Achieving Secrecy Capacity of Minimum Storage Regenerating Codes for all Feasible (n, k, d) Parameter Values

V. Arvind Rameshwar
Department of Electrical Communication Engineering
Indian Institute of Science, Bengaluru
Email: vrameshwar@iisc.ac.in

Navin Kashyap
Department of Electrical Communication Engineering
Indian Institute of Science, Bengaluru
Email: nkashyap@iisc.ac.in

***Abstract*—This paper addresses the problem of constructing secure exact-repair regenerating codes at the MSR point for all feasible values of the parameters. The setting involves a passive eavesdropper who is allowed to observe the stored contents of, and the downloads into, an l-subset of the n nodes of a distributed storage system (DSS). The objective is to achieve perfect secrecy between the eavesdropped symbols and the file stored on the DSS. Previous secure code constructions (most notably that by Rawat et al.) tackle the problem only for the restricted case wherein the number, d, of helper nodes aiding in the recovery of a failed node is equal to $n-1$. This paper builds on Rawat's work, by combining Gabidulin pre-coding and an MSR construction by Ye and Barg to prove the achievability of secrecy capacity at the MSR point for all allowed values of d.**

I. Introduction

A distributed storage system (DSS) stores a file $\mathbf{f}$ of size M (symbols over a finite field $\mathbb{F}$) on n storage nodes. The system possesses the "k-out-of-n" property, in that a data collector (DC) can recover the file by connecting to any k-subset of the nodes. The nodes, however, are prone to failure and the objective is to design schemes that allow for failed-node repair by contacting any d helper nodes, while preserving the "k-out-of-n" property. The work by Dimakis et al. [1] introduced the concept of *regenerating codes*, which address the problem of simultaneous repair and reconstruction while ensuring that each node stores no more than α independent symbols and each helper node passes on no more than β independent symbols to the failed node. Then from [1],

$$M \leq \sum_{i=1}^{k} \min\{\alpha, (d-i+1)\beta\}. \tag{1}$$

The upper bound describes a tradeoff between the parameters α and β, for a fixed M. Two extremal points of this trade-off curve are the minimum storage regeneration (MSR) and the minimum bandwidth regneration (MBR) points. The MSR point, which is of interest to us, is where α is minimized for a given M. From [1] and the tradeoff curve (1), we have

$$(\alpha_{MSR}, \beta_{MSR}) = \left(\frac{M}{k}, \frac{M}{k(d-k+1)}\right).$$

978-1-5386-9286-8/19/$31.00

Since MSR codes are equivalent to standard MDS array codes, the goal is to suitably augment MDS array code constructions with repair schemes. MSR codes that meet the capacity upper bound of (1) are described in [2], [6], [8] and [11]. In particular, Ye and Barg's constructions in [11] allow for the parameter k to take on all feasible values (from 1 to n), and similarly, d to take any value in its permissible range of $k+1$ to $n-1$.

Now consider the *passive eavesdropper* setting, where an eavesdropper, Eve, is allowed to observe, over a long time, the stored contents of, and the downloads into, an l-subset of the n nodes. We need to ensure that Eve obtains no information about the file stored in the DSS.

Capacity upper bounds for *perfect secrecy* at the MSR and MBR points are provided in [5]. For the MSR point, work towards tightening the bound in [5] can be found in [3], [4] and [8]. While secure codes meeting the capacity upper bound at the MBR point in [5] have been constructed for all values of n, k, d [9], the task of constructing secure codes at the MSR point that achieve the improved capacity upper bound in [3] has been tackled only for the restricted case of $d = n-1$.

In this work, we provide a secure code construction at the MSR point that achieves the capacity upper bound in [3] and [8] for all values of n, k, d, effectively closing the open problem of achieving secrecy capacity at the MSR point.

In Section II, we provide a formal description of the system model and discuss related literature. Section III describes the MSR code construction, and provides a proof of secrecy.

II. Background and Related Work

In this section, we formally describe the system model, and provide details of Gabidulin-based pre-coding, and an overview of the MSR construction by Ye and Barg [11]. In what follows, the notation $[a:b]$ denotes the set of integers between a and b, both inclusive, i.e., $[a:b] = \{i \in \mathbb{Z} : a \leq i \leq b\}$. We use $[n]$ as shorthand for $[1:n]$.

A. System Model

An (n, k, d) DSS consists of n storage nodes, indexed from 1 to n, that store in a distributed, coded fashion, the M

symbols (over a field $\mathbb{F}$) of a file $\mathbf{f}$. The symbols are drawn independently and uniformly at random from the field.

The recovery of the stored file follows the k-*out-of-*n property, i.e., it is sufficient to contact any subset of k nodes, to recover $\mathbf{f}$. Let $\mathbf{c}_i$, $i \in [n]$, denote the coded symbols stored in node i. Firstly, we require that each node stores no more than α independent symbols.

We assume that node failures in the system occur in stages, with no more than one failure at any stage. At stage t, we say that a node j is *active* if it does not fail in that stage. We operate in the *exact-repair* setting, wherein the downloads from d active helper nodes ($k+1 \leq d \leq n-1$) can exactly recover the contents of the failed node. In keeping with [1], our second constraint is that the failed node downloads no more than β independent symbols from any one helper node.

Now, suppose that node i has failed. Let $D_{j,i}$ denote the collection of random symbols sent by helper node j to i. If $H(X)$ represents the entropy of a random variable X, then

$$H(\mathbf{c}_i) \leq \alpha, \tag{2}$$

$$H(D_{j,i}) \leq \beta. \tag{3}$$

From [10], we know that exact-repair codes that satisfy (1) with equality must also satisfy (2) and (3) with equality.

Since at the MSR point, α is minimized for a given M, we have from (1) that $\alpha = M/k$. The minimum value of β then is $\beta = \alpha/(d-k+1)$. Hence,

$$(\alpha, \beta) = \left(\frac{M}{k}, \frac{M}{k(d-k+1)} \right).$$

Now consider the case where an eavesdropper, Eve, observes the downloaded symbols into an arbitrary l-subset $\mathcal{E}$ of the nodes. We assume that each node in $\mathcal{E}$ may fail multiple times, and in order to repair the same node over and over again, information is possibly downloaded from different sets of helper nodes. Thus, over time, Eve knows the stored contents of the nodes in $\mathcal{E}$, and has observed repair information for nodes in $\mathcal{E}$ from all nodes not in $\mathcal{E}$. Let the random vector $\mathbf{e}$ denote the symbols observed by Eve. Thus, $\mathbf{e}$ consists of $\mathbf{c}_i$, $i \in \mathcal{E}$, as well as all the $D_{j,i}$, $i \in \mathcal{E}$, $j \notin \mathcal{E}$. If $\mathbf{f}^{(s)}$ ((s) for "secure") is the file that we desire to store on the DSS, and $M^{(s)}$ is its size, the *perfect secrecy* condition then is: $I(\mathbf{f^{(s)}}; \mathbf{e}) = 0$, where $I(\mathbf{x}; \mathbf{y})$ is the mutual information between the random vectors $\mathbf{x}$ and $\mathbf{y}$.

B. Related Work

The setting of the passive adversary was first discussed in [5], and an upper bound on the secrecy capacity for functional repair was derived to be

$$M^{(s)} \leq \sum_{i=l+1}^{k} \min\{\alpha, (d-i+1)\beta\},$$

where $l = |\mathcal{E}|$. Later work by Shah et al. [9] employed the Product-Matrix (PM) code construction to design a secure MSR coding scheme that achieved a maximum file size of $(k-l)(\alpha - l\beta)$. This was improved upon in [8] and [3], wherein the secrecy capacity was shown to be bounded as

$$M^{(s)} \leq (k-l)(1 - 1/(d-k+1))^l \alpha. \tag{4}$$

This upper bound was shown to be achievable in [8], for the case $n = d+1$, using the concept of zigzag codes. Another achievability scheme, due to Rawat [7], uses a construction in Ye and Barg's paper [11] to show the capacity upper bound in (4) being met, again when $n = d+1$. In this paper, we build upon Rawat's work to prove the achievability of the capacity upper bound in (4) for *all* feasible values of d, using an alternative construction from [11].

The tools we need are provided by the work of Huang et al. [4]. Recall that in a DSS, a given helper node j may in general belong to multiple repair groups (sets of helper nodes) for a given failed node i. A distributed storage code operating at the MSR point is said to be *stable* [4, Definition 7] if for each pair of nodes i and j, the information downloaded from node j to repair node i is the same across all repair groups for i containing the node j. In Lemma 7 of [4], it is shown that for DSSs based on a stable MSR code, the secrecy capacity of an $(n = d+1, k, d)$ DSS is the same as that of an $(n > d+1, k, d)$ DSS, when all other parameters are identical. This result, along with the observations of Rawat [7], in fact suffices to establish that the construction we describe in Section III achieves the secrecy capacity upper bound in (4). We, however, give a more direct proof, involving some ideas from hypergraph theory that may be of independent interest.

C. Preliminaries

Given a DSS that can store M symbols when $l = 0$, we augment our file of size $M^{(s)}$ with random symbols $\mathbf{r}$, where $\mathbf{r}$ is a random vector of length $R = M - M^{(s)}$. Each random symbol in $\mathbf{r}$ is drawn i.i.d. and uniformly at random from the field $\mathbb{F}$. We shall now describe the ingredients of our construction, namely the Gabidulin pre-coding procedure and the *d-optimal repair* MSR construction (for all parameters n, k, d), by Ye and Barg [11].

1) Gabidulin Pre-coding: Assume that we have an M-length vector $\mathbf{m} = (m_1, \ldots, m_M)$, where each m_i, $1 \leq i \leq M$, is drawn from a finite field $\mathbb{F}$. Let $\mathbb{B}$ be some sub-field of $\mathbb{F}$. Further, let points $y_1, y_2, \ldots, y_M$, be elements of $\mathbb{F}$ that are linearly independent over $\mathbb{B}$ ($\dim_{\mathbb{B}}(\mathbb{F}) \geq M$).

The procedure for Gabidulin coding is:

- First, a linearized polynomial $p_{\mathbf{m}}(x)$ is constructed:

$$p_{\mathbf{m}}(x) = \sum_{i=0}^{M-1} m_{i+1} x^{|\mathbb{B}|^i}$$

- The polynomial is evaluated at the collection $\mathcal{Y}$ of points $y_1, y_2, \ldots, y_M$, yielding

$$p(\mathbf{m}, \mathcal{Y}) := (p_m(y_1), p_m(y_2), \ldots, p_m(y_M)).$$

2) Ye and Barg construction: Here, we provide a brief description of the *d-optimal repair* construction.

Construction 1: First we shall introduce some notation: let $s = d-k+1$ and let $\alpha = s^n$. Let $\mathbb{F}$ be a finite field of size $|\mathbb{F}| \geq sn$ and let $\{e_a : a \in [0 : \alpha - 1]\}$ be the standard basis of $\mathbb{F}^\alpha$ over $\mathbb{F}$.

For an integer $a \in [0:\alpha-1]$, let $\mathbf{a} = (a_n, a_{n-1}, \ldots, a_1)$ denote its $s-$ary representation so that $\mathbf{a} = \sum_{i=1}^{n} a_i s^{i-1}$. Suppose that $\{\lambda_{i,j}\}$ for $i \in [n]$ and $j \in [0:s-1]$ are sn distinct elements in $\mathbb{F}$. We define the matrices A_i, $i \in [n]$, to be $\alpha \times \alpha$ diagonal matrices with the $(a,a)^{th}$ entry being λ_{i,a_i}. In other words,

$$A_i = \sum_{a=0}^{\alpha-1} \lambda_{i,a_i} e_a e_a^T. \tag{5}$$

We shall construct an $(n-k)\alpha \times n\alpha$ parity check matrix H for the MSR code $\mathcal{C}$ as:

$$H = \begin{pmatrix} I & \ldots & I & I \\ A_1 & \ldots & A_{n-1} & A_n \\ A_1^2 & \ldots & A_{n-1}^2 & A_n^2 \\ \vdots & \ddots & \vdots & \vdots \\ A_1^{n-k-1} & \ldots & A_{n-1}^{n-k-1} & A_n^{n-k-1} \end{pmatrix} \tag{6}$$

where I is the $\alpha \times \alpha$ identity matrix.

In [11], the authors prove that the code $\mathcal{C}$ obeys the "k-out-of-n" property, while storing exactly α independent symbols in each node. In addition, it is proved that the exact-repair requirement of the DSS is also met, ensuring that the contents of any one failed node can be exactly recovered from $\beta = \frac{\alpha}{d-k+1} = \frac{s^n}{s} = s^{n-1}$ symbols from each of d other active nodes.

We shall now describe the repair scheme. Let $\mathbf{c} = \{\mathbf{c}_1, \ldots, \mathbf{c}_n\}$ be a codeword of $\mathcal{C}$. Since $\alpha = s^n$, we shall index the symbols in $\mathbf{c}_i$ by n-tuples from $[0:s-1]^n$. Let $\mathcal{S} = [0:s-1]$ and $\mathcal{S}^n = [0:s-1]^n$. The symbols in $\mathbf{c}_i$ are indexed by the vectors $\mathbf{a} \in \mathcal{S}^n$, in lexicographic order, starting with the vector $\mathbf{0}$ and ending with the vector $\mathbf{z}$ (the s-ary representation of $\alpha - 1$). In vectorized form, $\mathbf{c}_i$ is the $\alpha \times 1$ column vector given as

$$\mathbf{c}_i = (c_{i,\mathbf{0}}, \ldots, c_{i,\mathbf{z}})^T.$$

For a vector $\mathbf{a} \in \mathcal{S}^n$, let $(a_n, a_{n-1}, \ldots, a_1)$ be its s-ary representation. Now, let $\mathbf{a}(i,u) \in \mathcal{S}^n$ be the vector obtained by substituting the symbol a_i in the s-ary representation of $\mathbf{a}$, with u, for $i \in [n]$ and $u \in [0:s-1]$. Thus,

$$\mathbf{a}(i,u) \equiv (a_n, a_{n-1}, \ldots, a_{i+1}, u, a_{i-1}, \ldots, a_1).$$

Assume that node i has failed and hence, $\mathbf{c}_i$ needs to be recovered. Recall that each helper node j, $j \neq i$, sends exactly $\beta = s^{n-1}$ independent symbols to node i. For some $\mathbf{a} \in \mathcal{S}^n$, we define the set $\mathcal{S}^n_{\mathbf{a},i}$ as

$$\mathcal{S}^n_{\mathbf{a},i} := \{\mathbf{a}(i,u) : \ u \in \mathcal{S}\}.$$

Note that $|\mathcal{S}^n_{\mathbf{a},i}| = s$, for any $\mathbf{a} \in \mathcal{S}^n$. Furthermore,

$$\bigcup_{\mathbf{a}} \mathcal{S}^n_{\mathbf{a},i} = \mathcal{S}^n.$$

Thus, there exist $\beta = s^{n-1}$ distinct sets $\mathcal{S}^n_{\mathbf{a},i}$, the union of which is the entire set of s-ary n-tuples. We shall use these β distinct sets to index the symbols sent by node j to failed node i.

Now, let $D_{j,i}$ represent the symbols contributed by helper node j towards the repair of node i ($j \neq i$). Thus, from [11], $D_{j,i}$ is the row vector of the β symbols

$$\mu_{j,i}^{(\mathcal{S}^n_{\mathbf{a},i})} = \sum_{u=0}^{s-1} c_{j,\mathbf{a}(i,u)} \tag{7}$$

for *distinct* sets $\mathcal{S}^n_{\mathbf{a},i}$, $\mathbf{a} \in \mathcal{S}^n$. Observe that $\mathcal{C}$ is a stable MSR code, in the sense of [4, Definition 7].

III. Secure MSR codes for all parameters

In this section, we describe our construction of secure MSR codes for all feasible values of d, using arguments from [7].

Construction 2: Consider a file $\mathbf{f^{(s)}}$, that we intend storing on the DSS, of size

$$M^{(s)} = (k-l)(1 - 1/(d-k+1))^l \alpha \tag{8}$$

symbols, over a field $\mathbb{F}$. The file size in (8) meets the secrecy capacity upper bound at the MSR point, derived in [3]. As in the Ye and Barg construction in Section II-C2, we take $\alpha = s^n$, where $s = d-k+1$, for $k+1 \leq d \leq n-1$. We now describe our coding scheme:

1) **Gabidulin pre-coding:** To the information set of size $M^{(s)}$, we add $R = M - M^{(s)}$ random symbols (denoted by the vector $\mathbf{r}$), drawn i.i.d. and uniformly from the field $\mathbb{F}$, where $M = k\alpha$. Let this overall message $\mathbf{m} = (\mathbf{f^{(s)}}, \mathbf{r})$ be Gabidulin pre-coded by the procedure described in II-C1. Let

$$\mathbf{f} := p(\mathbf{m}, \mathcal{Y}) = (p_m(y_1), p_m(y_2), \ldots, p_m(y_M)).$$

2) **Ye and Barg encoding:** Let H be the parity check matrix of the Ye and Barg code specified in Construction 1 of Section II-C2. The $k\alpha \times n\alpha$ generator matrix of the code, G, satisfies

$$GH^T = \mathbf{0},$$

where $\mathbf{0}$ denotes the $k\alpha \times (n-k)\alpha$ zero matrix. The code vector $\mathbf{c} = (\mathbf{c}_1, \ldots, \mathbf{c}_n)$ to be stored in the nodes of the DSS is obtained as

$$\mathbf{c} = \mathbf{f}G \in \mathcal{C},$$

where $\mathbf{f}$ is the pre-coded vector from Step 1. The i^{th} node of the DSS stores the vector $\mathbf{c}_i$ of α symbols.

From the discussion in Section II-C2 and from [11], we know that the coding scheme described above is MSR and satisfies the exact-repair property for all values of d. The proof of secrecy follows next.

A. Proof of secrecy for $k+1 \leq d \leq n-1$

We shall follow the line of argument presented in [7]. Let the set of nodes that Eve eavesdrops on be $\mathcal{E} = \{i_1, i_2, \ldots, i_l\}$. In the worst case, all nodes in $\mathcal{E}$ have failed at least once. Note that, as before, we require $|\mathcal{E}| = l < k$. Further, let $D_{j,\mathcal{E}}$ represent the symbols sent by the j^{th} active

node ($j \in \mathcal{D} \subset [n] \setminus \mathcal{E}$, such that $|\mathcal{D}| = d$), for the repair of the nodes in $\mathcal{E}$. Hence,

$$D_{j,\mathcal{E}} = [D_{j,i_1} | D_{j,i_2} | \dots | D_{j,i_l}],$$

where the solid vertical lines represent concatenation.

Without loss of generality, we assume that $\mathcal{E} = \{n-l+1, n-l+2, \dots, n\}$, for, if otherwise, we can always reorder the nodes prior to the first node failure. To characterize the symbols downloaded by the nodes in $\mathcal{E}$, we make the following definition.

Definition III.1. (Symbol Matrix): A symbol matrix P corresponding to the repair scheme (7) is a 0-1 matrix of dimension $l\beta \times \alpha$ such that $D_{j,\mathcal{E}}^T = P\mathbf{c}_j$ for all $j \in [n] \setminus \mathcal{E}$.

In order to explicitly describe the entries of P, we require some notation. Recall from Section II that $\mathcal{S}^n$ represents the set of vectors in $[0 : s-1]^n$, and that $\mathcal{S}^n_{\mathbf{a},i} = \{\mathbf{a}(i,u) : u \in \mathcal{S}\}$, for $\mathbf{a} \in \mathcal{S}^n$. Now, define

$$\mathcal{S}^n_{i\leftarrow *} := \{(a_n, \dots, a_1) : a_i = *, \ a_j \in \mathcal{S} \text{ for } j \neq i\}.$$

Note that for any i, $|\mathcal{S}^n_{i\leftarrow *}| = s^{n-1}$.

For a vector $\mathbf{a} \in \mathcal{S}^n$ (or $\mathcal{S}^n_{j\leftarrow *}$ for some j), let $\mathbf{a}_{\setminus i}$ denote the vector obtained by puncturing $\mathbf{a}$ in its i^{th} coordinate:

$$\mathbf{a}_{\setminus i} = (a_n, \dots, a_{i+1}, a_{i-1}, \dots, a_1).$$

Now, from the definition of the symbol matrix P, we have

$$P = \begin{bmatrix} \frac{P_n}{P_{n-1}} \\ \hline \vdots \\ \hline P_{n-l+1} \end{bmatrix} \tag{9}$$

where each P_i, $i \in [n-l+1 : n]$ is a 0-1 matrix of dimensions $\beta \times \alpha$, such that $P_i\mathbf{c}_j = D_{j,i}^T = \left(\mu_{j,i}^{(\mathcal{S}^n_{\mathbf{0},i})}, \dots, \mu_{j,i}^{(\mathcal{S}^n_{\mathbf{z},i})}\right)^T$, with $\mu_{j,i}^{(\mathcal{S}^n_{\mathbf{a},i})}$ as in (7).

We now seek to characterize P_i, $i \in [n-l+1 : n]$, completely. Let the columns of P_i ($i \in [n-l+1 : n]$) be indexed by all the vectors in $\mathcal{S}^n$, listed in lexicographic order and let the rows of P_i be indexed by the vectors $\mathbf{b} \in \mathcal{S}^n_{i\leftarrow *}$, in lexicographic order.

From (7), we see that the row in P_i indexed by some $\mathbf{b} \in \mathcal{S}^n_{i\leftarrow *}$ contains exactly s 1's — these are in the columns indexed by the vectors $\mathbf{b}(i,u) = (b_n, \dots, b_{i+1}, u, b_{i-1}, \dots, b_1)$, for $u \in [0 : s-1]$. All other entries of P_i are 0's. Note that the column indices containing a 1 entry differ in exactly their i^{th} coordinate. Explicitly,

$$[P_i]_{\mathbf{r},\mathbf{t}} = \begin{cases} 1, & \text{if } \mathbf{t}_{\setminus i} = \mathbf{r}_{\setminus i}, \\ 0, & \text{otherwise} \end{cases} \tag{10}$$

for $\mathbf{r} \in \mathcal{S}^n_{i\leftarrow *}$ and $\mathbf{t} \in \mathcal{S}^n$.

Further, equation (9) coupled with equation (10) above, implies that each column of P contains exactly l 1's, one in each P_i.

Equations (9) and (10) completely characterize P. We add that $H(D_{j,\mathcal{E}}) = \text{rank}(P)$, since the symbols in $\mathbf{c}_j$ are independent of one another.

As an example, consider the $(n,k,d,l) = (4,2,3,1)$ DSS wherein Eve eavesdrops on the last (fourth) node. The symbol matrix P in this case is:

$$P = \begin{bmatrix} 1&0&0&0&0&0&0&0&1&0&0&0&0&0&0&0 \\ 0&1&0&0&0&0&0&0&0&1&0&0&0&0&0&0 \\ 0&0&1&0&0&0&0&0&0&0&1&0&0&0&0&0 \\ 0&0&0&1&0&0&0&0&0&0&0&1&0&0&0&0 \\ 0&0&0&0&1&0&0&0&0&0&0&0&1&0&0&0 \\ 0&0&0&0&0&1&0&0&0&0&0&0&0&1&0&0 \\ 0&0&0&0&0&0&1&0&0&0&0&0&0&0&1&0 \\ 0&0&0&0&0&0&0&1&0&0&0&0&0&0&0&1 \end{bmatrix}.$$

It is easy to verify that rank(P) above is 8, which in turn is equal to $H(D_{j,\mathcal{E}})$.

We intend to obtain a handle on the rank of P, in general. To this end, we claim that the following theorem holds true:

Theorem III.1. *For $s = d-k+1 \geq 2$,*

$$H(D_{j,\mathcal{E}}) = s^{n-l}(s^l - (s-1)^l). \tag{11}$$

In other words, the rank of the symbol matrix P is $s^{n-l}(s^l - (s-1)^l)$.

We shall defer the proof of Theorem III.1 until later, and prove the following theorem, based on the validity of Theorem III.1.

Theorem III.2. *The coding scheme of Construction 2 is secure, for $k+1 \leq d \leq n-1$, against a passive eavesdropper that has access to a set $\mathcal{E} \subset [n]$ of nodes, with $|\mathcal{E}| = l$.*

Proof. The proof of the theorem is similar to the proof of Proposition 1 in [7]. We intend showing that $H(\mathbf{e}) \leq H(\mathbf{r})$ and $H(\mathbf{r}|\mathbf{e}, \mathbf{f}^{(\mathbf{s})}) = 0$, thereby implying (from the perfect secrecy lemma of [9]) that $I(\mathbf{f}^{(\mathbf{s})}; \mathbf{e}) = 0$. Let $\mathcal{T}$ represent a group of $k-l$ nodes such that $\mathcal{T} \cap \mathcal{E} = \emptyset$. We know that

$$\mathbf{e} = (\mathbf{c}_i : i \in \mathcal{E}) \cup \left(\bigcup_{i \in \mathcal{E}} \bigcup_{j \in [n]\setminus\mathcal{E}} \{D_{j,i}\} \right).$$

Now, using the notation $\mathbf{c}_{\mathcal{E}} := (\mathbf{c}_i : i \in \mathcal{E})$, we have

$$\begin{aligned} H(\mathbf{e}) &= l\alpha + H\left(\bigcup_{i \in \mathcal{E}} \bigcup_{j \in [n]\setminus\mathcal{E}} \{D_{j,i}\} \ \middle| \ \mathbf{c}_{\mathcal{E}} \right) \\ &= l\alpha + H\left(\bigcup_{i \in \mathcal{E}} \bigcup_{j \in \mathcal{T}} \{D_{j,i}\} \ \middle| \ \mathbf{c}_{\mathcal{E}} \right) \\ &\leq l\alpha + H\left(\bigcup_{i \in \mathcal{E}} \bigcup_{j \in \mathcal{T}} \{D_{j,i}\} \right) \\ &\leq l\alpha + \sum_{j \in \mathcal{T}} H(D_{j,\mathcal{E}}) \\ &= ls^n + (k-l)(1-(1-1/s)^l)s^n \\ &= ks^n - (k-l)(1-1/s)^l s^n \\ &= M - M^{(s)} = H(\mathbf{r}). \end{aligned} \tag{12}$$

The equality in (12) follows from the fact that

$$\begin{aligned} & H\Big(\bigcup_{i\in\mathcal{E}}\bigcup_{j\in[n]\setminus(\mathcal{T}\cup\mathcal{E})}\{D_{j,i}\}\ \Big|\ \mathbf{c}_{\mathcal{E}}, \bigcup_{i\in\mathcal{E}}\bigcup_{j\in\mathcal{T}}\{D_{j,i}\}\Big) \\ & \leq \sum_{i\in\mathcal{E}} H\Big(\bigcup_{j\in[n]\setminus(\mathcal{T}\cup\mathcal{E})}\{D_{j,i}\}\ \Big|\ \mathbf{c}_{\mathcal{E}}, \bigcup_{j\in\mathcal{T}}\{D_{j,i}\}\Big) \qquad (13) \\ & = 0, \end{aligned}$$

the last equality holding since each summand in (13) equals 0 by the arguments used in the proof of Lemma 7 in [4].

Using the MDS array property of the Ye and Barg code and from Remark 8 of [8], it is possible to show that $H(\mathbf{r}|\mathbf{e}, \mathbf{f}^{(\mathbf{s})}) = 0$. We refer the reader to the proof of Proposition 1 in [7], for more details.

Now, from the perfect secrecy lemma in [9], the two conditions above imply that $I(\mathbf{f}^{(\mathbf{s})}; \mathbf{e}) = 0$, thereby proving that perfect secrecy holds. □

We shall now proceed to the proof of Theorem III.1, beginning with the definitions of a few notions related to hypergraphs.

Definition III.2. (Incidence matrix) The incidence matrix (or vertex-edge incidence matrix) V of a hypergraph (X, E) is a 0-1 matrix of dimension $|V|\times|E|$, with the rows representing nodes and columns representing hyperedges, such that $V_{i,j} = 1$ if edge j is incident on vertex i, and 0 otherwise.

For a vector $\mathbf{v}$, we define its support to be the set of coordinates in which $\mathbf{v}$ takes on non-zero values.

Definition III.3. (Connected hypergraph) A hypergraph (X, E) is said to be connected, if for every pair of nodes $(u, w) \in X \times X$, $u \neq w$, there exists an alternating sequence of nodes and hyperedges, $v_0, h_0, v_1, h_1, \ldots, v_{m-1}, h_{m-1}, v_m$, $(m \in \mathbb{Z}_+)$ with $v_0 = u$ and $v_m = w$, such that for $i \in [0 : m-1]$, h_i is incident on both v_i and v_{i+1}. We call the sequence of hyperedges $h_0, h_1, \ldots, h_{m-1}$ as a *path* from u to w.

Now, we denote by $\mathcal{G}_{s,n}$, the n-dimensional regular hypergraph (X, E) with $|X| = s^n$ and $E \subset X^s$, with incidence matrix $V_{\mathcal{G}_{s,n}}$ defined as follows: let the rows of $V_{\mathcal{G}_{s,n}}$ be indexed by all the vectors in $\mathcal{S}^n$, listed in lexicographic order. Further, let the columns of $V_{\mathcal{G}_{s,n}}$ be indexed by the vectors $\mathbf{b} \in \mathcal{S}^n_{i\leftarrow *}$, ($i$ ranging from n down to 1), where for any i, the vectors $\mathbf{b}$ are listed in lexicographic fashion. Hence, the first s^{n-1} columns of $V_{\mathcal{G}_{s,n}}$ are indexed in lexicographic order by vectors in $\mathcal{S}^n_{n\leftarrow *}$, the next s^{n-1} columns are indexed by vectors in $\mathcal{S}^n_{(n-1)\leftarrow *}$, and so on. Thus, there are ns^{n-1} columns overall. The entries of $V_{\mathcal{G}_{s,n}}$ are

$$[V_{\mathcal{G}_{s,n}}]_{\mathbf{r},\mathbf{t}} = \begin{cases} 1, & \text{if } \mathbf{r}_{\setminus i} = \mathbf{t}_{\setminus i} \text{ and } t_i = * \\ 0, & \text{otherwise} \end{cases} \qquad (14)$$

where $\mathbf{r} \in \mathcal{S}^n$ and $\mathbf{t} \in \mathcal{S}^n_{i\leftarrow *}$, $i \in [n]$.

The column in $V_{\mathcal{G}_{s,n}}$ indexed by the vector $\mathbf{b} \in \mathcal{S}^n_{i\leftarrow *}$ for some i, has exactly s 1's in precisely those rows $\mathbf{t}$ for which $t_i = u$, $u \in [0 : s-1]$ and $t_j = b_j$, for $j \neq i$. Moreover, each row of $V_{\mathcal{G}_{s,n}}$ has exactly n 1's.

With the aid of the definitions above, we wish to prove Theorem III.1. Two lemmas follow. The first (Lemma III.3) establishes that the transpose of the symbol matrix P can be thought of as the incidence matrix of a regular hypergraph having exactly s^{n-l} connected components. Our second lemma (Lemma III.4) proves that the rank of the incidence matrix corresponding to each of these connected components is $(s^l - (s-1)^l)$. We conclude by showing that the rank of P is precisely the sum of the ranks of the incidence matrices of these connected components.

Lemma III.3. *P^T is the incidence matrix of a subgraph $\mathcal{H}$ of $\mathcal{G}_{s,n}$. Further, the number of connected components in $\mathcal{H}$ is s^{n-l}.*

Proof. Recall that the symbol matrix P has a block matrix form, given in equation (9). Taking the transpose of each P_i, $i \in [n-l+1 : n]$ thus gives us:

$$[P_i^T]_{\mathbf{r},\mathbf{t}} = \begin{cases} 1, & \text{if } \mathbf{t}_{\setminus i} = \mathbf{r}_{\setminus i}, \\ 0, & \text{otherwise} \end{cases} \qquad (15)$$

where $\mathbf{r} \in \mathcal{S}^n$ and $\mathbf{t} \in \mathcal{S}^n_{i\leftarrow *}$.

Since i ranges from $n-l+1$ to n only, by comparing the equation above with (14), we get that P^T is a submatrix of $V_{\mathcal{G}_{s,n}}$, containing only the first ls^{n-1} columns of $V_{\mathcal{G}_{s,n}}$. We denote by $\mathcal{H}$, the subgraph induced by this submatrix; thus, $\mathcal{H}$ is a sub-hypergraph of $\mathcal{G}_{s,n}$.

Now, let t be the number of connected components in $\mathcal{H}$ and let h_i, $1 \leq i \leq t$, be the number of hyperedges in each component $\mathcal{H}_i$. Then,

$$\sum_{i=1}^{t} h_i = l\beta = ls^{n-1}. \qquad (16)$$

Pick some node indexed by vector $\mathbf{v}$ so that $\mathbf{v}$ belongs to connected component $\mathcal{H}_j$. Consider the collection of nodes $\mathcal{W}$, where $\mathcal{W} = \{\mathbf{w} \in \mathcal{S}^n : \ w_i = v_i, \ i \in [n-l]$ and $w_j \in [0 : s-1], j \in [n-l+1 : n]\}$. Note that the set $\mathcal{W}$ includes the node $\mathbf{v}$. From equation (10), it is easy to verify that the sub-hypergraph of $\mathcal{H}$ induced by the nodes in $\mathcal{W}$ forms the connected component $\mathcal{H}_j$. This can be seen by choosing some node $\mathbf{x} \notin \mathcal{W}$. Node $\mathbf{x}$ differs from any node in $\mathcal{W}$ in at least one position $j \in [n-l]$. Hence, the row corresponding to $\mathbf{x}$ in P^T will have as support, those columns where none of the nodes in $\mathcal{W}$ have a 1 entry, thereby implying that there does not exist a path from $\mathbf{x}$ to any of the nodes in $\mathcal{W}$. Observe that the number of nodes in $\mathcal{H}_j$ is s^l.

Since the coordinates of nodes $\mathbf{w}$ in $\mathcal{H}_j$ in positions $i \in [n-l]$ are fixed, we puncture the vectors $\mathbf{w}$ at these positions, to form the vectors $\mathbf{w}'$. Thus, the sets $\mathcal{S}^n_{i\leftarrow *}$, $i \in [n-l+1 : n]$ can be written as the sets $\mathcal{S}^l_{i\leftarrow *}$, $i \in [l]$, each set now containing vectors $\mathbf{w}'$.

The incidence matrix $V_{\mathcal{H}_j}$ has entries

$$[V_{\mathcal{H}_j}]_{\mathbf{r},\mathbf{t}} = \begin{cases} 1, & \text{if } \mathbf{t}_{\setminus i} = \mathbf{r}_{\setminus i}, \\ 0, & \text{otherwise} \end{cases} \qquad (17)$$

where $\mathbf{r} \in \mathcal{S}^l$ and $\mathbf{t} \in \mathcal{S}^l_{i\leftarrow *}$, $i \in [l]$.

$\mathcal{H}_j$ is precisely the hypergraph $\mathcal{G}_{s,l}$, having $h_j = ls^{l-1}$ edges. Since this is true for any $j \in [t]$, substituting in equation (16), we get that the number of connected components in $\mathcal{H}$ equals s^{n-l}. □

We shall now introduce an $(s-1)^l$-dimensional code $\mathcal{C}^{\otimes l}$, of block length s^l, the parity check matrix of which will aid us in obtaining a handle on the rank of $V_{\mathcal{G}_{s,l}}$.

Consider the single parity check code $\mathcal{C}_s$ of block length s over the field $\mathbb{F}$, having the $(s-1) \times s$ generator matrix G_s given by

$$G_s = \begin{bmatrix} 1 & 0 & \cdots & 0 & -1 \\ 0 & 1 & \cdots & 0 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -1 \end{bmatrix}.$$

The parity check matrix, H_s, of $\mathcal{C}_s$ is simply the $1 \times s$ all-ones matrix. We then define the direct product code $\mathcal{C}_s^{\otimes l}$ (where $\mathcal{C}_s$ is the underlying code), as the code generated by

$$G_s^{\otimes l} = \underbrace{G_s \otimes G_s \otimes \cdots \otimes G_s}_{l \text{ times}}$$

where $\otimes$ denotes the Kronecker product. A codeword in $\mathcal{C}_s^{\otimes l}$ is of length s^l and can be described by an l-dimensional array. Each entry of the array (which is a coordinate of the codeword) is indexed by an s-ary l-tuple $\mathbf{v} = (v_1, v_2, \ldots, v_l) \in \mathcal{S}^l$.

The code $\mathcal{C}_s^{\otimes l}$ has the property that each array element $\mathbf{v}$ is involved in l parity check equations, one along each coordinate $i \in [l]$. In other words, for every symbol $\mathbf{v} \in \mathcal{S}^l$, there exists a parity check equation indexed by a vector $\mathbf{b} \in S^l_{i \leftarrow *}$, $i \in [l]$, such that $\mathbf{b}_{\setminus i} = \mathbf{v}_{\setminus i}$. Moreover, the parity check equation along coordinate j is the sum of those codeword symbols that differ from $\mathbf{v}$ in only their j^{th} coordinate.

Formally, the code $\mathcal{C}_s^{\otimes l}$ can be described by the $ls^{l-1} \times s^l$ parity check matrix H having entries

$$H_{\mathbf{r},\mathbf{t}} = \begin{cases} 1, & \text{if } \mathbf{t}_{\setminus i} = \mathbf{r}_{\setminus i}, \\ 0, & \text{otherwise} \end{cases} \tag{18}$$

where $\mathbf{r} \in S^l_{i \leftarrow *}$, $i \in [l]$ and $\mathbf{t} \in \mathcal{S}^l$. Thus, each row of H corresponds to a parity check equation $\mathbf{b}$, that finds the sum of symbols $\mathbf{v} \in \mathcal{S}^l$, which differ only in that coordinate of the l-tuple, i, in which $b_i = *$.

Lemma III.4. *The parity check matrix H of $\mathcal{C}_s^{\otimes l}$ is equal to $V^T_{\mathcal{G}_{s,l}}$. Further, rank$(V^T_{\mathcal{G}_{s,l}}) = s^l - (s-1)^l$.*

Proof. The first part of the lemma is obvious from equations (14) and (18). We observe that $\text{rank}(G_s^{\otimes l}) = \prod_{i=1}^{l} \text{rank}(G_s) = \text{rank}(G_s)^l = (s-1)^l$. Thus, the rank of the parity check matrix H is equal to $s^l - (s-1)^l$. □

Using Lemmas III.3 and III.4, we shall now prove Theorem III.1.

Proof. Recall from Lemma III.3 that for any node $\mathbf{v} \in \mathcal{S}^n$ of the hypergraph $\mathcal{H}$, the connected component $\mathcal{H}_j$ containing $\mathbf{v}$ consists of nodes in the set $\mathcal{W} = \{\mathbf{w} \in \mathcal{S}^n : \ w_i = v_i, \ i \in [n-l]$ and $w_j \in [0 : s-1], j \in [n-l+1 : n]\}$.

It is now possible to permute the rows $\mathbf{v}$ of P^T in lexicographic order of $\mathbf{v}' = (v_{n-l}, \ldots, v_1)$ so that all the rows $\mathbf{v} \in \mathcal{S}^n$ corresponding to a fixed value of $\mathbf{v}'$ occur together. Thus, the first s^l rows of P^T are indexed by vectors $\mathbf{v}$ such that $\mathbf{v}' = (0, 0, \ldots, 0)$, the next s^l rows are indexed by vectors $\mathbf{v}$ with $\mathbf{v}' = (0, 0, \ldots, 1)$ and so on. Each collection of s^l rows corresponding to a particular value of $\mathbf{v} = (v_{n-l}, \ldots, v_1)$ forms the incidence matrix of a connected component.

From (17), we observe that the supports of the rows corresponding to any two connected components $\mathcal{H}_i$ and $\mathcal{H}_j$, $i \neq j$, are disjoint. Hence, the rank of P^T is equal to the sum of the ranks of the incidence matrices of the connected components of hypergraph $\mathcal{H}$, induced by P^T. Since each connected component is precisely the hypergraph $\mathcal{G}_{s,l}$ (from the proof of Lemma III.3), we get that

$$\text{rank}(P^T) = s^{n-l}(\text{rank}(V_{\mathcal{G}_{s,l}})) = s^{n-l}(s^l - (s-1)^l),$$

where the first equality follows from Lemma III.3 and the second from Lemma III.4. □

ACKNOWLEDGEMENTS

N. Kashyap's work was supported by a Swarnajayanti Fellowship awarded by the Department of Science and Technology, Government of India.

REFERENCES

[1] A. G. Dimakis, P. B. Godfrey, Y. Wu, M. J. Wainwright, and K. Ramchandran. Network coding for distributed storage systems. *IEEE Transactions on Information Theory*, 56(9):4539–4551, Sept 2010.

[2] S. Goparaju, A. Fazeli, and A. Vardy. Minimum storage regenerating codes for all parameters. *IEEE Transactions on Information Theory*, 63(10):6318–6328, Oct 2017.

[3] S. Goparaju, S. El Rouayheb, A. R. Calderbank, and H. V. Poor. Data secrecy in distributed storage systems under exact repair. *CoRR*, abs/1304.3156, 2013.

[4] K. Huang, U. Parampalli, and M. Xian. On secrecy capacity of minimum storage regenerating codes. *IEEE Transactions on Information Theory*, 63(3):1510–1524, March 2017.

[5] S. Pawar, S. El Rouayheb, and K. Ramchandran. On secure distributed data storage under repair dynamics. In *2010 IEEE International Symposium on Information Theory*, pages 2543–2547, June 2010.

[6] 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 Transactions on Information Theory*, 57(8):5227–5239, Aug 2011.

[7] A. S. Rawat. Secrecy capacity of minimum storage regenerating codes. In *2017 IEEE International Symposium on Information Theory (ISIT)*, pages 1406–1410, June 2017.

[8] A. S. Rawat, O. O. Koyluoglu, N. Silberstein, and S. Vishwanath. Optimal locally repairable and secure codes for distributed storage systems. *IEEE Transactions on Information Theory*, 60(1):212–236, Jan 2014.

[9] N. B. Shah, K. V. Rashmi, and P. V. Kumar. Information-theoretically secure regenerating codes for distributed storage. In *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, pages 1–5, Dec 2011.

[10] N. B. Shah, K. V. Rashmi, P. V. Kumar, and K. Ramchandran. Distributed storage codes with repair-by-transfer and nonachievability of interior points on the storage-bandwidth tradeoff. *IEEE Transactions on Information Theory*, 58(3):1837–1852, March 2012.

[11] M. Ye and A. Barg. Explicit constructions of high-rate mds array codes with optimal repair bandwidth. *IEEE Transactions on Information Theory*, 63(4):2001–2014, April 2017.