{c}_j$ must read the i-th sub-symbol of $\boldsymbol{c}_j$ if and only if the i-th column of $\boldsymbol{W}_j$ is nonzero. Thus, $\mathsf{nz}(\boldsymbol{W}_j)$ specifies the I/O cost incurred at the node storing $\boldsymbol{c}_j$.

Lemma 3 can be understood as follows. If $\boldsymbol{c}_j = f(\boldsymbol{\alpha}_j)$ is read in full, ℓ sub-symbols will need to be accessed. However, it is possible that fewer than ℓ sub-symbols of $\boldsymbol{c}_j$ need to be read. Each *zero* column of $\boldsymbol{W}_j$ indicates a sub-symbol in the vector representation of $\boldsymbol{c}_j$ that does not need to be read.

III. On the I/O Cost for $[2^m, 2^m - 2]_{2^\ell}$ RS Codes

In this section we consider $[2^m, 2^m - 2]_{2^\ell}$ RS codes and assume that the nodes use a *common* F-basis of E. We first establish a lower bound on the I/O cost, thus generalizing [19, Thm. 1]. Then we present one construction generalizing [19, Cnstr. 1], and one based on a lifting transformation.

Theorem 1. *For a Reed-Solomon code $RS_{\mathbb{F}_{2^\ell}}(\mathcal{S}_m, 2^m - 2)$ where $\mathcal{S}_m$ is an m-dimensional $\mathbb{F}_2$-subspace of $\mathbb{F}_{2^\ell}$, the I/O cost (in bits) of an arbitrary linear repair scheme $\mathcal{R}$ always satisfies the following inequality.*

$$\mathsf{ic}(\mathcal{R}) \geq (n-1)\ell - (\ell - m + 1)2^{m-1}. \tag{4}$$

We first prove this theorem and then provide two repair schemes. Although not achieving the lower bound, they may provide a hint for constructing an I/O-optimal one.

As $r = 2$ and $\mathcal{S}_m$ is an m-dimensional subspace of $\mathbb{F}_{2^\ell}$, a dual codeword of $\mathscr{C}_m = \mathrm{RS}_{\mathbb{F}_{2^\ell}}(\mathcal{S}_m, 2^m - 2)$ can be obtained by evaluating a polynomial of degree at most one at all the elements of $\mathcal{S}_m$. A linear repair scheme, therefore, is based on a set of ℓ polynomials $g_i(x) = \boldsymbol{a}_i x + \boldsymbol{b}_i$, $\boldsymbol{a}_i, \boldsymbol{b}_i \in \mathbb{F}_{2^\ell}$, $i \in [\ell]$. Set $A = \{\boldsymbol{a}_i\}_{i=1}^{\ell}$ and $B = \{\boldsymbol{b}_i\}_{i=1}^{\ell}$. We say the repair scheme *is defined by A and B*. Moreover, by generalizing [18, Lem. 8] to the case when the evaluation points form a subspace, it suffices to consider repairing $\boldsymbol{c}_1 = f(0)$, the first codeword component corresponding to the evaluation point $\boldsymbol{\alpha}_1 = 0$. As a repair scheme for $\boldsymbol{c}_1$, it is required that $\mathsf{rank}_{\mathbb{F}_2}(\{g_i(0)\}_{i=1}^{\ell}) = \ell$. In other words, B *must be an $\mathbb{F}_2$-basis* of $\mathbb{F}_{2^\ell}$. We henceforth set $r_A \triangleq \mathsf{rank}_{\mathbb{F}_2}(A)$ and $A\boldsymbol{\gamma} + B \triangleq \{\boldsymbol{a}_i\boldsymbol{\gamma} + \boldsymbol{b}_i\}_{i=1}^{\ell}$.

It will later become clear in the proof of Theorem 1 that the set of "good" helpers w.r.t. a fixed basis element $\boldsymbol{\beta} = \boldsymbol{\beta}_i \in \mathcal{B}$, defined in Lemma 4, consists of the helpers where one bit of I/O can be saved in the repair scheme defined by A and B.

Lemma 4. *Let $\mathcal{S}_m$ be an m-dimensional $\mathbb{F}_2$-subspace of $\mathbb{F}_{2^\ell}$. Suppose $A = \{\boldsymbol{a}_i\}_{i=1}^{\ell} \subset \mathbb{F}_{2^\ell}$, $B = \{\boldsymbol{b}_i\}_{i=1}^{\ell}$ is an $\mathbb{F}_2$-basis of $\mathbb{F}_{2^\ell}$, and $\boldsymbol{\beta} \in \mathbb{F}_{2^\ell}^*$. We define $G_{A,B,\boldsymbol{\beta}}^m \triangleq \mathcal{S}_m \cap G_{A,B,\boldsymbol{\beta}}$ where*

$$G_{A,B,\boldsymbol{\beta}} \triangleq \{\boldsymbol{\gamma} \in \mathbb{F}_{2^\ell}\colon \mathsf{Tr}((\boldsymbol{a}_i\boldsymbol{\gamma} + \boldsymbol{b}_i)\boldsymbol{\beta}) = 0,\ \forall i \in [\ell]\}. \tag{5}$$

Then $G_{A,B,\boldsymbol{\beta}}^m$, which is called the set of "good" helpers w.r.t. $\boldsymbol{\beta}$, has size at most $2^{\min\{m-1, \ell - r_A\}}$.

Proof. According to [19, Lem. 4], $G_{A,B,\boldsymbol{\beta}}$ is the solution set of a non-homogeneous system of linear equations where the coefficient matrix has $\mathbb{F}_2$-rank r_A. Hence, it has size zero or $2^{\ell-r_A}$. Therefore, $|G^m_{A,B,\boldsymbol{\beta}}| \leq 2^{\ell-r_A}$. It remains to show that $|G^m_{A,B,\boldsymbol{\beta}}| \leq 2^{m-1}$.

As $G_{A,B,\boldsymbol{\beta}}$ is an affine $\mathbb{F}_2$-subspace of $\mathbb{F}_{2^\ell}$, we can write $G_{A,B,\boldsymbol{\beta}} = \boldsymbol{g}^* + G$ where G is an $(\ell - r_A)$-dimensional $\mathbb{F}_2$-subspace of $\mathbb{F}_{2^\ell}$ and $\boldsymbol{g}^* \notin G$. We aim to show that $|\mathcal{S}_m \cap (\boldsymbol{g}^* + G)| \leq 2^{m-1}$. Suppose, for the sake of contradiction, that there are $2^{m-1}+1$ distinct elements $\boldsymbol{s}_1, \ldots, \boldsymbol{s}_{2^{m-1}+1}$ in $\mathcal{S}_m \cap (\boldsymbol{g}^* + G)$. Set $\boldsymbol{t}_i = \boldsymbol{s}_1 - \boldsymbol{s}_{i+1}$, for $i \in [2^{m-1}]$. Then $\boldsymbol{t}_i \in \mathcal{S}_m \cap G$. As $G \cap (\boldsymbol{g}^* + G) = \varnothing$, the sets $\{\boldsymbol{s}_i\}_{i=1}^{2^{m-1}+1}$ and $\{\boldsymbol{t}_i\}_{i=1}^{2^{m-1}}$ are disjoint and both belong to $\mathcal{S}_m$. This leads to a contradiction because $|\mathcal{S}_m| = 2^m < (2^{m-1}+1) + 2^{m-1}$. ∎

As we shall see in the proof of Theorem 1, the set of "good" basis elements defined as in Lemma 5 determines how many $\boldsymbol{\beta}_i \in \mathcal{B}$ can give rise to the saving of *at least* one bit in I/O in a repair scheme defined by A and B.

Lemma 5. *Let $\mathcal{S}_m$ be an m-dimensional $\mathbb{F}_2$-subspace of $\mathbb{F}_{2^\ell}$. Given $A = \{\boldsymbol{a}_i\}_{i=1}^{\ell} \subset \mathbb{F}_{2^\ell}$ and $B = \{\boldsymbol{b}_i\}_{i=1}^{\ell}$ an $\mathbb{F}_2$-basis of $\mathbb{F}_{2^\ell}$, we define the set of "good" basis elements as*

$$G^m_{A,B} \triangleq \left\{\boldsymbol{\beta} \in \mathbb{F}^*_{2^\ell} \colon |G^m_{A,B,\boldsymbol{\beta}}| > 0\right\}, \tag{6}$$

where $G^m_{A,B,\boldsymbol{\beta}}$ is defined as in Lem. 4. Then $\mathsf{rank}_{\mathbb{F}_2}(G^m_{A,B}) \leq r_A$.

Proof. If we define $G_{A,B} \triangleq \left\{\boldsymbol{\beta} \in \mathbb{F}^*_{2^\ell} \colon |G_{A,B,\boldsymbol{\beta}}| > 0\right\}$, where $G_{A,B,\boldsymbol{\beta}}$ is defined as in Lemma 4, then according to [19, Lem. 7], $\mathsf{rank}_{\mathbb{F}_2}(G_{A,B}) \leq r_A$. Note that $G_{A,B}$ was defined slightly different in [19, Lem. 7], however, the two definitions are equivalent, as explained by the first sentence in the proof of that lemma. As $G^m_{A,B,\boldsymbol{\beta}} = \mathcal{S}_m \cap G_{A,B,\boldsymbol{\beta}} \subseteq G_{A,B,\boldsymbol{\beta}}$, we have $\mathsf{rank}_{\mathbb{F}_2}(G^m_{A,B}) \leq \mathsf{rank}_{\mathbb{F}_2}(G_{A,B}) \leq r_A$, as desired. ∎

We are now ready to prove Theorem 1. Its proof generalizes the proof of [19, Thm. 1].

Proof of Theorem 1. Consider a scheme $\mathcal{R}_{A,B}$ repairing the first codeword component $\boldsymbol{c}_1 = f(0)$ of $\mathscr{C}_m = \mathrm{RS}_{\mathbb{F}_{2^\ell}}(\mathcal{S}_m, 2^m - 2)$, defined by $A = \{\boldsymbol{a}_i\}_{i=1}^{\ell}$ and $B = \{\boldsymbol{b}_i\}_{i=1}^{\ell}$, where B is an $\mathbb{F}_2$-basis of $\mathbb{F}_{2^\ell}$. The corresponding dual codewords are obtained by evaluating the polynomials $g_i(x) = \boldsymbol{a}_i x + \boldsymbol{b}_i$, $i \in [\ell]$ at all $\boldsymbol{\gamma} \in \mathcal{S}_m$. Set $\mathcal{S}^*_m \triangleq \mathcal{S}_m \setminus \{\boldsymbol{0}\}$. By Lemma 3, the I/O cost of $\mathcal{R}_{A,B}$ w.r.t. a fixed basis $\mathcal{B}$ is

$$\mathsf{ic}_{\mathcal{B}}(\mathcal{R}_{A,B}) = \sum_{\boldsymbol{\gamma} \in \mathcal{S}^*_m} \mathsf{nz}(\boldsymbol{W}_{\boldsymbol{\gamma}}) = (n-1)\ell - \sum_{\boldsymbol{\gamma} \in \mathcal{S}^*_m} \mathsf{z}(\boldsymbol{W}_{\boldsymbol{\gamma}}),$$

where $\mathsf{nz}(\boldsymbol{W}_{\boldsymbol{\gamma}})$ and $\mathsf{z}(\boldsymbol{W}_{\boldsymbol{\gamma}})$ denote the counts of *nonzero columns* and *zero columns*, respectively, in the $\ell \times \ell$ matrix $\boldsymbol{W}_{\boldsymbol{\gamma}}$ whose rows are $\boldsymbol{w}^{\boldsymbol{a}_i\boldsymbol{\gamma}+\boldsymbol{b}_i,\mathcal{B}}$, $i \in [\ell]$, given as follows

$$\begin{pmatrix} \mathsf{Tr}((\boldsymbol{a}_1\boldsymbol{\gamma}+\boldsymbol{b}_1)\boldsymbol{\beta}_1) \cdots \mathsf{Tr}((\boldsymbol{a}_1\boldsymbol{\gamma}+\boldsymbol{b}_1)\boldsymbol{\beta}_i) \cdots \mathsf{Tr}((\boldsymbol{a}_1\boldsymbol{\gamma}+\boldsymbol{b}_1)\boldsymbol{\beta}_\ell) \\ \mathsf{Tr}((\boldsymbol{a}_2\boldsymbol{\gamma}+\boldsymbol{b}_2)\boldsymbol{\beta}_1) \cdots \mathsf{Tr}((\boldsymbol{a}_2\boldsymbol{\gamma}+\boldsymbol{b}_2)\boldsymbol{\beta}_i) \cdots \mathsf{Tr}((\boldsymbol{a}_2\boldsymbol{\gamma}+\boldsymbol{b}_2)\boldsymbol{\beta}_\ell) \\ \vdots \quad \ddots \quad \vdots \quad \ddots \quad \vdots \\ \mathsf{Tr}((\boldsymbol{a}_\ell\boldsymbol{\gamma}+\boldsymbol{b}_\ell)\boldsymbol{\beta}_1) \cdots \mathsf{Tr}((\boldsymbol{a}_\ell\boldsymbol{\gamma}+\boldsymbol{b}_\ell)\boldsymbol{\beta}_i) \cdots \mathsf{Tr}((\boldsymbol{a}_\ell\boldsymbol{\gamma}+\boldsymbol{b}_\ell)\boldsymbol{\beta}_\ell) \end{pmatrix}.$$

Each *zero* column of $\boldsymbol{W}_j$ indicates a bit in the vector representation of $\boldsymbol{c}_j$ that does not need to be read, therefore, leads to a *saving of one bit* in I/O at the j-th node. The more zero columns, the larger the saving, and hence the lower the I/O cost. Thus, we are interested in maximizing the total number of zero columns in all $\boldsymbol{W}_j, j \in [n] \backslash \{j^*\}$.

Instead of finding $\mathsf{z}(\boldsymbol{W}_{\boldsymbol{\gamma}})$ and sum that up over all $\boldsymbol{\gamma} \in \mathbb{F}^*_{2^m}$, we first fix an index i and count the number of zero i-th columns of $\boldsymbol{W}_{\boldsymbol{\gamma}}$ when $\boldsymbol{\gamma}$ varies over $\mathbb{F}^*_{2^m}$, and then sum that up over all $i \in [\ell]$. By Lemma 4, for each $\boldsymbol{\beta}_i \in \mathcal{B}$, the number of zero i-th columns in $\boldsymbol{W}_{\boldsymbol{\gamma}}$, $\boldsymbol{\gamma} \in \mathbb{F}^*_{2^m}$, is $|G^m_{A,B,\boldsymbol{\beta}_i}| \leq 2^{\min\{m-1,\ell-r_A\}}$. The set $G^m_{A,B,\boldsymbol{\beta}_i}$ consists of those "good" helpers $\boldsymbol{\gamma}$ where the i-th column of $\boldsymbol{W}_{\boldsymbol{\gamma}}$ is zero, corresponding to a saving of one bit in I/O when $\boldsymbol{c}_{\boldsymbol{\gamma}}$ is accessed. If $|G^m_{A,B,\boldsymbol{\beta}_i}| > 0$, $\boldsymbol{\beta}_i$ is a "good" basis element. The set of "good" $\boldsymbol{\beta}_i$ is denoted by $G^m_{A,B}$ as in Lemma 5. Furthermore, as $\mathsf{rank}_{\mathbb{F}_2}(G^m_{A,B}) \leq r_A$, the basis $\mathcal{B}$ contains at most r_A "good" $\boldsymbol{\beta}_i$. Clearly, if $\boldsymbol{\beta}_i$ is not "good", then there does not exist any $\boldsymbol{\gamma} \in \mathbb{F}_{2^\ell}$ such that the i-th column of $\boldsymbol{W}_{\boldsymbol{\gamma}}$ is zero. Thus,

$$\mathsf{ic}_{\mathcal{B}}(\mathcal{R}_{A,B}) \geq (n-1)\ell - \max_{0 \leq r_A \leq \ell} \{r_A 2^{\min\{m-1,\ell-r_A\}}\},$$

where the term $r_A 2^{\min\{m-1,\ell-r_A\}}$ corresponds to the maximum possible reduction/saving in I/O if $\mathsf{rank}_{\mathbb{F}_2}(A) = r_A$. The factor r_A accounts for the maximum number of "good" $\boldsymbol{\beta}_i \in \mathcal{B}$ while $2^{\min\{m-1,\ell-r_A\}}$ accounts for the maximum saving in I/O w.r.t. each of such "good" $\boldsymbol{\beta}_i$. It can be easily verified that $r_A = \ell - m + 1$ minimizes the right-hand side of the above inequality. That concludes the proof. ∎

According to the proof of Theorem 1, we should choose A of rank $\ell - m + 1$ to achieve the I/O cost specified in the lower bound. However, it seems challenging to design a repair scheme in that case. The following construction generalizes Construction I in [19] to RS codes of length 2^m over $\mathbb{F}_{2^\ell}$ and uses a set A of rank one. For a set $K \subseteq \mathbb{F}_{2^\ell}$ and an element $\boldsymbol{\beta} \in \mathbb{F}^*_{2^\ell}$, we use the notation $K/\boldsymbol{\beta}$ for the set $\{\boldsymbol{\kappa}/\boldsymbol{\beta} : \boldsymbol{\kappa} \in K\}$.

Construction I. Let $\mathcal{B} = \{\boldsymbol{\beta}_i\}_{i=1}^{\ell}$ be an $\mathbb{F}_2$-basis of $\mathbb{F}_{2^\ell}$ and fix an arbitrary element $\boldsymbol{\beta} \in \mathcal{B}$. Let K be the kernel of $\mathsf{Tr}_{\mathbb{F}_{2^\ell}/\mathbb{F}_2}(\cdot)$ and $\{\boldsymbol{h}_i\}_{i=1}^{\ell-1}$ be an $\mathbb{F}_2$-basis of K.

Step 1 Select an arbitrary $\boldsymbol{a} \in \mathbb{F}_{2^\ell}$ satisfying $\boldsymbol{a}\mathcal{S}_m \not\subseteq K/\boldsymbol{\beta}$.
Step 2 Set $A = \{\boldsymbol{a}, \boldsymbol{a}, \ldots, \boldsymbol{a}\}$.
Step 3 Select an arbitrary $\boldsymbol{b}_1 = \boldsymbol{b} \in \mathbb{F}_{2^\ell}$ satisfying $\boldsymbol{b} \notin K/\boldsymbol{\beta}$.
Step 4 Set $B = \{\boldsymbol{b}_i\}_{i=1}^{\ell}$, where $\boldsymbol{b}_i = \boldsymbol{b} + \boldsymbol{h}_{i-1}$, for $2 \leq i \leq \ell$.

The output of this construction is the repair scheme for $\boldsymbol{c}_1 = f(0)$ of the code $\mathrm{RS}_{\mathbb{F}_{2^\ell}}(\mathcal{S}_m, 2^m - 2)$ defined by A and B. Repair schemes with the same I/O cost for other $\boldsymbol{c}_j$ can be obtained by modifying this repair scheme (by generalizing [18, Lem. 8] in a straightforward way to the case of subspace).

Theorem 2. *For every $1 \leq m \leq \ell$, the repair scheme in Construction I, which repairs $\boldsymbol{c}_1 = f(0)$ of the code $RS_{\mathbb{F}_{2^\ell}}(\mathcal{S}_m, 2^m - 2)$ where $\mathcal{S}_m$ is an m-dimensional subspace of $\mathbb{F}_{2^\ell}$, incurs an I/O cost of $(n-1)\ell - 2^{m-1}$ bits.*

Proof. Suppose $\boldsymbol{\beta} = \boldsymbol{\beta}_{i^*}$, for some $i^* \in [\ell]$. We aim to show that $|G^m_{A,B,\boldsymbol{\beta}}| = 2^{m-1}$, which implies that there are precisely 2^{m-1} elements $\boldsymbol{\gamma} \in \mathcal{S}_m$ satisfying $\mathsf{Tr}((\boldsymbol{a}_i\boldsymbol{\gamma} + \boldsymbol{b}_i)\boldsymbol{\beta}_{i^*}) = 0$, for all $i \in [\ell]$. Following the proof of Theorem 1, this means there are precisely 2^{m-1} elements $\boldsymbol{\gamma} \in \mathcal{S}_m$ where the i^*-th column in the corresponding matrix $\boldsymbol{W}_{\boldsymbol{\gamma}}$ is zero. This implies a saving of 2^{m-1} bits in the I/O cost. Therefore, the I/O cost of the repair scheme in Construction I is $(n-1)\ell - 2^{m-1}$.

It remains to prove that $|G^m_{A,B,\boldsymbol{\beta}}| = 2^{m-1}$. Note that if $\boldsymbol{a}\boldsymbol{\gamma} + \boldsymbol{b} \in K/\boldsymbol{\beta}$ then $\boldsymbol{a}_i\boldsymbol{\gamma} + \boldsymbol{b}_i = \boldsymbol{a}\boldsymbol{\gamma} + \boldsymbol{b} + \boldsymbol{h}_{i-1} \in K/\boldsymbol{\beta}$, $2 \leq i \leq \ell$. Therefore, we can write

$$G_{A,B,\boldsymbol{\beta}} = \{\boldsymbol{\gamma} \in \mathbb{F}_{2^\ell} \colon \boldsymbol{a}\boldsymbol{\gamma} + \boldsymbol{b} \in K/\boldsymbol{\beta}\} = \frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}.$$

Therefore, $G^m_{A,B,\boldsymbol{\beta}} = \mathcal{S}_m \cap \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)$. By Step 3, $\frac{\boldsymbol{b}}{\boldsymbol{a}} \notin \frac{K}{\boldsymbol{a}\boldsymbol{\beta}}$. Hence, $\mathbb{F}_{2^\ell} = \frac{K}{\boldsymbol{a}\boldsymbol{\beta}} \cup \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)$. According to Step 1, $\mathcal{S}_m \not\subseteq \frac{K}{\boldsymbol{a}\boldsymbol{\beta}}$, which then implies that there exists $\boldsymbol{s} \in \mathcal{S}_m \cap \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)$. We then have $\boldsymbol{s} + (\mathcal{S}_m \cap \frac{K}{\boldsymbol{a}\boldsymbol{\beta}}) \subseteq G^m_{A,B,\boldsymbol{\beta}} = \mathcal{S}_m \cap \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)$. This implies that $|G^m_{A,B,\boldsymbol{\beta}}| \geq |\mathcal{S}_m \cap \frac{K}{\boldsymbol{a}\boldsymbol{\beta}}|$. Since

$$2^m = |\mathcal{S}_m| = \left|\mathcal{S}_m \cap \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)\right| + \left|\mathcal{S}_m \cap \frac{K}{\boldsymbol{a}\boldsymbol{\beta}}\right|,$$

we deduce that $|G^m_{A,B,\boldsymbol{\beta}}| = |\mathcal{S}_m \cap \left(\frac{K}{\boldsymbol{a}\boldsymbol{\beta}} + \frac{\boldsymbol{b}}{\boldsymbol{a}}\right)| \geq 2^{m-1}$. By Lemma 4, we conclude that $|G^m_{A,B,\boldsymbol{\beta}}| = 2^{m-1}$. ∎

Next, we propose another construction based on Proposition 1, which transforms a repair scheme for a full-length code over $\mathbb{F}_{q^m}$ into a repair scheme for the code of the same length $n = q^m$ but over a larger field $\mathbb{F}_{q^\ell} \supset \mathbb{F}_{q^m}$. The proposition applies to arbitrary number of parities and arbitrary q.

Proposition 1 (Lifting transformation)**.** *Let $\mathcal{R}_m$ be a repair scheme of $\mathscr{C}_m = RS_{\mathbb{F}_{q^m}}(\mathcal{A}, k)$ where $\mathcal{A} = \mathbb{F}_{q^m}$. If $m \mid \ell$ then there exists a repair scheme $\mathcal{R}_\ell$ of $\mathscr{C}_\ell = RS_{\mathbb{F}_{q^\ell}}(\mathcal{A}, k)$ so that $\mathsf{ic}(\mathcal{R}_\ell) = \frac{\ell}{m}\mathsf{ic}(\mathcal{R}_m)$, where $\mathsf{ic}(\cdot)$ denotes the I/O cost of a repair scheme w.r.t. an appropriately selected set of bases.*

Proof. Suppose $\mathcal{R}_m$ repairs $\boldsymbol{c}_{j^*}$ of the codeword $\boldsymbol{c} \in \mathscr{C}_m$ and is based on the set of dual codewords $\{\boldsymbol{g}^{(i)}\}_{i=1}^m$ obtained by evaluating the polynomials $\{g^{(i)}(x)\}_{i=1}^m$ of degree less than k at the points of $\mathcal{A}$. Moreover, suppose $\mathcal{B}^m = \{\boldsymbol{\beta}_i\}_{i=1}^m$ is the *common* $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$ used by every node in $\mathcal{R}_m$.

Let $\{\boldsymbol{\tau}_t\}_{t=1}^{\ell/m}$ be an $\mathbb{F}_{q^m}$-basis of $\mathbb{F}_{q^\ell}$ and $\{\boldsymbol{\eta}_t\}_{t=1}^{\ell/m}$ its dual (trace-orthogonal) basis. Consider the scheme $\mathcal{R}_\ell$ that repairs $\overline{\boldsymbol{c}}_{j^*}$ of the codeword $\overline{\boldsymbol{c}} \in \mathscr{C}_\ell$ based on the set of dual codewords $\{\boldsymbol{\tau}_t \boldsymbol{g}^{(i)}\}_{t\in[\ell/m], i\in[m]}$. Note that $\boldsymbol{\tau}_t \boldsymbol{g}^{(i)}$ is obtained by evaluating the polynomial $\boldsymbol{\tau}_t g^{(i)}(x)$ of degree less than k on $\mathcal{A}$ and is, therefore, a dual codeword of $\mathrm{RS}_{\mathbb{F}_{q^\ell}}(\mathcal{A}, k)$. Let $\mathcal{B}^\ell = \{\boldsymbol{\eta}_t \boldsymbol{\beta}_i\}_{t\in[\ell/m], i\in[m]}$ be the *common* basis used by every node in $\mathcal{R}_\ell$. Note that since $\mathsf{rank}_q(\{\boldsymbol{g}_{j^*}^{(i)}\}_{i=1}^m) = m$, we have $\mathsf{rank}_q(\{\boldsymbol{\tau}_t \boldsymbol{g}_{j^*}^{(i)}\}_{t\in[\ell/m], i\in[m]}) = \ell$, and hence, $\mathcal{R}_\ell$ is indeed a repair scheme for $\overline{\boldsymbol{c}}_{j^*}$.

We subsequently show that to repair $\overline{\boldsymbol{c}}_{j^*}$, the I/O cost incurred at the node storing $\overline{\boldsymbol{c}}_j$, $j \neq j^*$, w.r.t. $\mathcal{B}^\ell$, is ℓ/m times more than the I/O cost incurred at the node storing $\boldsymbol{c}_j$ when repairing $\boldsymbol{c}_{j^*}$ w.r.t. $\mathcal{B}^m$. To this end, let $\boldsymbol{W}_j^m$ and $\boldsymbol{W}_j^\ell$ denote the I/O matrices corresponding to $\boldsymbol{c}_j$ and $\overline{\boldsymbol{c}}_j$, respectively. The rows and columns of $\boldsymbol{W}_j^m$ correspond to the sets $\{\boldsymbol{g}^{(i)}\}_{i=1}^m$ and $\{\boldsymbol{\beta}_i\}_{i=1}^m$, respectively, and its typical (i, i')-entry is $\mathsf{Tr}_{\mathbb{F}_{q^m}/\mathbb{F}_q}(\boldsymbol{g}_j^{(i)}\boldsymbol{\beta}_{i'})$. The rows and columns of $\boldsymbol{W}_j^\ell$ correspond to the sets $\{\boldsymbol{\tau}_t \boldsymbol{g}^{(i)}\}_{t\in[\ell/m], i\in[m]}$ and $\{\boldsymbol{\eta}_t \boldsymbol{\beta}_i\}_{t\in[\ell/m], i\in[m]}$, respectively, and its typical $\big((t,i),(t',i')\big)$-entry is

$$\begin{aligned}
\mathsf{Tr}_{\mathbb{F}_{q^\ell}/\mathbb{F}_q}\Big(\boldsymbol{\tau}_t \boldsymbol{g}_j^{(i)} \boldsymbol{\eta}_{t'} \boldsymbol{\beta}_{i'}\Big) &= \mathsf{Tr}_{\mathbb{F}_{q^m}/\mathbb{F}_q}\Big(\mathsf{Tr}_{\mathbb{F}_{q^\ell}/\mathbb{F}_{q^m}}\big(\boldsymbol{\tau}_t \boldsymbol{g}_j^{(i)} \boldsymbol{\eta}_{t'} \boldsymbol{\beta}_{i'}\big)\Big) \\
&= \mathsf{Tr}_{\mathbb{F}_{q^m}/\mathbb{F}_q}\Big(\boldsymbol{g}_j^{(i)} \boldsymbol{\beta}_{i'} \mathsf{Tr}_{\mathbb{F}_{q^\ell}/\mathbb{F}_{q^m}}\big(\boldsymbol{\tau}_t \boldsymbol{\eta}_{t'}\big)\Big) \\
&= \begin{cases} \mathsf{Tr}_{\mathbb{F}_{q^m}/\mathbb{F}_q}\big(\boldsymbol{g}_j^{(i)} \boldsymbol{\beta}_{i'}\big), & \text{if } t = t', \\ 0, & \text{if } t \neq t', \end{cases}
\end{aligned}$$

because $\boldsymbol{g}_j^{(i)}, \boldsymbol{\beta}_{i'} \in \mathbb{F}_{q^m}$ and $\{\boldsymbol{\tau}_t\}_{t=1}^{\ell/m}$ and $\{\boldsymbol{\eta}_t\}_{t=1}^{\ell/m}$ are dual bases of $\mathbb{F}_{q^\ell}$ over $\mathbb{F}_{q^m}$. Therefore, the $\ell \times \frac{\ell}{m}$ submatrix of $\boldsymbol{W}_j^\ell$ consisting of columns indexed by the set $\{(t', i')\}_{t'=1}^{\ell/m}$ can be obtained by taking the Kronecker product of the identity matrix of order ℓ/m and the i'-th column of $\boldsymbol{W}_j^m$. Therefore, in Lemma 3's notation, $\mathsf{nz}(\boldsymbol{W}_j^\ell) = \frac{\ell}{m}\mathsf{nz}(\boldsymbol{W}_j^m)$, which implies that $\mathsf{ic}(\mathcal{R}_\ell) = \frac{\ell}{m}\mathsf{ic}(\mathcal{R}_m)$. ∎

When $m \mid \ell$, the lifting transformation provided in Proposition 1 yields a lower I/O cost than the one in Construction I. However, a specific basis must be used.

Corollary 1. *If $m \mid \ell$ then there exists a repair scheme of the code $RS_{\mathbb{F}_{2^\ell}}(\mathcal{A}, 2^m - 2)$ where $\mathcal{A} = \mathbb{F}_{2^m} \subset \mathbb{F}_{2^\ell}$ with an I/O cost of $(n-1)\ell - \frac{\ell}{m}2^{m-1}$ bits.*

Proof. As shown in [19], there exists a repair scheme $\mathcal{R}_m$ of $\mathscr{C}_m = \mathrm{RS}_{\mathbb{F}_{2^m}}(\mathcal{A}, 2^m - 2)$ where $\mathcal{A} = \mathbb{F}_{2^m}$ with an I/O cost of $(2^m - 1)m - 2^{m-1}$. By Proposition 1, there exists a repair scheme $\mathcal{R}_\ell$ of $\mathscr{C}_\ell = \mathrm{RS}_{\mathbb{F}_{2^\ell}}(\mathcal{A}, 2^m - 2)$ with the I/O cost

$$\mathsf{ic}(\mathcal{R}_\ell) = \frac{\ell}{m}\big((2^m - 1)m - 2^{m-1}\big) = (n-1)\ell - \frac{\ell}{m}2^{m-1}. \quad ∎$$

IV. On the I/O Costs of Other Families of RS Codes

In this section we consider arbitrary number of parities. We show that with a careful selection of the basis at each node, known repair schemes for two families of RS codes can achieve low I/O costs. For the first family, the I/O cost equals the repair bandwidth. The second family achieves the optimal repair bandwidth, and we show that the I/O cost is at most twice the bandwidth. We henceforth assume that different nodes may use different bases but known to all other nodes. We use a different but equivalent way to compute the I/O cost.

A. Another Formulation of the I/O Cost

The dual of $\mathrm{RS}_E(\mathcal{A}, k)$ where $\mathcal{A}$ is not necessarily a subspace is a Generalized Reed-Solomon (GRS) code, given by $\Big\{\big(\mu_{\boldsymbol{\alpha}} g(\boldsymbol{\alpha})\big)_{\boldsymbol{\alpha}\in\mathcal{A}} \colon g \in E[x], \deg(g) < n - k\Big\}$, where $\mu_{\boldsymbol{\alpha}}$ is the (nonzero) *column multiplier* corresponding to the evaluation point $\boldsymbol{\alpha} \in \mathcal{A}$. Moreover, $\mu_{\boldsymbol{\alpha}} = 1/\prod_{\boldsymbol{\alpha}'\in\mathcal{A}\setminus\{\boldsymbol{\alpha}\}}(\boldsymbol{\alpha} - \boldsymbol{\alpha}')$ [21, p. 167], which can be different in general. When $\mathcal{A}$ is a subspace of E, $\mu_{\boldsymbol{\alpha}} = 1/\prod_{\boldsymbol{\alpha}'\in\mathcal{A}^*}\boldsymbol{\alpha}'$ and are the same for every $\alpha \in \mathcal{A}$, which is the case in the previous sections.

If $E = \mathbb{F}_{q^\ell}$ and $F = \mathbb{F}_q$ then each repair scheme for $\boldsymbol{c}_{\boldsymbol{\alpha}^*} = f(\boldsymbol{\alpha}^*)$ is based on a set of ℓ polynomials $g_i(x) \in E[x], i \in [\ell]$, $\deg(g_i) < r = n - k$, and the corresponding repair equations are

$$\mathsf{Tr}_{E/F}\big(\mu_{\boldsymbol{\alpha}^*} g_i(\boldsymbol{\alpha}^*) f(\boldsymbol{\alpha}^*)\big) = -\sum_{\boldsymbol{\alpha}\in\mathcal{A}\setminus\{\boldsymbol{\alpha}^*\}} \mathsf{Tr}_{E/F}\big(\mu_{\boldsymbol{\alpha}} g_i(\boldsymbol{\alpha}) f(\boldsymbol{\alpha})\big).$$

Let $\mathcal{B}_{\boldsymbol{\alpha}} = \{\boldsymbol{\beta}_{\boldsymbol{\alpha},i}\}_{i=1}^\ell$ be the F-basis of E used by the node storing $f(\boldsymbol{\alpha})$. Note that this basis should be *fixed* for α and *independent* of the failed node. By Lemma 2 (c), the computation of $\mathsf{Tr}_{E/F}\big(\mu_{\boldsymbol{\alpha}} g_i(\boldsymbol{\alpha}) f(\boldsymbol{\alpha})\big)$ incurs an I/O cost equal to the Hamming weight of the vector $\boldsymbol{w}^{\mu_{\boldsymbol{\alpha}} g_i(\boldsymbol{\alpha}), \mathcal{B}_{\boldsymbol{\alpha}}}$, which is

$$(\mathsf{Tr}_{E/F}(g_i(\boldsymbol{\alpha})\mu_{\boldsymbol{\alpha}}\boldsymbol{\beta}_{\boldsymbol{\alpha},1}), \ldots, \mathsf{Tr}_{E/F}(g_i(\boldsymbol{\alpha})\mu_{\boldsymbol{\alpha}}\boldsymbol{\beta}_{\boldsymbol{\alpha},\ell})).$$

Let $\mathcal{B}'_{\boldsymbol{\alpha}} = \{\boldsymbol{\beta}'_{\boldsymbol{\alpha},i}\}_{i=1}^\ell$ be the dual (trace-orthogonal) basis of $\mu_{\boldsymbol{\alpha}}\mathcal{B}_{\boldsymbol{\alpha}} = \{\mu_{\boldsymbol{\alpha}}\boldsymbol{\beta}_{\boldsymbol{\alpha},i}\}_{i=1}^\ell$. Then, the I/O cost of each of the trace functionals on the RHS of the repair equations above is the Hamming weight of $\phi_{\mathcal{B}'_{\boldsymbol{\alpha}}}(g_i(\boldsymbol{\alpha}))$, the vector representation of $g_i(\boldsymbol{\alpha})$ w.r.t. the basis $\mathcal{B}'_{\boldsymbol{\alpha}}$. Therefore, the I/O cost incurred at the node storing $f(\boldsymbol{\alpha})$ is equal to the number of elements in $\mathcal{B}'_{\boldsymbol{\alpha}}$ required to linearly generate $g_1(\boldsymbol{\alpha}), \ldots, g_\ell(\boldsymbol{\alpha})$ over $\mathbb{F}_q$. We henceforth use this formulation of the I/O cost.

B. Two-Coset Codes

In this subsection, we slightly modify the construction in [17, Thm. 2] to obtain repair schemes with the I/O costs as low as the repair bandwidths established in their work.

Suppose $m \mid \ell$ and $\boldsymbol{\beta}$ and $\boldsymbol{\gamma}$ are the primitive elements of $\mathbb{F}_{q^\ell}$ and $\mathbb{F}_{q^m}$, respectively. Consider an even $n \leq 2(q^m - 1)$ and $k > 0$ so that $\ell/m \leq r = n - k$. For $0 < m_1 < m_2$ satisfying $m_2 - m_1 = q^{sm}$ (the reason to have this will be clear later), we consider the so-called two-coset code $\mathrm{RS}_{\mathbb{F}_{q^\ell}}(\mathcal{A}, k)$ where $\mathcal{A}$ consists of $n/2$ points from the coset $\boldsymbol{\beta}^{m_1}\mathbb{F}_{q^m}^*$ and $n/2$ points from the coset $\boldsymbol{\beta}^{m_2}\mathbb{F}_{q^m}^*$ of $\mathbb{F}_{q^m}^*$ in $\mathbb{F}_{q^\ell}^*$. As $\boldsymbol{\beta}$ is a primitive element of $\mathbb{F}_{q^\ell}$, these two cosets are disjoint.

Suppose that the node storing $\boldsymbol{c}_{\boldsymbol{\alpha}} = f(\boldsymbol{\alpha})$ uses an $\mathbb{F}_q$-basis $\mathcal{B}_{\boldsymbol{\alpha}}$ so that $\mathcal{B}'_{\boldsymbol{\alpha}} = \{\boldsymbol{\gamma}^{t_1}\boldsymbol{\beta}^{t_2} : 0 \leq t_1 \leq m-1,\ 0 \leq t_2 \leq \ell/m - 1\}$, where $\mathcal{B}_{\boldsymbol{\alpha}}$ and $\mathcal{B}'_{\boldsymbol{\alpha}}$ are defined in Section IV-A. Note that even if $\mathcal{B}'_{\boldsymbol{\alpha}}$ does not depend on α, the column multiplier $\mu_{\boldsymbol{\alpha}}$ does, and so does the basis $\mathcal{B}_{\boldsymbol{\alpha}}$ used by the node. We consider the repair scheme $\mathcal{R}_{\boldsymbol{\alpha}^*}$ for $\boldsymbol{c}_{\boldsymbol{\alpha}^*}$ based on the set of polynomials

$$g_{\boldsymbol{\alpha}^*,(t_1,t_2)}(x) = \begin{cases} \boldsymbol{\gamma}^{t_1}\big(\frac{x}{\boldsymbol{\beta}^{m_2}}\big)^{t_2}, & \text{if } \boldsymbol{\alpha}^* \in \boldsymbol{\beta}^{m_1}\mathbb{F}_{q^m}^*, \\ \boldsymbol{\gamma}^{t_1}\big(\frac{x}{\boldsymbol{\beta}^{m_1}}\big)^{t_2}, & \text{if } \boldsymbol{\alpha}^* \in \boldsymbol{\beta}^{m_2}\mathbb{F}_{q^m}^*, \end{cases}$$

for $0 \leq t_1 \leq m - 1$ and $0 \leq t_2 \leq \ell/m - 1$.

We now analyze the I/O cost of $\mathcal{R}_{\boldsymbol{\alpha}^*}$. Without loss of generality, we assume that $\boldsymbol{\alpha}^* \in \boldsymbol{\beta}^{m_2}\mathbb{F}^*_{q^m}$. First, for $\boldsymbol{\alpha} \in \boldsymbol{\beta}^{m_1}\mathbb{F}^*_{q^m}$, i.e., $\boldsymbol{\alpha} = \boldsymbol{\beta}^{m_1}\boldsymbol{\gamma}^{t_0}$ for some $0 \le t_0 \le q^m - 2$, we have $g_{\boldsymbol{\alpha}^*,(t_1,t_2)}(\boldsymbol{\alpha}) = \boldsymbol{\gamma}^{t_1+t_0t_2} \in \mathbb{F}^*_{q^m}$, for every t_1, t_2. Therefore, $g_{\boldsymbol{\alpha}^*,(t_1,t_2)}(\boldsymbol{\alpha})$ can always be generated as linear combinations of the m elements $\{\boldsymbol{\gamma}^0, \boldsymbol{\gamma}^1, \ldots, \boldsymbol{\gamma}^{m-1}\} \subset \mathcal{B}'_{\boldsymbol{\alpha}}$. Thus, the I/O cost incurred at the node storing $\boldsymbol{c}_{\boldsymbol{\alpha}}$ if $\boldsymbol{\alpha} \in \boldsymbol{\beta}^{m_1}\mathbb{F}^*_{q^m}$ is at most m. Next, suppose that $\boldsymbol{\alpha} = \boldsymbol{\beta}^{m_2}\boldsymbol{\gamma}^{t_0} \in \boldsymbol{\beta}^{m_2}\mathbb{F}^*_{q^m}$. Then, $g_{\boldsymbol{\alpha}^*,(t_1,t_2)}(\boldsymbol{\alpha}) = \boldsymbol{\beta}^{(m_2-m_1)t_2}\boldsymbol{\gamma}^{t_1+t_0t_2}$. As $\{\boldsymbol{\beta}^{t_2}\}_{t_2=0}^{\ell/m-1}$ is an $\mathbb{F}_{q^m}$-basis of $\mathbb{F}_{q^\ell}$, $\{\boldsymbol{\beta}^{(m_2-m_1)t_2}\}_{t_2=0}^{\ell/m-1} = \{\boldsymbol{\beta}^{q^{sm}t_2}\}_{t_2=0}^{\ell/m-1}$ is also an $\mathbb{F}_{q^m}$-basis of $\mathbb{F}_{q^\ell}$ (cf. [17, Lem. 2]). As $\boldsymbol{\gamma}^{t_0t_2} \in \mathbb{F}_{q^m}$, this implies that $\{\boldsymbol{\beta}^{(m_2-m_1)t_2}\boldsymbol{\gamma}^{t_0t_2}\}_{t_2=0}^{\ell/m-1}$ is also an $\mathbb{F}_{q^m}$-basis of $\mathbb{F}_{q^\ell}$. Moreover, since $\{\boldsymbol{\gamma}^{t_1}\}_{t_1=0}^{m-1}$ is an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^m}$, the set $\{g_{\boldsymbol{\alpha}^*,(t_1,t_2)}(\boldsymbol{\alpha})\colon 0 \le t_1 \le m-1, 0 \le t_2 \le \ell/m-1\} = \{\boldsymbol{\beta}^{(m_2-m_1)t_2}\boldsymbol{\gamma}^{t_0t_2}\boldsymbol{\gamma}^{t_1} : 0 \le t_1 \le m-1, 0 \le t_2 \le \ell/m-1\}$ forms an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^\ell}$. Hence, the I/O cost incurred at the node storing $\boldsymbol{c}_{\boldsymbol{\alpha}}$ if $\boldsymbol{\alpha} \in \boldsymbol{\beta}^{m_2}\mathbb{F}^*_{q^m}$ is ℓ. Summing up the two cases,

$$\mathsf{ic}(\mathcal{R}_{\boldsymbol{\alpha}^*}) \le (n/2)m + (n/2-1)\ell,$$

which is the same as the repair bandwidth in [16, Thm. 2].

C. Tamo-Ye-Barg Codes

We analyze the I/O cost of a slightly modified construction in [10], whose repair bandwidth matches the cut-set bound of regenerating codes. Let $p_1 < p_2 < \cdots < p_n$ be n prime numbers satisfying $r \mid (p_j - 1)$, $j \in [n]$, where $r \triangleq n-k$ (note that here we only consider the repair schemes where $d = n-1$ helpers participate). Let $\boldsymbol{\alpha}_j$ be an element of order p_j over $\mathbb{F}_q$, so that adjoining $\boldsymbol{\alpha}_j$ to $\mathbb{F}_q$ results in $\mathbb{F}_{q^{p_j}} = \mathbb{F}_q(\boldsymbol{\alpha}_j)$, for every $j \in [n]$. Set $\mathcal{A} = \{\boldsymbol{\alpha}_j\}_{j=1}^n$ to be the set of n (distinct) evaluation points of the code. Let $\boldsymbol{\beta}$ be an element of order r over $\mathbb{F}_{q^{p_1p_2\cdots p_n}}$. The code $\mathrm{RS}_{\mathbb{F}_{q^\ell}}(\mathcal{A}, k)$ is defined over $\mathbb{F}_{q^\ell} = \mathbb{F}_q(\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_n, \boldsymbol{\beta})$ where $\ell = r\prod_{j=1}^n p_j$.

Assume that the failed node stores $\boldsymbol{c}_{j^*} = f(\boldsymbol{\alpha}_{j^*})$ for $f \in \mathbb{F}_{q^\ell}[x]$, $\deg(f) < k$. We consider the repair scheme $\mathcal{R}_{j^*}$ based on a set of ℓ polynomials $\{\boldsymbol{\eta}_{j^*,t}x^{w-1}\colon t \in [\ell/r], w \in [r]\}$, which are of degrees at most $r-1$. The coefficients $\boldsymbol{\eta}_{j^*,t} \in \mathbb{F}_{q^\ell}$ are constructed as in [10] so that the set $\{\boldsymbol{\eta}_{j^*,t}\boldsymbol{\alpha}_{j^*}^{w-1} : t \in [\ell/r], w \in [r]\}$ forms an $\mathbb{F}_q$-basis of $\mathbb{F}_{q^\ell}$, which then qualifies $\mathcal{R}_{j^*}$ as a repair scheme for $\boldsymbol{c}_{j^*}$. Note that while in [10], the repair scheme for $\boldsymbol{c}_{j^*}$ is considered over the base field $\mathbb{F}_{q^{\prod_{j\neq j^*} p_j}}$, we go further down to $\mathbb{F}_q$ instead and hence new factors will be introduced to reflect that change.

We define $\{\boldsymbol{\eta}_{j^*,t}\}_{t=1}^{\ell/r}$ as the set

$$\Big\{\boldsymbol{\beta}^u\boldsymbol{\alpha}_{j^*}^{u+vr}\prod_{j\neq j^*}\boldsymbol{\alpha}_j^{m_j} : 0 \le u \le r-1, 0 \le v \le \frac{p_{j^*}-1}{r} - 1, 0 \le m_j \le p_j - 1\Big\} \bigcup \Big\{\Big(\sum_{s=0}^{r-1}\boldsymbol{\beta}^s\Big)\boldsymbol{\alpha}_{j^*}^{p_{j^*}-1}\prod_{j\neq j^*}\boldsymbol{\alpha}_j^{m_j} : 0 \le m_j \le p_j - 1\Big\}.$$

As the first and the second sets in the union have sizes $\frac{\ell}{rp_{j^*}}(p_{j^*} - 1)$ and $\frac{\ell}{rp_{j^*}}$, respectively, the union has size $\frac{\ell}{r}$ for a fixed j^*, as claimed. Compared to [10], the set $\{\prod_{j\neq j^*}\boldsymbol{\alpha}_j^{m_j} : 0 \le m_j \le p_j - 1\}$ includes our newly introduced factors and serves as an $\mathbb{F}_q$-basis for $\mathbb{F}_{q^{\prod_{j\neq j^*} p_j}}$. The helper node storing $\boldsymbol{c}_{j'} = f(\boldsymbol{\alpha}_{j'})$, $j' \neq j^*$, chooses a basis $\mathcal{B}_{\boldsymbol{\alpha}_{j'}}$ so that

$$\mathcal{B}'_{\boldsymbol{\alpha}_{j'}} = \Big\{\boldsymbol{\beta}^u\prod_{j=1}^n\boldsymbol{\alpha}_j^{m_j} : 0 \le u \le r-1, 0 \le m_j \le p_j - 1\Big\}.$$

Similar to the two-coset construction, even if $\mathcal{B}'_{\boldsymbol{\alpha}_{j'}}$ does not depend on $\boldsymbol{\alpha}_{j'}$, the bases of the nodes can be different due to the different column multipliers. As discussed in Section IV-A, the I/O cost incurred at this node is equal to the number of elements in $\mathcal{B}'_{\boldsymbol{\alpha}_{j'}}$ required to generate $\{\boldsymbol{\eta}_{j^*,t}\boldsymbol{\alpha}_{j'}^{w-1}\colon t \in [\ell/r], w \in [r]\}$. Note that as $\boldsymbol{\alpha}_{j'}$ is of order $p_{j'}$ over $\mathbb{F}_q$, all of its powers can be represented as an $\mathbb{F}_q$-linear combination of $\{\boldsymbol{\alpha}_{j'}^{m_{j'}} : 0 \le m_{j'} \le p_{j'} - 1\}$. Therefore, every element in the set $\{\boldsymbol{\eta}_{j^*,t}\boldsymbol{\alpha}_{j'}^{w-1}\colon t \in [\ell/r], w \in [r]\}$ can be linearly generated over $\mathbb{F}_q$ by the set

$$\Big\{\boldsymbol{\beta}^u\boldsymbol{\alpha}_{j^*}^{u+vr}\prod_{j\neq j^*}\boldsymbol{\alpha}_j^{m_j} : 0 \le u \le r-1, 0 \le v \le \frac{p_{j^*}-1}{r} - 1, 0 \le m_j \le p_j - 1\Big\} \bigcup \Big\{\boldsymbol{\beta}^u\boldsymbol{\alpha}_{j^*}^{p_{j^*}-1}\prod_{j\neq j^*}\boldsymbol{\alpha}_j^{m_j} : 0 \le u \le r-1, 0 \le m_j \le p_j - 1\Big\}.$$

Clearly, this is a subset of $\mathcal{B}'_{\boldsymbol{\alpha}_{j'}}$ of cardinality $\frac{p_{j^*}+r-1}{p_{j^*}}\frac{\ell}{r}$, which is the I/O cost incurred at the node storing $\boldsymbol{c}_{j'} = f(\boldsymbol{\alpha}_{j'})$. Therefore, the I/O cost in repairing $\boldsymbol{c}_{j^*} = f(\boldsymbol{\alpha}_{j^*})$ is $\big(1 + \frac{r-1}{p_{j^*}}\big)\frac{(n-1)\ell}{r}$ sub-symbols over $\mathbb{F}_q$, where $\frac{(n-1)\ell}{r}$ is the repair bandwidth of this repair scheme. Since $r \le p_{j^*} - 1$ due to the assumption that $r \mid (p_{j^*} - 1)$, the I/O cost is at most twice the (optimal) repair bandwidth for any $\boldsymbol{c}_{j^*}$, and the average I/O cost over the failed nodes is fairly close to the bandwidth.

REFERENCES

[1] I. S. Reed and G. Solomon, "Polynomial codes over certain finite fields," *J. Soc. Ind. Appl. Math.*, vol. 8, no. 2, pp. 300–304, 1960.

[2] H. Dau, I. Duursma, H. M. Kiah, and O. Milenkovic, "Repairing Reed-Solomon codes with multiple erasures," *IEEE Trans. Inform. Theory*, vol. 54, no. 10, pp. 6567–6582, 2018.

[3] "HDFS Erasure Coding," https://hadoop.apache.org/docs/r3.0.0/hadoop-project-dist/hadoop-hdfs/HDFSErasureCoding.html/.

[4] K. Shanmugam, D. S. Papailiopoulos, A. G. Dimakis, and G. Caire, "A repair framework for scalar MDS codes," *IEEE J. Selected Areas Comm. (JSAC)*, vol. 32, no. 5, pp. 998–1007, 2014.

[5] V. Guruswami and M. Wootters, "Repairing Reed-Solomon codes," in *Proc. Annu. Symp. Theory Comput. (STOC)*, 2016.

[6] ——, "Repairing Reed-Solomon codes," *IEEE Trans. Inform. Theory*, vol. 63, no. 9, pp. 5684–5698, 2017.

[7] M. Ye and A. Barg, "Explicit constructions of MDS array codes and RS codes with optimal repair bandwidth," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, 2016, pp. 1202–1206.

[8] H. Dau and O. Milenkovic, "Optimal repair schemes for some families of Reed-Solomon codes," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, 2017, pp. 346–350.

[9] I. Duursma and H. Dau, "Low bandwidth repair of the RS(10,4) Reed-Solomon code," in *Proc. Inform. Theory Applicat. Workshop (ITA)*, 2017.

[10] I. Tamo, M. Ye, and A. Barg, "Optimal repair of Reed-Solomon codes: Achieving the cut-set bound," in *Proc. 58th Annual IEEE Symp. Foundations Computer Sci. (FOCS)*, 2017.

[11] ——, "The repair problem for Reed-Solomon codes: Optimal repair of single and multiple erasures with almost optimal node size," in *IEEE Trans. Inform. Theory*, 2018, to appear.

[12] H. Dau, I. Duursma, H. M. Kiah, and O. Milenkovic, "Repairing Reed-Solomon codes with two erasures," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, 2017, pp. 351–355.

[13] B. Bartan and M. Wootters, "Repairing multiple failures for scalar MDS codes," in *Proc. 55th Annual Allerton Conf. Comm Control Comput. (Allerton)*, 2017.

[14] J. Mardia, B. Bartan, and M. Wootters, "Repairing multiple failures for scalar MDS codes," *IEEE Trans. Inform. Theory*, vol. 65, no. 5, pp. 2661–2672, 2018.

[15] A. Chowdhury and A. Vardy, "Improved schemes for asymptotically optimal repair of MDS codes," in *Proc. 55th Annual Allerton Conf. Comm Control Comput. (Allerton)*, 2017.

[16] W. Li, Z. Wang, and H. Jafarkhani, "A tradeoff between the sub-packetization size and the repair bandwidth for Reed-Solomon code," in *Proc. 55th Annual Allerton Conf. Comm. Control Comput. (Allerton)*, 2017, pp. 942–949.

[17] ——, "A tradeoff between the sub-packetization size and the repair bandwidth for Reed-Solomon codes," *arXiv:1806.00496*, 2018.

[18] H. Dau, I. Duursma, and H. Chu, "On the I/O costs of some repair schemes for full-length Reed-Solomon codes," in *Proc. IEEE Int. Symp. Inform. Theory (ISIT)*, 2018, pp. 1700–1704.

[19] H. Dau and E. Viterbo, "Repair schemes with optimal I/O costs for full-length Reed-Solomon codes with two parities," in *Proc. IEEE Inform. Theory Workshop (ITW)*, 2018, pp. 590–594.

[20] R. Lidl and H. Niederreiter, *Introduction to Finite Fields and Their Applications*. Cambridge University Press, 1986.

[21] R. Roth, *Introduction to Coding Theory*. New York, NY, USA: Cambridge University Press, 2006.