

Explicit Codes Minimizing Repair Bandwidth for Distributed Storage

Nihar B. Shah[†], K. V. Rashmi[†], P. Vijay Kumar[†], Kannan Ramchandran[#]

[†] Dept. of ECE, Indian Institute Of Science, Bangalore. {rashmikv, nihar, vijay}@ece.iisc.ernet.in

[#] Dept. of EECS, University of California, Berkeley. kannanr@eecs.berkeley.edu

Abstract—We consider the setting of data storage across n nodes in a distributed manner. A data collector (DC) should be able to *reconstruct* the entire data by connecting to any k out of the n nodes and downloading all the data stored in them. When a node fails, it has to be *regenerated* back using the existing nodes. An obvious means of accomplishing this is to use a Reed-Solomon type MDS code where each node stores a single finite field symbol and where one downloads the entire file for regeneration of a failed node. However, storing vectors in place of symbols makes it easy to extract partial information from a node, and helps in reducing the amount of download required for regeneration of a failed node, termed as *repair bandwidth*.

Recently, there has been additional interest in storing data in *systematic* form as no post processing is required when the DC connects to the k systematic nodes. On failure of a systematic node, there is a need to regenerate it back quickly and exactly due to their preferred status. Replacement of a failed node by an exact replica is termed *exact regeneration*.

In this paper, we consider the problem of minimizing the repair bandwidth for exact regeneration of the systematic nodes. The file to be stored is of size B and each node can store $\alpha = B/k$ units of data. A failed systematic node is regenerated by downloading β units of data each from d existing nodes. We give a lower bound for the repair bandwidth for exact regeneration of the systematic nodes which matches with the bound given by Wu et al. For $d \geq 2k - 1$ we give an explicit code construction which achieves the lower bound on repair bandwidth when the existing $k - 1$ systematic nodes participate in the regeneration. We show the existence and construction of codes that achieve the bound for $d \geq 2k - 3$. Here we also establish the necessity of *interference alignment*. We prove that the bound is not achievable for $d \leq 2k - 4$ when $\beta = 1$, except for the case when $\alpha = 1$ for which any $[n, k]$ MDS code will trivially achieve the bound. We also give a coding scheme which can be used for any d and k , which is optimal for $d \geq 2k - 1$.

I. INTRODUCTION

Consider a scenario where a file of size B is to be stored in a distributed manner across n storage nodes. A data collector (DC) should be able to *reconstruct* the entire file by downloading data stored in any k out of n nodes. Each node can store α units of data (symbols) given by

$$\alpha = B/k \quad (1)$$

When a node fails, the failed node has to be *regenerated* back by downloading β symbols each from d existing nodes as shown in Figure 1. We consider the problem of minimizing the repair bandwidth for exact regeneration of the systematic nodes.

Consider the exact regeneration of a systematic node, say node l by connecting to some set of d nodes. Each symbol




Fig. 1. An illustration of exact regeneration: On failure of node 5, data from nodes 1 to 4 is used to regenerate back the same data that node 5 earlier had.

stored is a linear function of the source symbols. By (1) the linear functionals associated with the symbols stored in any $k - 1$ of the d nodes are linearly independent of those associated with the symbols of node l . Hence an additional α linear functionals are necessary to exactly regenerate node l . From this it follows that a lower bound on the repair bandwidth $d\beta$ is given by

$$d\beta \geq \alpha + (k - 1)\beta \quad (2)$$

In particular, for $\beta = 1$ we have

$$d \geq \alpha + k - 1 \quad (3)$$

Our focus in the current paper is on the case $\beta = 1$. Given a construction for $\beta = 1$, constructions for larger β can be obtained by partitioning the data into smaller chunks, and encoding them individually using the construction for $\beta = 1$. As reconstruction and regeneration are performed separately on these smaller chunks, additional processing and storage required is greatly reduced.

In general, it is an open problem whether this lower bound is achievable for the problem of exact regeneration of the systematic nodes, and we address this issue in the present paper. For $\alpha = 1$, we get $B = k$ and the lower bound as $d \geq k$. In this case, any $[n, k]$ -MDS code will achieve the lower bound for exact regeneration. Hence, we will consider $\alpha > 1$ throughout. We have categorized the (k, d) parameter set with respect to the lower bound on repair bandwidth in Figure 2.

In an independent work [2], authors consider the same setting and provide constructions for codes corresponding to a repair bandwidth that is significantly higher than the lower bound on repair bandwidth.

We say a code is *optimal* exact regenerating if it achieves the lower bound on repair bandwidth for the exact regeneration of systematic nodes. The non-systematic nodes are not the




Fig. 2. Categorization of the (k, d) parameter set

focus and hence their regeneration is not considered in detail. We assume the naive strategy of downloading the entire file for the exact regeneration of the non-systematic nodes.

The pioneering paper in this area [1] considers a more general setting in which each node stores slightly more data than the minimum required, namely $\alpha = (B/k)\nu$ for some $\nu \geq 1$, in order to reduce the repair bandwidth. In this scheme, the regenerated node need not be identical to the failed node as long as it maintains all the properties of the system. The authors establish a tradeoff between the amount of storage in each node α and the repair bandwidth $d\beta$. In the present paper, we are interested only in the case $\nu = 1$ which corresponds to the Minimum Storage Regeneration (MSR) point on the tradeoff. The bound given in equation (2) matches with the MSR point on the tradeoff.

In our previous work [3], we also considered the problem of exact regeneration, however for that value of $\nu > 1$ which minimized the repair bandwidth. We gave explicit codes for the other extreme point on the tradeoff, and an approximately exact regenerating code for the MSR point, both of which minimized the repair bandwidth at the respective points.

In the rest of the paper, the results are presented in terms of k and α , as this leads to a more intuitive understanding of the codes. In Section II we give a subspace viewpoint which will be used throughout the paper. Explicit and optimal code constructions for $k \leq \alpha$ are given in Section III. The existence and construction of optimal codes for $k \leq \alpha + 2$ is given in Section IV. In Section V we prove that the lower bound in not achievable for $k \geq \alpha + 3$ with $\beta = 1$. A coding scheme for any (k, α) parameter set is provided in Section VI which is optimal for $k \leq \alpha$.

II. SUBSPACE VIEWPOINT FOR LINEAR CODES

We consider only linear codes in this paper. By a linear code, we mean that any symbol stored is a linear combination of the source symbols, and only linear operations are allowed on them. Define a vector \underline{z} of length B consisting of the source symbols. Let

$$\underline{z} = \begin{bmatrix} z_1 \\ \vdots \\ z_k \end{bmatrix}$$

where z_i is a column vector of length α . Each source symbol can independently take values from \mathbb{F}_q , a finite field of size q . Hence, the B source symbols can be thought of as forming a B -dimensional vector space over \mathbb{F}_q .

Since the code is linear, any stored symbol can be written as $\underline{\ell}^t \underline{z}$ for some column vector $\underline{\ell}$. These vectors which specify the kernels for the stored symbols define the code, and the actual symbols stored depend on the instantiation of \underline{z} . Since a node stores α symbols, it can be considered as storing α




Fig. 3. Exact regeneration of systematic node l

vectors of the code, and hence can be represented by a $\alpha \times B$ matrix. We will say that the node *stores* this matrix.

Linear operations performed on the stored symbols are equivalent to the same operations performed on the corresponding vectors. Hence storing an $\alpha \times B$ matrix is equivalent to *storing a subspace* of dimension at most α . However, from (1) it is clear that each node must store a subspace of dimension at least α .

For $m = 1, \dots, n$ denote the matrix stored by node m as $\mathbf{G}^{(m)} = [G_1^{(m)} \ G_2^{(m)} \ \dots \ G_k^{(m)}]$, where $G_l^{(m)}, l = 1, \dots, k$ are $\alpha \times \alpha$ matrices. The α symbols stored by node m are $\mathbf{G}^{(m)} \underline{z} = \sum_{l=1}^k G_l^{(m)} z_l$. We will denote the j^{th} row of $G_l^{(m)}$ as $\underline{g}_{jl}^{(m)}$.

There are n storage nodes, out of which k are systematic and store α data symbols each in uncoded form. For $m = 1, \dots, k$, systematic node m stores the symbol set \underline{z}_m . Thus for $l = 1, \dots, k$,

$$G_l^{(m)} = \begin{cases} I_\alpha & \text{if } l = m \\ 0_\alpha & \text{if } l \neq m \end{cases} \quad (4)$$

where 0_α is $\alpha \times \alpha$ zero matrix, and I_α is $\alpha \times \alpha$ identity matrix. Hence, for any non-systematic node m , $G_l^{(m)}$ denotes the *components* along systematic node l that are stored in node m .

For regeneration of a failed systematic node, d other nodes provide one symbol each. We say that each node *passes a vector* for the regeneration of the failed node. In the vectors passed by the non-systematic nodes, the components along the existing systematic nodes constitute *interference*.

Let \mathbb{D} denote the set of d existing nodes used for regeneration of systematic node l . Let $\underline{v}_{\mathbb{D}}^{(m,l)} = [v_{\mathbb{D},1}^{(m,l)} \ \dots \ v_{\mathbb{D},k}^{(m,l)}]$ represent the vector passed by node $m \in \mathbb{D}$ for the regeneration of node $l \notin \mathbb{D}$ (as shown in Figure 3) where $v_{\mathbb{D},i}^{(m,l)}, (i = 1, \dots, k)$ is an α -length row vector representing the component along the symbols of the systematic node i . Thus, $v_{\mathbb{D},i}^{(m,l)}, (i = 1, \dots, k, i \neq l)$ constitute interference. Let $\underline{x}_{\mathbb{D}}^{(m,l)} = [x_{\mathbb{D},1}^{(m,l)} \ \dots \ x_{\mathbb{D},\alpha}^{(m,l)}]$ be the coefficients of the linear combination of the rows of $\mathbf{G}^{(m)}$ to obtain the vector that node m passes for regeneration of systematic node l . For brevity, we will discard the subscript \mathbb{D} from the notation and the set of d nodes being used for regeneration will be clear from the context. Thus,

$$\underline{v}^{(m,l)} = \underline{x}^{(m,l)} \mathbf{G}^{(m)} \quad (5)$$

Throughout this paper, we use superscripts to refer to the node numbers, and subscripts to index the elements of any matrix. No distinction is made between row and column vectors and the orientation of the vector under consideration is clear from the context. \underline{e}_i represents an α -length unit vector with 1 in i^{th} position and 0 elsewhere. We say two vectors are *aligned* if they are linearly dependent.

III. OPTIMAL EXPLICIT CODE FOR $k \leq \alpha$

In this section an explicit linear construction is given, which achieves optimal exact regeneration of systematic nodes for $k \leq \alpha$. The construction assumes that when a systematic node fails, the existing $k - 1$ systematic nodes along with any α non-systematic nodes participate in the regeneration. First, we provide a code construction for $k = \alpha$. Codes for any $k < \alpha$ can be obtained by modifying the code for $k = \alpha$. Initially, a simple example is given to illustrate the code.

A. Example

Take $k = \alpha = 3$. This gives $d = 5$ and $B = 9$. Thus each node stores a 3×9 matrix. Let $n = 6$ and $q = 7$.

Let the first three nodes be systematic. Hence,

$$\mathbf{G}^{(1)} = [I_3 \ 0_3 \ 0_3] \quad (6)$$

$$\mathbf{G}^{(2)} = [0_3 \ I_3 \ 0_3] \quad (7)$$

$$\mathbf{G}^{(3)} = [0_3 \ 0_3 \ I_3] \quad (8)$$

Let $\Psi_3 = \begin{bmatrix} \psi_1^{(4)} & \psi_2^{(4)} & \psi_3^{(4)} \\ \psi_1^{(5)} & \psi_2^{(5)} & \psi_3^{(5)} \\ \psi_1^{(6)} & \psi_2^{(6)} & \psi_3^{(6)} \end{bmatrix}$ be a 3×3 Cauchy matrix [5]. Any submatrix of a Cauchy matrix is full rank.

The three non-systematic nodes store the matrices $\mathbf{G}^{(m)}$, $m = 4, 5, 6$, given by

$$\begin{bmatrix} 2\psi_1^{(m)} & 2\psi_2^{(m)} & 2\psi_3^{(m)} & \psi_2^{(m)} & 0 & 0 & \psi_3^{(m)} & 0 & 0 \\ 0 & \psi_1^{(m)} & 0 & 2\psi_1^{(m)} & 2\psi_2^{(m)} & 2\psi_3^{(m)} & 0 & \psi_3^{(m)} & 0 \\ 0 & 0 & \psi_1^{(m)} & 0 & 0 & \psi_2^{(m)} & 2\psi_1^{(m)} & 2\psi_2^{(m)} & 2\psi_3^{(m)} \end{bmatrix}$$

1) *Regeneration*: For the regeneration of systematic node l ($\in \{1, 2, 3\}$), each non-systematic node passes its l^{th} row. The choice of the non-systematic node matrices is such that in the vectors passed for regeneration of a systematic node, components along the existing systematic nodes (which constitute interference) are aligned. For example, consider regeneration of systematic node 1. Each non-systematic node passes its first row. First rows of $\mathbf{G}^{(m)}$, $m = 4, 5, 6$, have components along the systematic nodes (nodes 2 and 3) aligned in the direction $[1 \ 0 \ 0]$. Now, the second and third systematic nodes pass $[0 \ 0 \ 0 \ 1 \ 0 \ 0 \ 0 \ 0 \ 0]$, and $[0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 1 \ 0 \ 0]$ respectively and cancel out the interference leaving behind the matrix $[\Psi_3 \ 0_3 \ 0_3]$. Since Ψ_3 is invertible, systematic node 1 can be exactly regenerated.

2) *Reconstruction*: To reconstruct the entire data, DC can connect to any three nodes. For reconstruction to be possible, the 9×9 matrix formed by juxtaposing the node matrices of these three nodes one below the other should be non-singular.

Reconstruction is trivially satisfied when the data collector connects to all the three systematic nodes. Suppose, the data

collector connects to two systematic nodes and one non-systematic node. For example, suppose DC connects to nodes 2, 3 and 4. For reconstruction, we need the following matrix to be non-singular:

$$A_1 = G_1^{(4)} = \begin{bmatrix} 2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} \\ 0 & \psi_2^{(4)} & 0 \\ 0 & 0 & \psi_3^{(4)} \end{bmatrix} \quad (9)$$

which is full rank since the elements of a Cauchy matrix are non-zero.

Consider the data collector connecting to one systematic node and two non systematic nodes. For example, suppose it connects to nodes 1, 4, and 5. Since all symbols of node 1 are available, $G_1^{(4)}$ and $G_1^{(5)}$ can be cancelled out. Hence for reconstruction to be possible the matrix B_1 given below must be full rank.

$$B_1 = \begin{bmatrix} G_2^{(4)} & G_3^{(4)} \\ G_2^{(5)} & G_3^{(5)} \end{bmatrix}$$

Claim: The matrix B_1 is full rank.

Proof: For $i = 2, 3, 1$ (in this order), group the i^{th} rows of the two non-systematic nodes together to give matrix

$$B_2 = \begin{bmatrix} 2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} & 0 & \psi_3^{(4)} & 0 \\ 2\psi_1^{(5)} & 2\psi_2^{(5)} & 2\psi_3^{(5)} & 0 & \psi_3^{(5)} & 0 \\ \hline 0 & 0 & \psi_2^{(4)} & 2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} \\ 0 & 0 & \psi_2^{(5)} & 2\psi_1^{(5)} & 2\psi_2^{(5)} & 2\psi_3^{(5)} \\ \hline \psi_2^{(4)} & 0 & 0 & \psi_3^{(4)} & 0 & 0 \\ \psi_2^{(5)} & 0 & 0 & \psi_3^{(5)} & 0 & 0 \end{bmatrix}$$

$$\text{Let } \Psi_2 = \begin{bmatrix} \psi_2^{(4)} & \psi_3^{(4)} \\ \psi_2^{(5)} & \psi_3^{(5)} \end{bmatrix}$$

Ψ_2 is a submatrix of the Cauchy matrix Ψ_3 and hence is invertible. Multiply the three groups of two rows each by Ψ_2^{-1} to obtain

$$B_3 = \begin{bmatrix} \Psi_2^{-1} & 0_3 & 0_3 \\ 0_3 & \Psi_2^{-1} & 0_3 \\ 0_3 & 0_3 & \Psi_2^{-1} \end{bmatrix} B_2 \quad (10)$$

$$= \begin{bmatrix} \phi & 2 & 0 & 0 & 0 & 0 \\ \phi & 0 & 2 & 0 & 1 & 0 \\ \hline 0 & 0 & 1 & \phi & 2 & 0 \\ 0 & 0 & 0 & \phi & 0 & 2 \\ \hline 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \quad (11)$$

where ϕ is some arbitrary value. Rows 1, 4, 5, 6 (and columns 1, 2, 4, 6) are linearly independent of all others (this includes all the columns containing ϕ) and hence be eliminated to obtain,

$$B_4 = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}$$

which is full rank. \blacksquare

Now consider the case of DC connecting to three non-systematic nodes. Let C_1 be the matrix formed by juxtaposing the matrices stored in these three nodes one below the other.

Claim: The matrix C_1 is full rank.

Proof: In C_1 , group the i^{th} ($i = 1, 2, 3$) rows of all the three nodes together to obtain the matrix C_2 . Thus,

$$C_2 = \quad (12)$$

$$\begin{array}{c}
\text{grp1} \\
\text{grp2} \\
\text{grp3}
\end{array}
\left[\begin{array}{ccc|ccc|ccc}
2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} & \psi_2^{(4)} & 0 & 0 & \psi_3^{(4)} & 0 & 0 \\
2\psi_1^{(5)} & 2\psi_2^{(5)} & 2\psi_3^{(5)} & \psi_2^{(5)} & 0 & 0 & \psi_3^{(5)} & 0 & 0 \\
2\psi_1^{(6)} & 2\psi_2^{(6)} & 2\psi_3^{(6)} & \psi_2^{(6)} & 0 & 0 & \psi_3^{(6)} & 0 & 0 \\
\hline
0 & \psi_1^{(4)} & 0 & 2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} & 0 & \psi_3^{(4)} & 0 \\
0 & \psi_1^{(5)} & 0 & 2\psi_1^{(5)} & 2\psi_2^{(5)} & 2\psi_3^{(5)} & 0 & \psi_3^{(5)} & 0 \\
0 & \psi_1^{(6)} & 0 & 2\psi_1^{(6)} & 2\psi_2^{(6)} & 2\psi_3^{(6)} & 0 & \psi_3^{(6)} & 0 \\
\hline
0 & 0 & \psi_1^{(4)} & 0 & 0 & \psi_2^{(4)} & 2\psi_1^{(4)} & 2\psi_2^{(4)} & 2\psi_3^{(4)} \\
0 & 0 & \psi_1^{(5)} & 0 & 0 & \psi_2^{(5)} & 2\psi_1^{(5)} & 2\psi_2^{(5)} & 2\psi_3^{(5)} \\
0 & 0 & \psi_1^{(6)} & 0 & 0 & \psi_2^{(6)} & 2\psi_1^{(6)} & 2\psi_2^{(6)} & 2\psi_3^{(6)}
\end{array} \right]$$

Multiply the 3 groups of 3 rows each by Ψ_3^{-1} to get a matrix C_3 given by

$$C_3 = \begin{bmatrix} \Psi_3^{-1} & 0_3 & 0_3 \\ 0_3 & \Psi_3^{-1} & 0_3 \\ 0_3 & 0_3 & \Psi_3^{-1} \end{bmatrix} C_2$$

$$= \left[\begin{array}{ccc|ccc|ccc}
2 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\
0 & 2 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 2 & 0 & 0 & 0 & 1 & 0 & 0 \\
\hline
0 & 1 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 2 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 2 & 0 & 1 & 0 \\
\hline
0 & 0 & 1 & 0 & 0 & 0 & 2 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & 0 & 2 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 2
\end{array} \right]$$

Rows 1, 2 and 3 in the groups 1, 2 and 3 respectively (i.e. rows 1, 5 and 9) are clearly independent of all others (and so are the corresponding columns). The remaining 6×6 submatrix can be rearranged to get the following form:

$$C_4 = \left[\begin{array}{cc|cc|cc}
2 & 1 & 0 & 0 & 0 & 0 \\
1 & 2 & 0 & 0 & 0 & 0 \\
\hline
0 & 0 & 2 & 1 & 0 & 0 \\
0 & 0 & 1 & 2 & 0 & 0 \\
\hline
0 & 0 & 0 & 0 & 2 & 1 \\
0 & 0 & 0 & 0 & 1 & 2
\end{array} \right]$$

This is a block diagonal matrix, and since $q = 7$, is full rank. Thus the matrix C_1 is full rank. ■

B. Explicit General Code Construction for $k = \alpha$

Let Ψ be an $(n - k) \times \alpha$ Cauchy matrix [5] with elements drawn from \mathbb{F}_q , i.e.,

$$\Psi = \begin{bmatrix} \psi^{(k+1)} \\ \psi^{(k+2)} \\ \vdots \\ \psi^{(n)} \end{bmatrix} \quad (13)$$

where $\psi^{(i)} = [\psi_1^{(i)} \dots \psi_\alpha^{(i)}]$, $i = k + 1, \dots, n$ are α -length row vectors. Any submatrix of a Cauchy matrix is full rank. The minimum field size required for the construction of this Cauchy matrix is:

$$q \geq \alpha + n - k \quad (14)$$

Note that since $n - k \geq \alpha \geq 2$, we will have $q \geq 4$.

For $m = k + 1, \dots, n$, $i, j = 1, \dots, \alpha$, set

$$\underline{g}_{ij}^{(m)} = \begin{cases} \epsilon \psi_j^{(m)} & \text{if } i = j \\ \psi_j^{(m)} \underline{e}_i, & \text{if } i \neq j \end{cases} \quad (15)$$

where ϵ is any arbitrary value such that $\epsilon \neq 0$ and $\epsilon^2 \neq 1$. Note that there always exists such a value if $q \geq 4$.

As illustrated in the example, this choice makes the interference in the vectors passed by non-systematic nodes for the regeneration of a failed systematic node *aligned*. This enables the existing systematic nodes to cancel the interference by passing one symbol each.

1) *Regeneration*: Consider regeneration of systematic node $\hat{l} (\hat{l} \in \{1, \dots, k\})$. All non-systematic nodes who participate in the regeneration pass their \hat{l}^{th} row, i.e. if non-systematic node $m (\in \{k + 1, \dots, n\})$ participates in the regeneration, then it passes

$$\underline{\mathbf{v}}^{(m, \hat{l})} = [\underline{g}_{i,1}^{(m)} \dots \underline{g}_{i,k}^{(m)}] \quad (16)$$

The systematic node l ($l = 1, \dots, k$, $l \neq \hat{l}$) passes

$$\underline{\mathbf{v}}^{(l, \hat{l})} = [\underline{0} \dots \underline{0} \ \underline{e}_{\hat{l}} \ \underline{0} \dots \underline{0}] \quad (17)$$

with $\underline{e}_{\hat{l}}$ in the \hat{l}^{th} position.

From equation (15), $\underline{g}_{i, \hat{l}}^{(m)}$ are all aligned along the direction of $\underline{e}_{\hat{l}}$. Hence $\underline{\mathbf{v}}^{(l, \hat{l})}$ can be used to remove interference along the systematic node l from $\underline{\mathbf{v}}^{(m, \hat{l})}$, $\forall m$ participating in the regeneration.

Also from equation (15), $\underline{g}_{i, \hat{l}}^{(m)}$ are rows of the Cauchy matrix Ψ , and hence are linearly independent. Using these α linearly independent vectors, the systematic node \hat{l} can be regenerated.

2) *Reconstruction*: For reconstruction to be successful, the matrices stored in the k nodes to which the DC connects, when juxtaposed one below the other, should form a $B \times B$ full rank matrix. Call this the *reconstruction matrix* R . If the DC connects to the k systematic nodes, then reconstruction is trivially satisfied. Consider DC connecting to p non-systematic nodes, and $k - p$ systematic nodes, $1 \leq p \leq k$. Let $\delta_1, \dots, \delta_p$ be the p non-systematic nodes to which DC connects and let $\Omega_1, \dots, \Omega_p$ ($\Omega_1 < \dots < \Omega_p$) be the p systematic nodes to which DC does *not* connect.

Reconstruction is successful if and only if the $p\alpha \times p\alpha$ matrix R formed by components along systematic nodes $\Omega_1, \dots, \Omega_p$ in $\mathbf{G}^{(\delta_1)}, \dots, \mathbf{G}^{(\delta_p)}$ is non-singular.

$$R = \begin{bmatrix} \mathbf{G}^{(\delta_1)} \\ \vdots \\ \mathbf{G}^{(\delta_p)} \end{bmatrix} = \begin{bmatrix} G_{\Omega_1}^{(\delta_1)} & G_{\Omega_2}^{(\delta_1)} & \dots & G_{\Omega_p}^{(\delta_1)} \\ \vdots & \vdots & & \vdots \\ G_{\Omega_1}^{(\delta_p)} & G_{\Omega_2}^{(\delta_p)} & \dots & G_{\Omega_p}^{(\delta_p)} \end{bmatrix} \quad (18)$$

Theorem 1: R is full rank.

The proof of Theorem 1 is provided in Appendix. The steps followed in the proof are similar to the ones used in the example.

C. Explicit Code construction for $k < \alpha$

For a given (k, α) first construct the code for $k = \alpha$. The theorem given below shows the existence and construction for any $k < \alpha$.

Theorem 2: If there exists a (k, α) linear code for exact regeneration of the systematic nodes, then there also exists a (\hat{k}, α) linear code for any $\hat{k} \leq k$.

Proof: Suppose there exists an (k, α) code for exact regeneration of the systematic nodes. From each node matrix, remove the last $(k - \hat{k})\alpha$ columns so that now $\mathbf{G}^{(m)}$ is of the size $\alpha \times \hat{k}\alpha$. Thus, we will have $\hat{B} = \hat{k}\alpha$ data symbols. Consider only the set of first \hat{k} systematic nodes and all the non-systematic nodes. This forms a (\hat{k}, α) code.

Reconstruction: Suppose the DC connects to ℓ systematic nodes and $\hat{k} - \ell$ non-systematic nodes. This case is same as the case in original (k, α) code where the DC connects to the removed $k - \hat{k}$ systematic nodes along with the above ℓ systematic nodes and $\hat{k} - \ell$ non-systematic nodes. Hence, the DC can reconstruct \hat{B} data symbols.

Regeneration: During regeneration of a systematic node, we have $\hat{d} = \alpha + \hat{k} - 1 = d - (k - \hat{k})$. All the α non-systematic nodes participating in the regeneration pass exactly the same vector as in the original code. Since the last $(k - \hat{k})\alpha$ column sets have been removed from the non-systematic node matrices, there is no interference from the data symbols corresponding to the removed $k - \hat{k}$ systematic nodes. All the interference which is present are aligned and the components along the failed systematic node span an α -dimensional space as in the original code. Hence $k - \hat{k}$ lesser vectors will be able to regenerate the failed node. ■

Remark: The above construction is optimal for $\beta = 1$. For any higher β , the data to be stored can be split into smaller chunks, which can be encoded individually using this construction for $\beta = 1$. Hence, this construction is optimal for any value of β .

IV. EXISTENCE AND CONSTRUCTION FOR $k \leq \alpha + 2$

The existence and construction of exact regenerating codes which meet the bound given by (3) is shown for the parameter set $k \leq \alpha + 2$. This proof assumes that when a systematic node fails, the existing $k - 1$ systematic nodes participate in regeneration along with any α non-systematic nodes, passing one symbol each. The proof can be extended to the general case as well, where any d existing nodes can participate in the regeneration.

A. Approach

In the sequel the reconstruction and regeneration conditions will be cast as product of rational polynomials. We will need to show that there exists a set of non-zero values such that these polynomials are all well defined and non-zero. In [4] a similar problem is arises in proving the existence of capacity achieving multicast network codes, but with respect to polynomials. But the argument can be easily extended to rational polynomials. If $\frac{f_1(\underline{x})}{g_1(\underline{x})}, \dots, \frac{f_p(\underline{x})}{g_p(\underline{x})}$ are rational polynomials, then define $f_{p+1}(\underline{x}) = \gcd(g_1(\underline{x}), \dots, g_p(\underline{x}))$. There exists a solution to \underline{x} such that the product of the rational polynomials is well defined and non-zero if and only if there exists a solution to \underline{x} such that the product of the polynomials $f_1(\underline{x}), \dots, f_{p+1}(\underline{x})$ is non-zero. Hence, the algorithm given by Koetter and Medard in [4] can be used to find the values of the variables, provided the field size is large enough.

B. Necessary Properties

1) Necessary Properties for Reconstruction:

Lemma 3: For reconstruction property to hold, for any non-systematic node m , $G_l^{(m)}$ must be full rank $\forall l \in \{1, \dots, k\}$.

Proof: Given some m and l , suppose the DC connects to the $k - 1$ systematic nodes other than l , and to the non-systematic node m . From the $k - 1$ systematic nodes, the DC recovers $(k - 1)\alpha$ data symbols. Hence column sets corresponding to these $k - 1$ systematic nodes (i.e $G_{\hat{l}}^{(m)}$, $\hat{l} = 1, \dots, k$, $\hat{l} \neq l$) can be removed from $\mathbf{G}^{(m)}$ leaving behind only $G_l^{(m)}$. Thus, for successful reconstruction, $G_l^{(m)}$ should be full rank. ■

2) Necessary Properties for Exact Regeneration:

Lemma 4: For the regeneration of a failed systematic node l ($l \in \{1, \dots, k\}$), the components along node l in the vectors passed by the α non-systematic nodes participating in the regeneration must be linearly independent.

Proof: Consider the regeneration of a failed systematic node l , by connecting to $k - 1$ existing systematic nodes m_1, \dots, m_{k-1} and α non-systematic nodes $m_k, \dots, m_{k-1+\alpha}$. Let matrix

$$\mathbf{V} = \begin{bmatrix} \underline{\mathbf{v}}^{(m_1, l)} \\ \vdots \\ \underline{\mathbf{v}}^{(m_{k-1+\alpha}, l)} \end{bmatrix} = [V_1 \quad V_2 \quad \cdots \quad V_k] \quad (19)$$

where $V_i = \begin{bmatrix} v_i^{(m_1, l)} \\ \vdots \\ v_i^{(m_{k-1+\alpha}, l)} \end{bmatrix}$, ($i = 1, \dots, k$) is a $d \times \alpha$ matrix

representing the component of \mathbf{V} along the i^{th} systematic node. For successful regeneration of l , we need an $\alpha \times d$ matrix Y such that

$$Y\mathbf{V} = \mathbf{G}^{(l)} \quad (20)$$

Consider the component of $Y\mathbf{V}$ along node l . Since $G_l^{(l)} = I_\alpha$, we need $\text{rank}(YV_l) \geq \alpha$. Since the $k - 1$ other systematic nodes cannot provide any vector in the direction of l , we get $v_i^{(m, l)} = 0$ for $i, m = 1, \dots, k, i \neq m, l \neq m$. Thus the first $k - 1$ rows of V_l are $\underline{0}$. Hence, the remaining α rows of V_l , which are the components along the failed node in the vectors given by the non-systematic nodes, have to be linearly independent. ■

Remark: Since only the last α rows of V_l are non-zero, the last α columns of Y should also be linearly independent.

Lemma 5: (Need for Interference Alignment) For the regeneration of a failed systematic node l , and for any $\hat{l} \in \{1, \dots, k, \hat{l} \neq l\}$, the vectors $v_{\hat{l}}^{(m, l)}$, $\forall m \in \{k + 1, \dots, n\}$ should be aligned.

Proof: Using the same notations as in Lemma 4, consider the components along any other systematic node \hat{l} . Since $G_{\hat{l}}^{(l)} = 0_\alpha$, we need $YV_{\hat{l}} = 0$. Since the other $k - 2$ systematic nodes provide $\underline{0}$ component along node \hat{l} , the corresponding rows of $V_{\hat{l}}$ will be zero. Let $\tilde{V}_{\hat{l}}$ ($\alpha + 1 \times \alpha$) and \tilde{Y} ($\alpha \times \alpha + 1$) be sub-matrices of $V_{\hat{l}}$ and Y with the $k - 2$ zero rows in $V_{\hat{l}}$ and the corresponding columns in Y removed. Thus we need $\tilde{Y}\tilde{V}_{\hat{l}} = 0$. Since $\text{rank}(\tilde{Y}) \geq \alpha$, it forces $\text{rank}(\tilde{V}_{\hat{l}}) \leq 1$. Hence, for the regeneration of a systematic node, in the vectors passed by the

α non systematic nodes, the components along any existing systematic node should be aligned in the same direction. By choosing different sets of α non-systematic nodes, we get that alignment should hold for all the non-systematic nodes. ■

Theorem 6: A necessary and sufficient condition for exact regeneration of a failed systematic node l by connecting to the existing $k - 1$ systematic nodes and α non-systematic nodes is that the set of vectors passed by these non-systematic nodes satisfy Lemmas 4 and 5.

Proof: Necessity: Proved in Lemmas 4 and 5 itself. *Sufficiency:* Suppose Lemma 5 is satisfied. Then, in the vectors passed by the non-systematic nodes, the components along any other systematic node \hat{l} are aligned in the same direction, i.e. $\underline{v}_{\hat{l}}^{(m,l)} = \kappa_{\hat{l}}^{(m,l)} \underline{w}_{\hat{l}}^{(l)}$ where m is any non-systematic node, $\underline{w}_{\hat{l}}^{(l)}$ is a vector independent of m , and κ 's are some constants in \mathbb{F}_q . The systematic node \hat{l} passes $\underline{v}_{\hat{l}}^{(\hat{l},l)} = \underline{w}_{\hat{l}}^{(l)}$ with the components of $\underline{\mathbf{v}}^{(\hat{l},l)}$ along other nodes as $\underline{0}$. Hence, this can be used to subtract the component in $\underline{\mathbf{v}}^{(m,l)}$ along any other systematic node \hat{l} , to give a set of vectors

$$\tilde{\underline{\mathbf{v}}}^{(m,l)} = \underline{\mathbf{v}}^{(m,l)} - \sum_{i=1, i \neq l}^k \kappa_i^{(m,l)} \underline{\mathbf{v}}^{(i,l)} \quad (21)$$

Since $\underline{v}_{\hat{l}}^{(\hat{l},l)} = \underline{0}$ for $\hat{l} = 1, \dots, k$, $\hat{l} \neq l$, we get $\tilde{\underline{v}}_{\hat{l}}^{(m,l)} = \underline{v}_{\hat{l}}^{(m,l)}$. Since Lemma 4 is satisfied, the components of these α vectors along node l are independent, and hence span the α -dimensional subspace stored in node l . ■

C. Structure of the Code

For $m = k + 1, \dots, n$ let,

$$G_i^{(m)} = \Lambda_i^{(m)} H_i^{(m)}, \quad i = 1, \dots, k \quad (22)$$

where $\Lambda_i^{(m)} = \text{diag}\{\lambda_{1,i}^{(m)}, \dots, \lambda_{\alpha,i}^{(m)}\}$ is an $\alpha \times \alpha$ diagonal matrix and

$$H_i^{(m)} = \begin{bmatrix} h_{1,i}^{(m)} \\ h_{2,i}^{(m)} \\ \vdots \\ h_{\alpha,i}^{(m)} \end{bmatrix} \quad (23)$$

where $\underline{h}_{i,j}^{(m)}$ is an α -length row vector. Also set

$$\lambda_{i,i}^{(m)} = 1, \quad i = 1, \dots, k \quad (24)$$

Regeneration: For $m = k + 1, \dots, n$, $l = 1, \dots, k$, let

$$\underline{\mathbf{v}}^{(m,l)} = \underline{\mathbf{x}}^{(m,l)} \mathbf{G}^{(m)} \quad (25)$$

For $l = 1, \dots, \alpha$, set

$$\underline{\mathbf{x}}^{(m,l)} = \underline{e}_l \quad \forall m \in \{k + 1, \dots, n\} \quad (26)$$

i.e. for regeneration of the systematic node $l (\in \{1, \dots, \alpha\})$, each non-systematic node passes the l^{th} row of its node matrix.

Thus to satisfy Lemma 5 we choose,

$$\begin{aligned} \underline{h}_{i,j}^{(m)} &= \underline{h}_{i,j}, & m &= k + 1, \dots, n & (27) \\ & & i &= 1, \dots, \alpha, \\ & & j &= 1, \dots, k, \quad j \neq i \end{aligned}$$

Thus, for regeneration of systematic nodes $1, \dots, \alpha$, the interference is aligned, and hence can be subtracted out.

Reconstruction: The DC connects to any set of k nodes and downloads all the $k\alpha$ data symbols stored in them.

D. Existence and construction for $k = \alpha + 2$

Consider regeneration of the systematic node $\alpha + 1$. By Lemma 5, the component along the systematic node l , $\forall l \in \{1, \dots, \alpha\}$ in the vector passed by non-systematic nodes need to be aligned in one direction. This leads to the following set of $n - k - 1$ equations: for $m = k + 2, \dots, n$,

$$\kappa_l^{(m,\alpha+1)} \underline{x}^{(k+1,\alpha+1)} G_l^{(k+1)} = \underline{x}^{(m,\alpha+1)} G_l^{(m)} \quad (28)$$

Similarly, alignment for the regeneration of the systematic node $\alpha + 2$ leads to another set of $n - k - 1$ equations: $m = k + 2, \dots, n$,

$$\kappa_l^{(m,\alpha+2)} \underline{x}^{(k+1,\alpha+2)} G_l^{(k+1)} = \underline{x}^{(m,\alpha+2)} G_l^{(m)} \quad (29)$$

for some constants κ 's $\in \mathbb{F}_q$.

Set

$$\kappa_l^{(m,\alpha+1)} = \kappa_l^{(m,\alpha+2)} = \kappa_l^{(m)} \text{ (say)} \quad (30)$$

For all $m \in \{k + 2, \dots, n\}$, multiply equation (28) by $(x_i^{(m,\alpha+1)})^{-1}$ and (29) by $(x_i^{(m,\alpha+2)})^{-1}$ and subtract the two. $h_{l,l}^{(m)}$ gets eliminated and a homogeneous equation in terms of $\underline{h}_{1,l}, \dots, \underline{h}_{l-1,l}, \underline{h}_{l+1,l}, \dots, \underline{h}_{\alpha,l}$ remains. One way to satisfy this equation is to equate all the scalar coefficients to zero.

This gives, for $l = 1, \dots, \alpha$, $m = k + 2, \dots, n$ and $i = 1, \dots, \alpha, i \neq l$,

$$\begin{aligned} \lambda_{i,l}^{(m)} &= \kappa_l^{(m)} \lambda_{i,l}^{(k+1)} \left[(x_l^{(m,\alpha+1)})^{-1} x_i^{(k+1,\alpha+1)} - (x_l^{(m,\alpha+2)})^{-1} \right. \\ &\quad \left. x_i^{(k+1,\alpha+2)} \right] \left[(x_l^{(m,\alpha+1)})^{-1} x_i^{(m,\alpha+1)} - (x_l^{(m,\alpha+2)})^{-1} x_i^{(m,\alpha+2)} \right]^{-1} \end{aligned} \quad (31)$$

Equation (31) ensures that second set of equations (i.e. 29) are satisfied whenever the first set (i.e. 28) is satisfied.

Note that any polynomial containing a $\lambda_{i,l}^{(m)}$ ($i \neq l$) term will be a rational polynomial. For such polynomials, we will obtain an assignment which will simultaneously ensure that none of the inverted terms are zero, and the polynomial is also not zero.

Now only the first set of equations have to be satisfied, for which, using equation (28) make the following assignments, for $m = k + 2, \dots, n$

$$\begin{aligned} \underline{h}_{l,l}^{(m)} &= (x_l^{(m,\alpha+1)})^{-1} [\kappa_l^{(m)} \{ \underline{h}_{l,l}^{(k+1)} x_l^{(k+1,\alpha+1)} + \\ &\quad \sum_{i=1, i \neq l}^{\alpha} \lambda_{i,l}^{(k+1)} x_i^{(k+1,\alpha+1)} \underline{h}_{i,l} \} - \sum_{i=1, i \neq l}^{\alpha} \lambda_{i,l}^{(m)} x_i^{(m,\alpha+1)} \underline{h}_{i,l}] \end{aligned} \quad (32)$$

The component along systematic node $\alpha + 1$ needs to be aligned in the vector passed for the regeneration of systematic node $\alpha + 2$ and vice versa. Hence the alignment of systematic nodes $\alpha + 1$ and $\alpha + 2$ result only in one set of $n - k - 1$ equations each. Consider the regeneration of the $(\alpha + 2)^{th}$ systematic node. By Lemma 5, the the component along the $(\alpha + 1)^{th}$ systematic node in the vector passed by non-systematic nodes need to be aligned in one direction. This leads to the following set of $n - k - 1$ equations: For $m = k + 2, \dots, n$,

$$\kappa_{\alpha+1}^{(m)} \underline{x}^{(k+1, \alpha+2)} H_{\alpha+1}^{(k+1)} \Lambda_{\alpha+1}^{(k+1)} = \underline{x}^{(m, \alpha+2)} H_{\alpha+1}^{(m)} \Lambda_{\alpha+1}^{(m)} \quad (33)$$

By equation (27) we have

$$H_{\alpha+1}^{(m)} = H_{\alpha+1}^{(k+1)} = H_{\alpha+1} \text{ (say)} \quad (34)$$

Thus, equating the coefficients to zero, we get for $i = 1, \dots, \alpha$,

$$\lambda_{i, \alpha+1}^{(m)} = \kappa_{\alpha+1}^{(m)} x_i^{(k+1, \alpha+2)} \lambda_{i, \alpha+1}^{(k+1)} (x_i^{(m, \alpha+2)})^{-1} \quad (35)$$

Similarly, for regeneration of node $\alpha + 1$, we need to align components along node $\alpha + 2$ which leads to

$$\lambda_{i, \alpha+2}^{(m)} = \kappa_{\alpha+2}^{(m)} x_i^{(k+1, \alpha+1)} \lambda_{i, \alpha+2}^{(k+1)} (x_i^{(m, \alpha+1)})^{-1} \quad (36)$$

Regeneration: Exact regeneration of each one of the systematic nodes $l \in \{1, \dots, \alpha\}$ results in a condition

$$\det \begin{pmatrix} \underline{h}_{l, l}^{(m_1)} \\ \vdots \\ \underline{h}_{l, l}^{(m_\alpha)} \end{pmatrix} \neq 0 \quad (37)$$

where m_1, \dots, m_α are the α non-systematic nodes used for regeneration After substituting for $\underline{h}_{l, l}^{(m_i)}$, $i = 1, \dots, \alpha$ from equation (32), this condition evaluates to a rational polynomial, which can be shown to be not identically equal to zero by the following assignments:

$$\begin{aligned} \kappa_l^{(m)} &= 1, \quad \lambda_{i, l}^{(k+1)} = 1, \quad \underline{h}_{l, l}^{(k+1)} = \underline{e}_l, \quad \underline{h}_{i, l} = \underline{e}_i, \\ x_i^{(k+1, \alpha+1)} &= 0, \quad x_i^{(k+1, \alpha+2)} = 1, \quad x_l^{(k+1, \alpha+1)} = 1 \\ x_l^{(m, \alpha+1)} &= 1, \quad x_l^{(m, \alpha+2)} = 1, \quad x_i^{(m, \alpha+1)} = 1 \\ x_i^{(m, \alpha+2)} &= (m - k)^{-j} + 1 \end{aligned} \quad (38)$$

for $i = 1, \dots, \alpha$, $i \neq l$, $m \in \{m_1, \dots, m_\alpha\}$, $m \neq k + 1$ and $j = i$ if $i < l$, $j = i - 1$ if $i > l$. This set of assignments makes the matrix under consideration in equation (37) a Vandermonde matrix which is full rank, and ensures that equations (31), (35) and (36) remain valid, provided the field size is large enough.

Exact regeneration of systematic nodes $\alpha + 1$ and $\alpha + 2$ also result in conditions of rational polynomials being not equal to zero. For exact regeneration of $(\alpha + 1)^{th}$ systematic node, Lemma 4 should hold. Choose $H_{\alpha+1}$ to be a full rank matrix. Lemma 4 implies that the coefficients resulting from the linear combinations need to be linearly independent, i.e $\underline{x}^{(m, \alpha+1)} \Lambda_{\alpha+1}^{(m)}$ should be linearly independent for any α out of the $n - k$ non-systematic nodes. Express the determinant of this matrix as a polynomial. To show that this polynomial is not identically zero, we choose $\Lambda_{\alpha+1}^{(k+1)} = I$, $\kappa_{\alpha+1}^{(m)} =$

1 , $x_i^{(m, \alpha+2)} = 1$, $x_i^{(k+1, \alpha+2)} = 1$ for $i = 1, \dots, \alpha$, $m = k + 2, \dots, n$.

From (35), we get $\Lambda_{\alpha+1}^{(m)} = I$. Choose $x_i^{(m, \alpha+1)} = (m - k)^i$ for $m = k + 1, \dots, n$ to make it a Vandermonde matrix (also ensuring that equations (31), (35) and (36) remain valid), provided the field size is large enough. A similar argument can be used to obtain a condition for regeneration of node $\alpha + 2$.

Reconstruction: The condition for reconstruction property to hold can be expressed as a product of polynomials not being equal to zero by viewing each determinant as a polynomial. For reconstruction to be successful, the node matrices corresponding to the k nodes to which the data collector connects, when juxtaposed one below the other, should form a $B \times B$ full rank matrix. If the data collector connects to the k systematic nodes, then reconstruction is trivially satisfied. Consider DC connecting to p non-systematic nodes and $k - p$ systematic nodes, $1 \leq p \leq k$. Let m_1, \dots, m_p , ($m_1 < \dots < m_p$) be the non-systematic nodes to which it connects. Let l_1, \dots, l_p , ($l_1 < \dots < l_p$) be the p systematic nodes to which it does *not* connect. Due to the structure of node matrices of the systematic nodes, we will be left with the condition of the $l\alpha \times l\alpha$ matrix formed by the column sets l_1, \dots, l_p of the node matrices of the p non-systematic nodes being non-singular. Thus the polynomial corresponding to this choice of k nodes is

$$\det \begin{pmatrix} G_{l_1}^{(m_1)} & G_{l_2}^{(m_1)} & \dots & G_{l_p}^{(m_1)} \\ G_{l_1}^{(m_2)} & G_{l_2}^{(m_2)} & \dots & G_{l_p}^{(m_2)} \\ \vdots & \vdots & \dots & \vdots \\ G_{l_1}^{(m_p)} & G_{l_2}^{(m_p)} & \dots & G_{l_p}^{(m_p)} \end{pmatrix} \quad (39)$$

We will now show that there exists an assignment of the variables such that this polynomial is not identically zero. For $i, j = 1, \dots, p$, $m_1 \neq k + 1$, set

$$\kappa_{l_j}^{(m_i)} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \quad (40)$$

For $i = 1, \dots, \alpha$, $m = k + 2, \dots, n$, $l = 1, \dots, k$ set $H_l^{(k+1)} = I$, $\Lambda_l^{(k+1)} = I$, $x_i^{(m, \alpha+1)} = 1$, $x_i^{(k+1, \alpha+1)} = 1$, $x_i^{(k+1, \alpha+2)} = 1$, $x_i^{(m, \alpha+2)} = (m - k)^i$. Thus from (31), (32), (35) and (36), we get that this matrix is full rank for a large enough field size, and also the equations (31), (35) and (36) remain valid.

Hence, provided that the field size is large enough, one can find solutions for these variables such that both reconstruction and exact regeneration properties are satisfied.

E. Existence and construction for $k < \alpha + 2$

For a given α , the code described for the previous subsection can be modified using Theorem 2 to obtain a code for any $k < \alpha + 2$.

Remark: This achievability scheme is optimal for $\beta = 1$. For any higher β , the data to be stored can be split into smaller chunks, which can be encoded individually using this construction for $\beta = 1$. Hence, this for any value of β , optimal regeneration for the parameter set $k \leq \alpha + 2$ is achievable using this scheme.

V. NON-ACHIEVABILITY FOR $k \geq \alpha + 3$

We define two codes to be *equivalent* if the corresponding nodes of both codes store the same subspace, and pass the same vectors for regeneration of any node. The only difference may be in the representation of what is stored in a node, i.e. in the node matrices.

Lemma 7: If there exists an exact regenerating code for $k \geq \alpha + 1$, then there exists an equivalent code with the following property:

$$\underline{h}_{i,j}^{(m)} = \underline{h}_{i,j}, \text{ for } i = 1, \dots, \alpha, j = 1, \dots, k, j \neq i \quad (41)$$

for any non-systematic node m .

Proof: Consider a code which performs exact regeneration of the systematic nodes. For any non-systematic node m in this code, we will obtain linearly independent vectors in the subspace stored in it one by one, and set them as the rows of its node matrix. By induction we will prove that the matrix of any node will take the following form: $\mathbf{G}^{(m)} =$

$$\begin{bmatrix} \lambda_{1,1}^{(m)} \underline{h}_{1,1}^{(m)} & \cdots & \lambda_{1,\alpha}^{(m)} \underline{h}_{1,\alpha} & \lambda_{1,\alpha+1}^{(m)} \underline{h}_{1,\alpha+1} & \cdots & \lambda_{1,k}^{(m)} \underline{h}_{1,k} \\ \lambda_{2,1}^{(m)} \underline{h}_{2,1} & \cdots & \lambda_{2,\alpha}^{(m)} \underline{h}_{2,\alpha} & \lambda_{2,\alpha+1}^{(m)} \underline{h}_{2,\alpha+1} & \cdots & \lambda_{2,k}^{(m)} \underline{h}_{2,k} \\ \vdots & & \vdots & & & \vdots \\ \lambda_{\alpha,1}^{(m)} \underline{h}_{\alpha,1} & \cdots & \lambda_{\alpha,\alpha}^{(m)} \underline{h}_{\alpha,\alpha} & \lambda_{\alpha,\alpha+1}^{(m)} \underline{h}_{\alpha,\alpha+1} & \cdots & \lambda_{\alpha,k}^{(m)} \underline{h}_{\alpha,k} \end{bmatrix} \quad (42)$$

for $m = k+1, \dots, n$. Note that by Lemma 3, $H_l^{(m)}$ and $\Lambda_l^{(m)}$ have to be full rank matrices.

Let the first row of $\mathbf{G}^{(m)}$ represent the vector passed by the non-systematic node m for the regeneration of the first systematic node. From Lemma 5, as the interference from the remaining $k-1$ systematic nodes has to be aligned,

$$\underline{h}_{1,j}^{(m)} = \underline{h}_{1,j}, \quad j = 2, \dots, k \quad (43)$$

Hence, the first row has to be of the given form. Suppose the vectors passed for the regeneration of systematic nodes $1, \dots, p-1$ ($1 < p \leq \alpha$) are linearly independent in all the non-systematic nodes. By a similar argument, the first $p-1$ rows of $\mathbf{G}^{(m)}$ have to be of the given form.

Now consider the regeneration of the p^{th} systematic node. Suppose some of the non-systematic nodes (say type A) pass a vector linearly dependent on the first $p-1$ rows of their node matrix and some (say type B) pass a linearly independent vector. Each type B node will have this vector as a new row in their node matrices. In this set of vectors passed for regeneration of the p^{th} systematic node, consider the component along systematic node $(\alpha+1)$, i.e. $\underline{v}_{\alpha+1}^{(m,p)}$. In vectors passed by type A nodes, this vector is a linear combination of $\underline{h}_{i,\alpha+1}$, $i = 1, \dots, p-1$, whereas in vectors passed by type B nodes, it is linearly independent of $\underline{h}_{i,\alpha+1}$, $i = 1, \dots, p-1$ (by Lemma 4). But by Lemma 5, these vectors have to be aligned. Hence there is a contradiction.

Suppose all the non-systematic nodes are of type A. Then all the vectors passed by non-systematic nodes will be linearly dependent on the first $p-1$ rows of their node matrices. Hence in all the vectors, the component along systematic node p $\underline{v}_p^{(m,p)}$ will be a linear combination of $\underline{h}_{i,p}$, $i = 1, \dots, p-1$. These can span at most $p-1$ dimensions whereas p^{th} systematic node spans α dimensions. Hence the regeneration of p^{th} systematic node is not possible.

Hence all the non-systematic nodes should be of type B. i.e they pass linearly independent vectors for the generation of systematic nodes. Along with Lemma 5, this proves that all α rows of the node matrix have to be of the given form. ■

Henceforth in this section, we will consider all nodes to be of this form.

Remark: In this code, for the regeneration of the p^{th} systematic node (for $1 \leq p \leq \alpha$), each non-systematic node passes the p^{th} row of its node matrix.

Corollary 8: For $l = \alpha+1, \dots, k$, and any non-systematic nodes m and m' ,

$$H_l^{(m)} = H_l^{(m')} \quad (44)$$

Corollary 9: For any non-systematic node m and $k \geq \alpha+1$, any α out of $\underline{\mathbf{v}}^{(m,1)}, \dots, \underline{\mathbf{v}}^{(m,k)}$ are linearly independent.

Proof: In the code given in Lemma 7, the choice of the first α systematic nodes was arbitrary. Hence, for a given set of α systematic nodes, an equivalent code can be constructed considering these as the first α nodes. Thus, by Lemma 7 the vectors passed by any non-systematic node for regeneration of these systematic nodes will be the α rows of its matrix. Hence they are independent. ■

Lemma 10: For $k \geq \alpha+2$, for any non-systematic node m ,

$$x_i^{(m,l)} \neq 0 \quad \text{for } l = \alpha+1, \dots, k, i = 1, \dots, \alpha$$

Proof: Suppose for some non-systematic node m , $l \in \{\alpha+1, \dots, k\}$ and $i \in \{1, \dots, \alpha\}$, $x_i^{(m,l)} = 0$. Since for $j \in \{1, \dots, \alpha\}$, $\underline{\mathbf{v}}^{(m,j)}$ is the j^{th} row of the node matrix of node m , $\underline{\mathbf{v}}^{(m,l)}$ is a linear combination of $\underline{\mathbf{v}}^{(m,j)}$, $j = 1, \dots, \alpha$, $j \neq i$. This is a contradiction to Corollary 9. ■

Theorem 11: For a linear code with $\beta = 1$, exact regeneration of systematic nodes meeting the bound given by (3) is not possible for $k \geq \alpha+3$.

Proof: (By contradiction) Consider any code that achieves exact regeneration of systematic nodes meeting the bound in (3) for $k \geq \alpha+3$. By Corollary 8,

$$H_i^{(m)} = H_i^{(m')} \quad (45)$$

for $m, m' \in k+1, \dots, n$, $i = \alpha+1, \dots, k$. Call these matrices H_i .

Consider regeneration of systematic node $(\alpha+3)$. By Lemma 5, components corresponding to systematic nodes $(\alpha+1)$ and $(\alpha+2)$ are to be aligned. Hence we have

$$\underline{\mathbf{x}}^{(k+1,\alpha+3)} G_{\alpha+1}^{(k+1)} = \kappa_1 \underline{\mathbf{x}}^{(k+2,\alpha+3)} G_{\alpha+1}^{(k+2)} \quad (46)$$

$$\underline{\mathbf{x}}^{(k+1,\alpha+3)} G_{\alpha+2}^{(k+1)} = \kappa_2 \underline{\mathbf{x}}^{(k+2,\alpha+3)} G_{\alpha+2}^{(k+2)} \quad (47)$$

where κ_1 and κ_2 are some constants in \mathbb{F}_q . From equation (45) and since $H_{\alpha+1}$ is full rank (Lemma 3), this simplifies to

$$\underline{\mathbf{x}}^{(k+1,\alpha+3)} \Lambda_{\alpha+1}^{(k+1)} = \kappa_1 \underline{\mathbf{x}}^{(k+2,\alpha+3)} \Lambda_{\alpha+1}^{(k+2)} \quad (48)$$

$$\underline{\mathbf{x}}^{(k+1,\alpha+3)} \Lambda_{\alpha+2}^{(k+1)} = \kappa_2 \underline{\mathbf{x}}^{(k+2,\alpha+3)} \Lambda_{\alpha+2}^{(k+2)} \quad (49)$$

$$\begin{aligned} \implies \kappa_1 \underline{\mathbf{x}}^{(k+2,\alpha+3)} \Lambda_{\alpha+1}^{(k+2)} (\Lambda_{\alpha+1}^{(k+1)})^{-1} = \\ \kappa_2 \underline{\mathbf{x}}^{(k+2,\alpha+3)} \Lambda_{\alpha+2}^{(k+2)} (\Lambda_{\alpha+2}^{(k+1)})^{-1} \end{aligned} \quad (50)$$

Since no element of $\underline{x}^{(k+2,\alpha+3)}$ is zero (Lemma 10), and the Λ matrices are diagonal, we get

$$\kappa_1 \Lambda_{\alpha+1}^{(k+2)} (\Lambda_{\alpha+1}^{(k+1)})^{-1} = \kappa_2 \Lambda_{\alpha+2}^{(k+2)} (\Lambda_{\alpha+2}^{(k+1)})^{-1} \quad (51)$$

Since none of the elements of the diagonal Λ matrices are zero,

$$\kappa_1 \neq 0, \quad \kappa_2 \neq 0 \quad (52)$$

Similarly, on regeneration of systematic node $\alpha + 2$, the components along $\alpha + 1$ and $\alpha + 3$ have to be aligned. Hence

$$\tilde{\kappa}_1 \Lambda_{\alpha+1}^{(k+2)} (\Lambda_{\alpha+1}^{(k+1)})^{-1} = \tilde{\kappa}_2 \Lambda_{\alpha+3}^{(k+2)} (\Lambda_{\alpha+3}^{(k+1)})^{-1} \quad (53)$$

where $\tilde{\kappa}_1$ and $\tilde{\kappa}_2$ are some other non-zero constants in \mathbb{F}_q .

Now, for regeneration of systematic node $(\alpha + 3)$ the component provided along it by the first non-systematic node is,

$$\begin{aligned} \underline{x}^{(k+1,\alpha+3)} \Lambda_{\alpha+3}^{(k+1)} H_{\alpha+3} \\ = \kappa_1 \underline{x}^{(k+2,\alpha+3)} \Lambda_{\alpha+1}^{(k+2)} (\Lambda_{\alpha+1}^{(k+1)})^{-1} \Lambda_{\alpha+3}^{(k+1)} H_{\alpha+3} \end{aligned} \quad (54)$$

The right hand side of (54) is obtained by substituting for $\underline{x}^{(k+1,\alpha+3)}$ from (48).

For regeneration of systematic node $(\alpha + 3)$ the component provided along it by the second non-systematic node is,

$$\begin{aligned} \underline{x}^{(k+2,\alpha+3)} \Lambda_{\alpha+3}^{(k+2)} H_{\alpha+3} \\ = \tilde{\kappa}_1 \tilde{\kappa}_2^{-1} \underline{x}^{(k+2,\alpha+3)} \Lambda_{\alpha+1}^{(k+2)} (\Lambda_{\alpha+1}^{(k+1)})^{-1} \Lambda_{\alpha+3}^{(k+1)} H_{\alpha+3} \end{aligned} \quad (55)$$

The right hand side of (55) is obtained by substituting for $\Lambda_{\alpha+3}^{(k+2)}$ from (53).

From equations (54) and (55), it is clear that components along systematic node $(\alpha + 3)$ node in the vectors passed by the two non-systematic nodes are linearly dependent. Hence by Lemma 4 regeneration of node $(\alpha + 3)$ node is not possible. ■

Since regeneration is not possible by connecting to $k - 1$ systematic nodes and α non-systematic nodes, it will not be possible even in a general setting of using any d nodes for regeneration.

VI. A CODING SCHEME FOR ANY (k, α)

In this section, a coding scheme is described which can be used for any (k, α) parameter set. This scheme assumes that when a systematic node fails, the existing $k - 1$ systematic nodes and any α non-systematic nodes participate in the regeneration. This can be easily extended to a more general case.

A. Scheme Description

Divide the k systematic nodes into α groups. Similar to the scheme given by Wu et al. [2], for regeneration of a systematic node, the existing systematic nodes in the same group as the failed node pass all their α symbols. The remaining systematic nodes and some α non-systematic nodes pass one symbol each.

The structure of the code is as follows. Let $\mu(l) \in \{1, \dots, \alpha\}$ denote the group to which the systematic node l belongs. Consider a set of variables $a_i^{(m)}$ and $b_{i,j}^{(m)}$, for

$m = k + 1, \dots, n$, $i = 1, \dots, k$, $j = 1, \dots, \alpha$, $j \neq \mu(i)$. Let

$$\underline{b}_i^{(m)} = [b_{i,1}^{(m)} \cdots b_{i,\mu(i)-1}^{(m)} \ 0 \ b_{i,\mu(i)+1}^{(m)} \cdots b_{i,\alpha}^{(m)}] \quad (56)$$

Let matrix $B_i^{(m)}$ be an $\alpha \times \alpha$ matrix such that it has $\underline{b}_i^{(m)}$ as its $\mu(i)^{th}$ row, and zeros elsewhere. Also let

$$\tilde{b}_i^{(m)} = [b_{i,1}^{(m)} \cdots b_{i,\mu(i)-1}^{(m)} \ a_i^{(m)} \ b_{i,\mu(i)+1}^{(m)} \cdots b_{i,\alpha}^{(m)}] \quad (57)$$

Let the node matrix of non-systematic node m ($\in \{k + 1, \dots, n\}$) be

$$G_i^{(m)} = a_i^{(m)} I_\alpha + B_i^{(m)} \quad (58)$$

for $i = 1, \dots, k$, where I_α is an $\alpha \times \alpha$ identity matrix.

For example, suppose $k = 5$, $\alpha = 3$ and the systematic nodes are grouped as follows: $\{1, 2\}$, $\{3\}$, $\{4, 5\}$. Then, the node matrix stored by non-systematic node m , ($\in \{k + 1, \dots, n\}$) is

$$\mathbf{G}^{(m)} = \begin{array}{c|c|c|c|c} \begin{array}{ccc|ccc} a_1^{(m)} & b_{1,2}^{(m)} & b_{1,3}^{(m)} & a_2^{(m)} & b_{2,2}^{(m)} & b_{2,3}^{(m)} \\ 0 & a_1^{(m)} & 0 & 0 & a_2^{(m)} & 0 \\ 0 & 0 & a_1^{(m)} & 0 & 0 & a_2^{(m)} \end{array} & \begin{array}{ccc|ccc} a_3^{(m)} & 0 & 0 & a_3^{(m)} & 0 & 0 \\ 0 & a_3^{(m)} & b_{3,3}^{(m)} & 0 & a_4^{(m)} & 0 \\ 0 & 0 & a_3^{(m)} & 0 & 0 & a_3^{(m)} \end{array} & \begin{array}{ccc|ccc} a_4^{(m)} & 0 & 0 & a_4^{(m)} & 0 & 0 \\ 0 & a_4^{(m)} & 0 & 0 & a_4^{(m)} & 0 \\ b_{4,1}^{(m)} & b_{4,2}^{(m)} & a_4^{(m)} & b_{4,1}^{(m)} & b_{4,2}^{(m)} & a_4^{(m)} \end{array} & \begin{array}{ccc|ccc} a_5^{(m)} & 0 & 0 & a_5^{(m)} & 0 & 0 \\ 0 & a_5^{(m)} & 0 & 0 & a_5^{(m)} & 0 \\ b_{5,1}^{(m)} & b_{5,2}^{(m)} & a_5^{(m)} & b_{5,1}^{(m)} & b_{5,2}^{(m)} & a_5^{(m)} \end{array} \end{array} \quad (59)$$

1) *Regeneration*: Consider regeneration of systematic node l ($\in \{1, \dots, k\}$). α non-systematic nodes, say m_1, \dots, m_α pass the $\mu(l)^{th}$ row of their node matrices. The systematic nodes in other groups, say node l' in group $\mu(l')$ ($\mu(l') \neq \mu(l)$), pass the vector $[0 \cdots 0 \ e_{\mu(l)} \ 0 \cdots 0]$ where the unit vector is in the position l' . Since the component along node l' in the vector passed by any non-systematic node is $a_{l'}^{(m)} e_{\mu(l)}$, it can be subtracted out. The existing systematic nodes in group $\mu(l)$ pass all their symbols and hence components along these nodes can also be cancelled out. Hence, for regeneration, the components given out by the non-systematic nodes along the direction of the l^{th} systematic node should be linearly independent. Thus, regeneration condition for systematic node l with this choice of α non-systematic nodes reduces to a polynomial being non-zero i.e

$$\det \begin{pmatrix} \tilde{b}_l^{(m_1)} \\ \tilde{b}_l^{(m_2)} \\ \vdots \\ \tilde{b}_l^{(m_\alpha)} \end{pmatrix} \quad (60)$$

Similar polynomials are obtained $\forall l$, and for all sets of α non-systematic nodes. Clearly, none of these polynomials are identically zero.

2) *Reconstruction*: If the data collector connects to the k systematic nodes, then reconstruction is trivially satisfied. Consider DC connecting to p non-systematic nodes, and $k - p$ systematic nodes, $1 \leq p \leq k$. Let m_1, \dots, m_p , ($m_1 < \dots < m_p$) be the non-systematic nodes to which it connects. Let l_1, \dots, l_p , ($l_1 < \dots < l_p$) be the p systematic nodes to which it does *not* connect. As in Section III-B2, reconstruction condition leads to following polynomial not equal to zero




Fig. 4. Average repair bandwidth(γ) required for exact regeneration of the systematic nodes with $\beta = 1$ is plotted for various values of d for $k = 9$.

condition.

$$\det \begin{pmatrix} G_{l_1}^{(m_1)} & G_{l_2}^{(m_1)} & \dots & G_{l_p}^{(m_1)} \\ G_{l_1}^{(m_2)} & G_{l_2}^{(m_2)} & \dots & G_{l_p}^{(m_2)} \\ \vdots & \vdots & \dots & \vdots \\ G_{l_1}^{(m_p)} & G_{l_2}^{(m_p)} & \dots & G_{l_p}^{(m_p)} \end{pmatrix} \quad (61)$$

We will now show that there exists an assignment of the variables such that this polynomial is not identically zero. Set

$$b_i^{(m)} = 0 \quad \forall i, m \quad (62)$$

$$a_{l_i}^{(m_j)} = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{if } i \neq j \end{cases} \quad (63)$$

By these assignments, the reconstruction matrix becomes an identity matrix, which is non-singular. Thus, the regeneration and reconstruction properties evaluate to the condition of the product of certain polynomials being non-zero. It is shown that none of these polynomials is identically zero. Assignment of values to the variables satisfying all the conditions can be obtained using the algorithm given by Koetter and Medard [4].

This scheme can be extended to regeneration using any combination of systematic and non-systematic nodes provided that the systematic nodes in the same group as the failed node participate in regeneration. The extended proof will involve a few more conditions of polynomials being non-zero.

B. Analysis

For $k \leq \alpha$, if all the α nodes are kept in different groups, this scheme achieves the minimum repair bandwidth and hence is optimal.

For $k > \alpha$, the amount of data to be downloaded for exact regeneration of a systematic node depends on the number of nodes in its group. If there are η nodes in a group, the total number of symbols required to regenerate a node in that group, is given by:

$$\gamma = (\eta - 1)\alpha + (d - \eta + 1) \quad (64)$$

Lemma 12: The average repair bandwidth for exact regeneration of systematic nodes using the above described scheme is minimum when the groups are uniformly divided.

Proof: Directly follows from equation (64) \blacksquare

Let

$$s = \lfloor k/\alpha \rfloor \quad (65)$$

Uniform division of groups would imply that out of the α groups, $k \bmod \alpha$ groups contain $s + 1$ nodes each and the rest contain s nodes each.

The average amount of download required for exact regeneration of the systematic nodes in our scheme is compared with the scheme proposed by Wu and Dimakis [2] (*Group interference alignment*) in Figure 4. The lower bound on the repair bandwidth is also plotted along side. It can be seen that for $d \geq 2k - 1$ (i.e. $k \leq \alpha$) our scheme achieves the lower bound. For smaller values of d , the amount of data downloaded is higher. However, whether this achieved value of repair bandwidth is optimal or not is not known for $d \leq 2k - 4$.

REFERENCES

- [1] Y. Wu, A. G. Dimakis and K. Ramchandran, "Deterministic Regenerating codes for Distributed Storage," *Proc. Allerton Conf.*, Sep. 2007.
- [2] Y. Wu and A. Dimakis, "Reducing Repair Traffic for Erasure Coding-Based Storage via Interference Alignment," *Proc. ISIT*, Jul. 2009.
- [3] K. V. Rashmi, Nihar B. Shah, P. Vijay Kumar and Kannan Ramchandran, "Explicit Construction of Optimal Exact Regenerating Codes for Distributed Storage," to appear in *Proc. Allerton Conf.*, Sep 2009. Available online at arXiv:0906.4913 [cs.IT]
- [4] Ralf Koetter and Muriel Medard, "An algebraic approach to network coding," *IEEE/ACM Transactions on Networking*, v.11 n.5, p.782-795, Oct. 2003.
- [5] Dennis S. Bernstein, *Matrix mathematics: Theory, facts, and formulas with application to linear systems theory*, Princeton University Press, Princeton, NJ, p.119, 2005.

APPENDIX

Proof of Theorem 1

Proof: Let $\omega_1, \dots, \omega_{k-p}$ ($\omega_1 < \dots < \omega_{k-p}$) be the systematic nodes to which DC connects, and $\Omega_1, \dots, \Omega_p$ ($\Omega_1 < \dots < \Omega_p$) be the p systematic nodes to which it does not connect. Thus the sets $\omega_1, \dots, \omega_{k-p}$ and $\Omega_1, \dots, \Omega_p$ are disjoint. Let $\delta_1, \dots, \delta_p$ be the p non-systematic nodes to which DC connects. The reconstruction matrix R is given by

$$R = \begin{bmatrix} \mathbf{G}'^{(\delta_1)} \\ \vdots \\ \mathbf{G}'^{(\delta_p)} \end{bmatrix} = \begin{bmatrix} G_{\Omega_1}^{(\delta_1)} & G_{\Omega_2}^{(\delta_1)} & \dots & G_{\Omega_p}^{(\delta_1)} \\ \vdots & \vdots & \dots & \vdots \\ G_{\Omega_1}^{(\delta_p)} & G_{\Omega_2}^{(\delta_p)} & \dots & G_{\Omega_p}^{(\delta_p)} \end{bmatrix} \quad (66)$$

Group the Ω_1^{th} rows of $\mathbf{G}'^{(\delta_m)}$ ($m = 1, \dots, p$) as the first p rows of a new matrix R' , then Ω_2^{th} rows as the next p rows, and so on. Hence, row number Ω_i of $\mathbf{G}'^{(\delta_m)}$ becomes the row number $p \times (i-1) + m$ in R' . Below these, group the ω_1^{th} rows, then the ω_2^{th} and so on. Row number ω_i of $\mathbf{G}'^{(\delta_m)}$ becomes the row number $p^2 + p \times (i-1) + m$ in R' . Hence there are α groups with p rows each in R' .

Let S be an $p \times \alpha$ matrix with elements $[S]_{i,j} = \psi_j^{(\delta_i)}$, $i = 1, \dots, p$, $j = 1, \dots, \alpha$. Let $T_{a,b}$ be an $p \times \alpha$ matrix with its b^{th} column as $[\psi_a^{(\delta_1)}, \dots, \psi_a^{(\delta_p)}]^t$, and rest of the elements zero. Thus, the b^{th} column of $T_{a,b}$ is identical to the a^{th} column of S .

The columns of R' are grouped into p groups of α columns each. Thus the matrix R' can be viewed as a block matrix, with each block of size $p \times \alpha$, and the dimension of R' being $\alpha \times p$ blocks.

Let $[R']_{(i,j)}$ represent the $(i,j)^{th}$ block of R' . For $i = 1, \dots, p$, $j = 1, \dots, p$ we get

$$[R']_{(i,j)} = \begin{cases} \epsilon S & \text{if } i = j \\ T_{\Omega_j, \Omega_i} & \text{if } i \neq j \end{cases} \quad (67)$$

For $i = p+1, \dots, \alpha$, $j = 1, \dots, p$,

$$[R']_{(i,j)} = T_{\Omega_j, \omega_{i-p}} \quad (68)$$

Thus,

$$R' = \left[\begin{array}{cccc|cccc} \epsilon S & T_{\Omega_2, \Omega_1} & T_{\Omega_3, \Omega_1} & \cdots & T_{\Omega_p, \Omega_1} & & & \\ T_{\Omega_1, \Omega_2} & \epsilon S & T_{\Omega_3, \Omega_2} & \cdots & T_{\Omega_p, \Omega_2} & & & \\ \vdots & \vdots & \vdots & & \vdots & & & \\ T_{\Omega_1, \Omega_p} & T_{\Omega_2, \Omega_p} & T_{\Omega_3, \Omega_p} & \cdots & \epsilon S & & & \\ \hline T_{\Omega_1, \omega_1} & T_{\Omega_2, \omega_1} & T_{\Omega_3, \omega_1} & \cdots & T_{\Omega_p, \omega_1} & & & \\ \vdots & \vdots & \vdots & & \vdots & & & \\ T_{\Omega_1, \omega_{k-p}} & T_{\Omega_2, \omega_{k-p}} & T_{\Omega_3, \omega_{k-p}} & \cdots & T_{\Omega_p, \omega_{k-p}} & & & \end{array} \right] \quad (69)$$

Let \tilde{S} be the $p \times p$ matrix formed by the columns $\Omega_1, \dots, \Omega_p$ of S . As \tilde{S} is a submatrix of Cauchy matrix Ψ , it is invertible. Let \hat{I} be an $p \times \alpha$ matrix with columns $\omega_1, \dots, \omega_{k-p}$ having some arbitrary values (denoted by ϕ), and the remaining p columns put together forming an identity matrix. Let $\hat{E}_{a,b}$ be an $p \times \alpha$ matrix with the element at position (a,b) as 1 and all other elements 0.

Multiply each of the α groups of p rows by \tilde{S}^{-1} . This will cause the following transformation in R' : S will be replaced by \hat{I} and $T_{\Omega_a, b}$ will be replaced by $\hat{E}_{a,b}$.

The resultant matrix will be of the form:

$$\left[\begin{array}{cccc|cccc} \epsilon \hat{I} & \hat{E}_{2, \Omega_1} & \hat{E}_{3, \Omega_1} & \cdots & \hat{E}_{p, \Omega_1} & & & \\ \hat{E}_{1, \Omega_2} & \epsilon \hat{I} & \hat{E}_{3, \Omega_2} & \cdots & \hat{E}_{p, \Omega_2} & & & \\ \vdots & \vdots & \vdots & & \vdots & & & \\ \hat{E}_{1, \Omega_p} & \hat{E}_{2, \Omega_p} & \hat{E}_{3, \Omega_p} & \cdots & \epsilon \hat{I} & & & \\ \hline \hat{E}_{1, \omega_1} & \hat{E}_{2, \omega_1} & \hat{E}_{3, \omega_1} & \cdots & \hat{E}_{p, \omega_1} & & & \\ \vdots & \vdots & \vdots & & \vdots & & & \\ \hat{E}_{1, \omega_{k-p}} & \hat{E}_{2, \omega_{k-p}} & \hat{E}_{3, \omega_{k-p}} & \cdots & \hat{E}_{p, \omega_{k-p}} & & & \end{array} \right] \quad (70)$$

In the groups of rows $p+1, \dots, \alpha$, every row has exactly one non-zero element. Hence these rows and the corresponding columns $(\omega_1, \dots, \omega_{k-p})$ are independent of all others and can be eliminated. Note that all the ϕ elements are present only in these columns and hence the actual values of ϕ do not matter. The resultant matrix will be a $p^2 \times p^2$ matrix of the following form:

$$\left[\begin{array}{cccc|cccc} \epsilon I_p & E_{2,1} & E_{3,1} & \cdots & E_{p,1} & & & \\ E_{1,2} & \epsilon I_p & E_{3,2} & \cdots & E_{p,2} & & & \\ \vdots & \vdots & \vdots & & \vdots & & & \\ E_{1,p} & E_{2,p} & E_{3,p} & \cdots & \epsilon I_p & & & \end{array} \right] \quad (71)$$

where I_p is a $p \times p$ identity matrix and $E_{a,b}$ is an $p \times p$ matrix with the element in the position (a,b) as 1 and all

other elements 0.

For $i = 1, \dots, p$, the i^{th} row(column) of the i^{th} row(column) group respectively contains exactly one non-zero element, and hence is linearly independent of all others. After eliminating these rows (and corresponding columns) the remaining matrix is rearranged by placing the i^{th} row(column) of the j^{th} group adjacent to the j^{th} row(column) of the i^{th} group to form:

$$\left[\begin{array}{cccc|cccc} \epsilon & 1 & 0 & 0 & \cdots & 0 & 0 & \\ 1 & \epsilon & 0 & 0 & \cdots & 0 & 0 & \\ 0 & 0 & \epsilon & 1 & \cdots & 0 & 0 & \\ 0 & 0 & 1 & \epsilon & \cdots & 0 & 0 & \\ \vdots & \vdots & \vdots & \vdots & & \vdots & \vdots & \\ 0 & 0 & 0 & 0 & \cdots & \epsilon & 1 & \\ 0 & 0 & 0 & 0 & \cdots & 1 & \epsilon & \end{array} \right] \quad (72)$$

This is a block diagonal matrix, and since $\epsilon^2 \neq 1$, is full rank. ■

In the example of $k = \alpha = 3$ considered in section III-A2 when the data collector connected to the first systematic node, and the first two non-systematic nodes, we have $p = 2$, $\omega_1 = 1$, $\Omega_1 = 2$, $\Omega_2 = 3$, $\delta_1 = 4$, $\delta_2 = 5$ and $\epsilon = 2$. Here,

$$\begin{aligned} D_2 &= R, \\ S &= \begin{bmatrix} \psi_1^{(4)} & \psi_2^{(4)} & \psi_3^{(4)} \\ \psi_1^{(5)} & \psi_2^{(5)} & \psi_3^{(5)} \end{bmatrix}, \\ T_{\Omega_1, \Omega_2} &= \begin{bmatrix} 0 & 0 & \psi_2^{(4)} \\ 0 & 0 & \psi_2^{(5)} \end{bmatrix}, \\ \hat{I} &= \begin{bmatrix} \phi & 2 & 0 \\ \phi & 0 & 2 \end{bmatrix}, \\ \hat{E}_{1, \Omega_2} &= \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \end{bmatrix} \text{ and} \\ E_{1,2} &= \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}. \end{aligned}$$