Shannon Inequalities in Distributed Storage - Part II

Birenjith Sasidharan and P. Vijay Kumar
(joint work with: Myna Vajha and Kaushik Senthoor)

Department of Electrical Communication Engineering,
Indian Institute of Science, Bangalore

Workshop on Advanced Information Theory
Commemorating the 100th Birthday of Claude Shannon

Organized by the IISc-IEEE ComSoc Student Chapter
(in association with IEEE Bangalore Section & ECE Department, IISc)

Indian Institute of Science, April 30, 2016

Outline of the Talk

1. Regenerating Codes
2. Interior Points of the Storage-Repair Bandwidth Tradeoff

An Example Erasure (RAID) Code

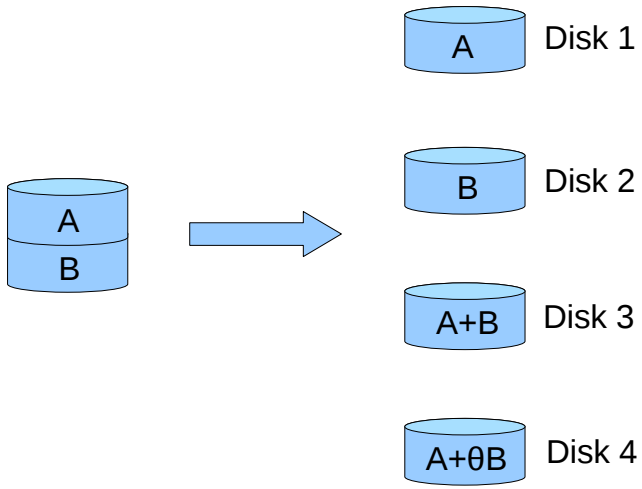

(4, 2) MDS code

Used in RAID 6

- $[4, 2]$ MDS code
- Can recover data by connecting to any 2 nodes

(RAID: Redundant Array of Independent Disks)

RAID Codes Not Very Efficient at Handling Node Repair

Approach to node repair:

- Connect to any k nodes,
- Reconstruct entire data file,
- Reconstruct data stored in the node

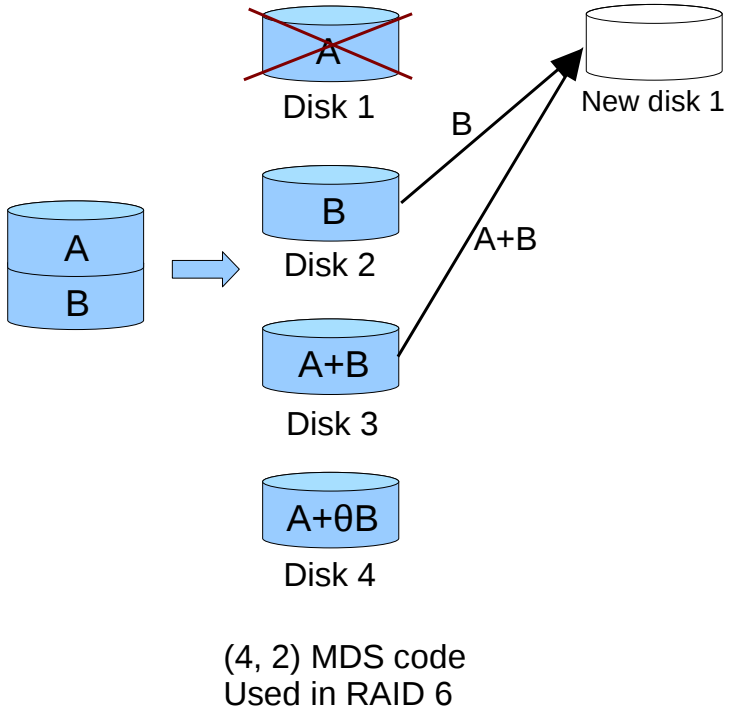

But downloading 2 units of data to revive a node that stores 1 unit of data is wasteful!

Regenerating Codes

Dimakis, Godfrey, Wu,Wainwright, Ramchandran, T-IT, Sep. 2010

An Example Regenerating Code

- Each node now stores two "half-symbols"
- We download 3 half-symbols as opposed to 2 full-symbols
  - vector symbol alphabet $\Rightarrow$ $\mathbb{F}_q^2$ versus $\mathbb{F}_{q^2}$

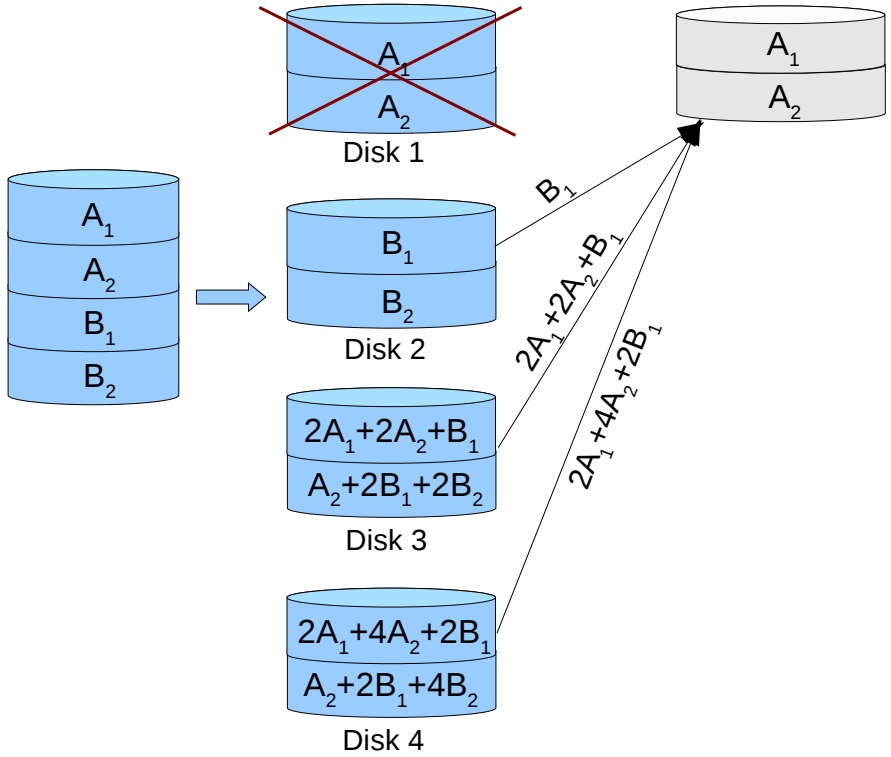

- File size $B = 4$
- $(n = 4, k = 2, d = 3)$
- $(\alpha = 2, \beta = 1)$

Regenerating Codes - Formal Definition

Parameters: $(\ (n,k,d),\ (\alpha,\beta),\ B,\ \mathbb{F}_q\ )$

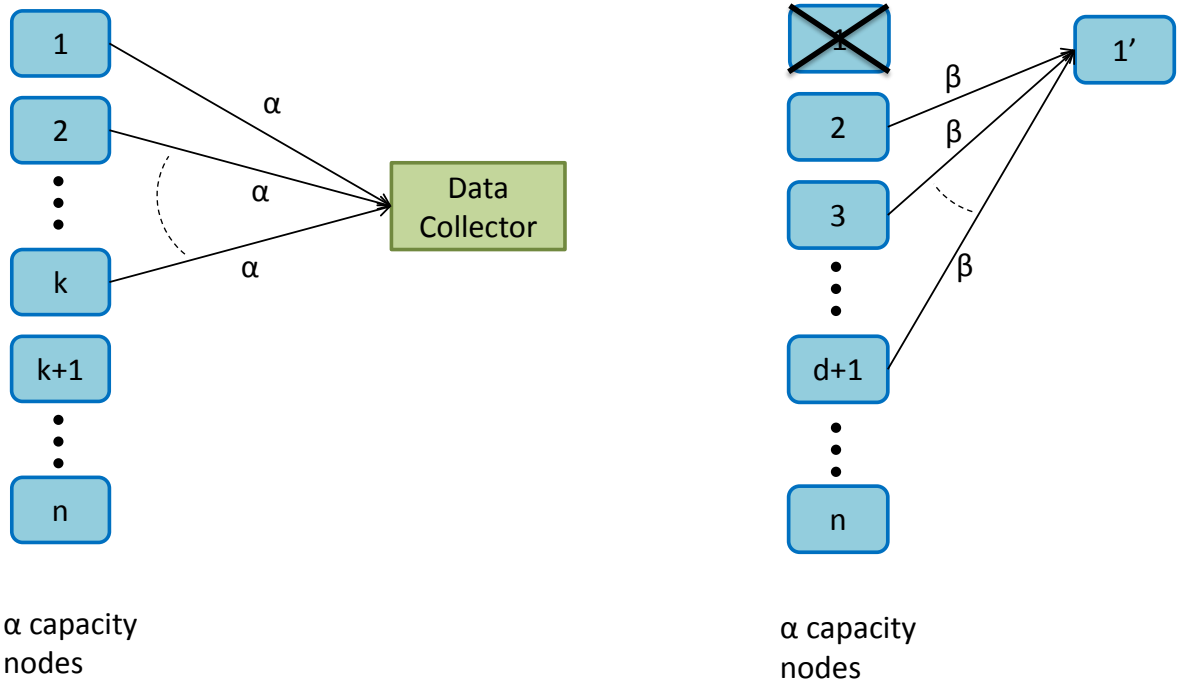

- Data to be recovered by connecting to any k of n nodes
- Nodes to be repaired by connecting to any d nodes, downloading β symbols from each node; ($d\beta << $ file size B)
- Difference between functional and exact repair (FR vs ER).

Cut-Set Bound from Network Coding

Given code parameters $\{(n, k, d), (\alpha, \beta)\}$:

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

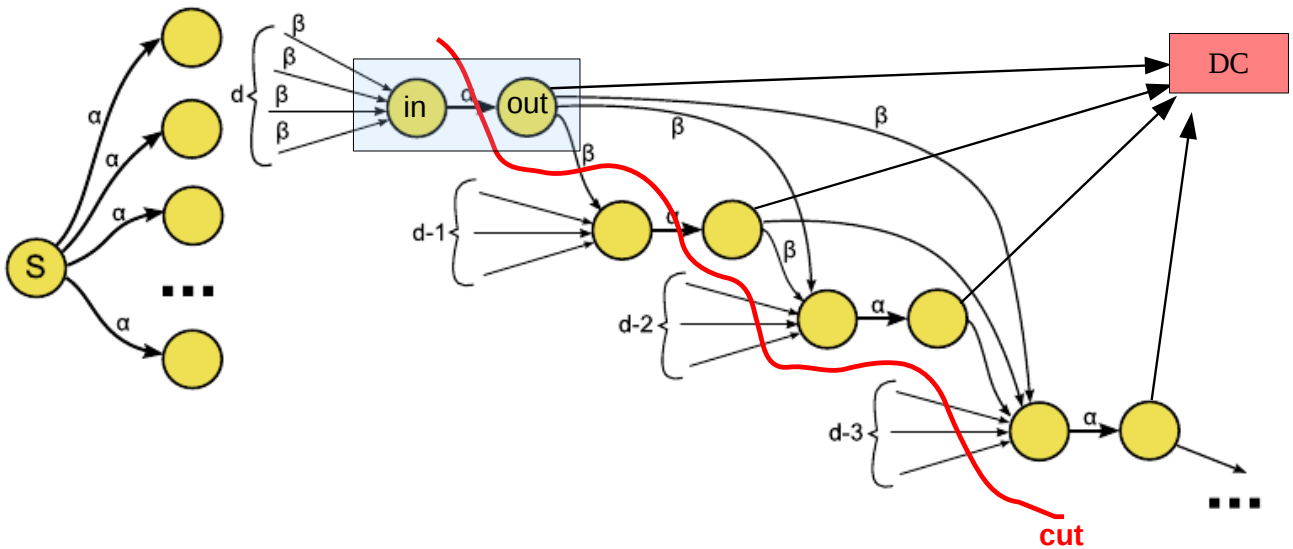

(can be shown to be achievable under functional repair)

Dimakis, Godfrey, Wu,Wainwright, Ramchandran, T-IT, Sep. 2010
Wu, JSAC, 2010

The Storage-Repair Bandwidth Tradeoff

The upper bound on file size:

$$B \leq \sum_{i=1}^{k} \min\{\alpha, (d-i+1)\beta\} \quad \text{(multiple } (\alpha, \beta) \text{ pairs can achieve bound)}$$

- Tradeoff curve drawn for fixed $(k, d), B$.
- Extreme points: MSR (Minimum Storage Regenerating) & MBR (Minimum Bandwidth Regenerating)
  - $\alpha_{\text{msr}} = (d-k+1)\beta_{\text{msr}}$, $\alpha_{\text{msr}} = \frac{B}{k}$
  - $\alpha_{\text{mbr}} = d\beta_{\text{mbr}}$, $\beta_{\text{mbr}} = \frac{B}{dk-\binom{k}{2}}$

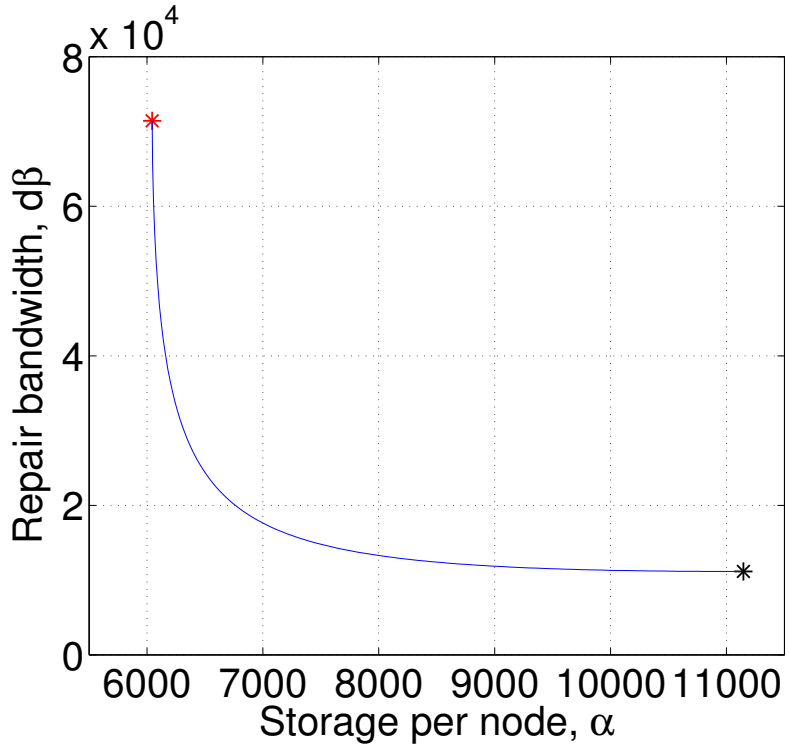

$(k, d) = (120, 129), B = 725360$

INTERIOR POINTS OF THE TRADEOFF

Interior Points Not-Achievable Under Exact Repair!

No exact-repair code can achieve an interior point on the tradeoff...

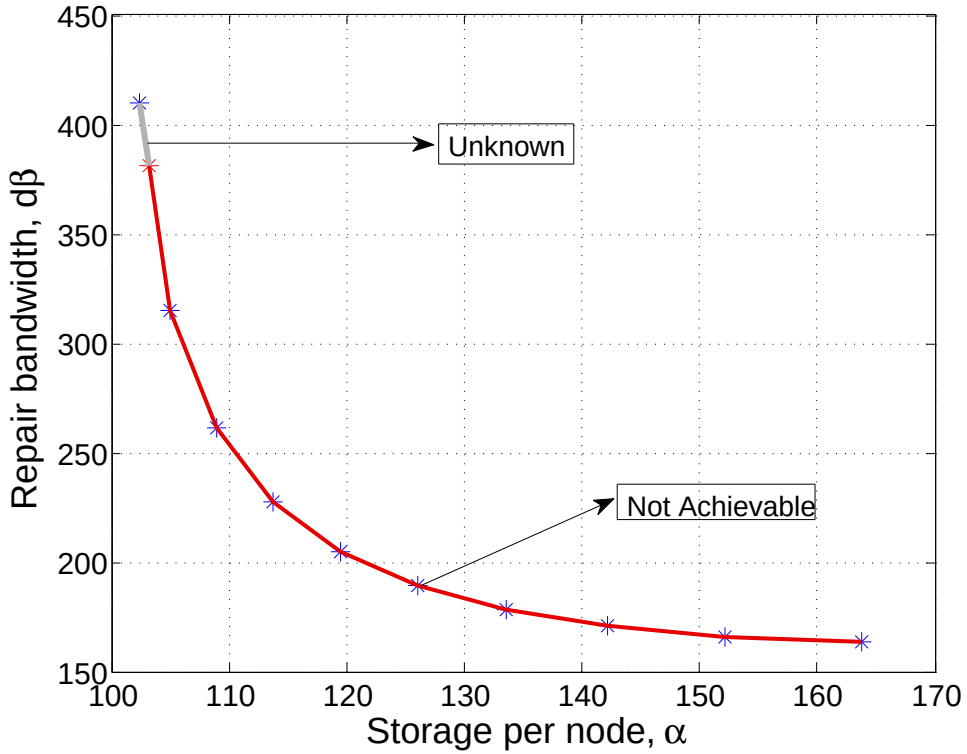

(Qn: Can exact-repair codes approach the tradeoff asymptotically, i.e., as $B \to \infty$?)

(Shah, Rashmi, K, Ramchandran, Mar. 2012)

Interior Points Not-Achievable Under Exact Repair!

No exact-repair code can achieve an interior point on the tradeoff...

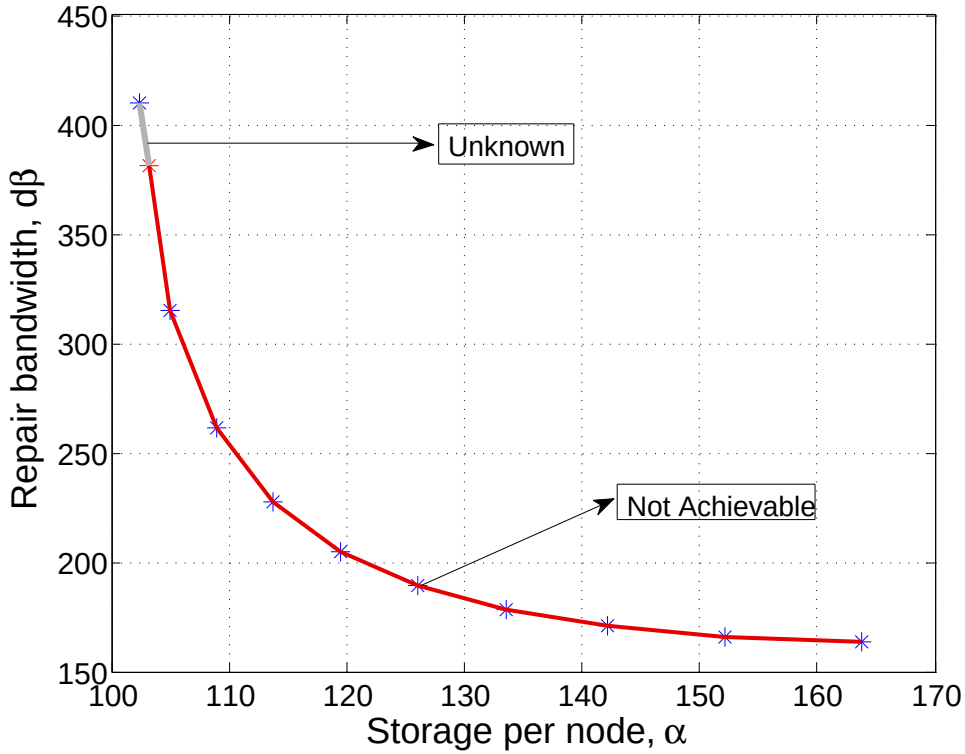

(Qn: Can exact-repair codes approach the tradeoff asymptotically, i.e., as $B \to \infty$?)

(Shah, Rashmi, K, Ramchandran, Mar. 2012)

Non-Existence Proof: Background

1. MSR $\Leftrightarrow$ $\alpha = (d - k + 1)\beta$, MBR $\Leftrightarrow$ $\alpha = d\beta$
2. Interior: $\Leftrightarrow$ $\alpha = (d - \mu)\beta$, $1 \leq \mu \leq (k - 2)$

We assume wolog $(n = d + 1)$, as restriction to $(d + 1)$ nodes is also a regenerating code:

Parameters: ($(n' = d + 1, k, d)$, (α, β), B, $\mathbb{F}_q$) (fixed throughout)

Notation

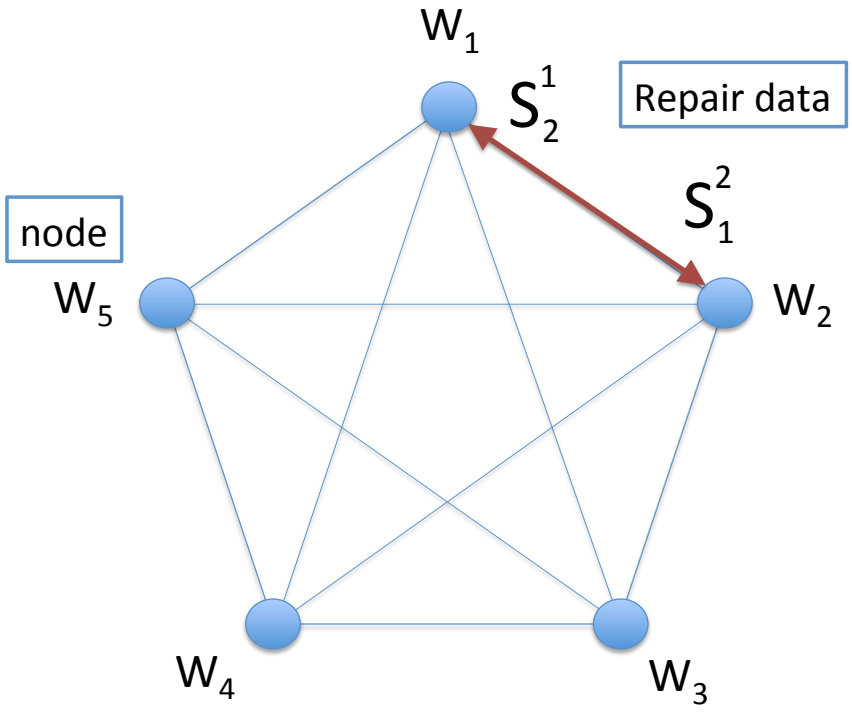

- n nodes
- ith node stores α symbols, random variable W_i
- $S_z^y \Rightarrow$ the β symbols sent from z to repair node y

The Repair Matrix $\mathcal{R}$

- S_x^y is the repair data sent from node x to node y

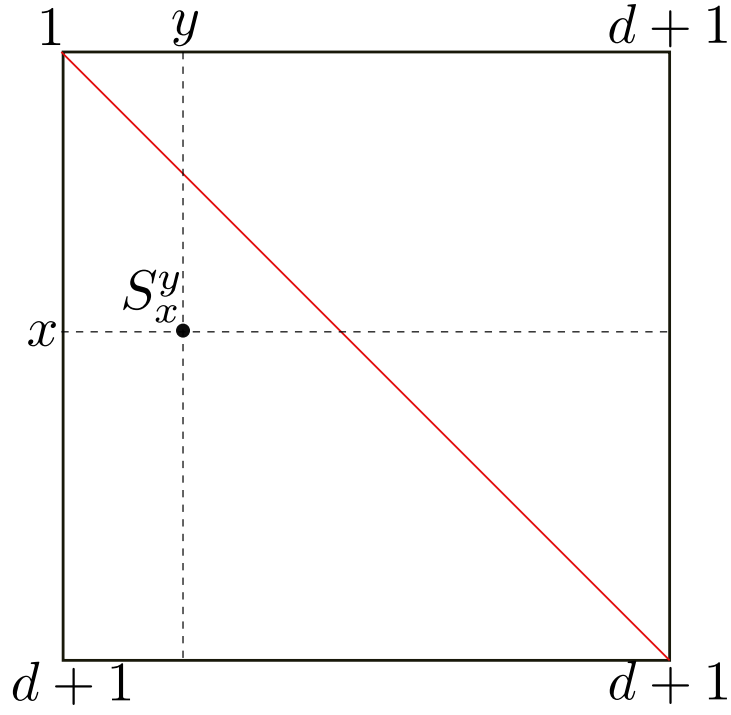

$$\left((d+1) \times (d+1) \right)$$

Setting Up the Random Variables

Alphabet of data file: $\mathcal{A}_B$

Alphabet of node data: $\mathcal{A}_W$

Alphabet of repair data: $\mathcal{A}_S$

$$
\begin{aligned}
\text{(Encoding)} \quad & e_i \quad : \quad \mathcal{A}_B \longrightarrow \mathcal{A}_W, \ i = 1, 2, \ldots, n. \\
\text{(Repair data)} \quad & h_i^j \quad : \quad \mathcal{A}_W \longrightarrow \mathcal{A}_R, \ i \neq j, \ i, j \in [n]. \\
\text{(Regeneration)} \quad & r_i \quad : \quad \mathcal{A}_R^d \longrightarrow \mathcal{A}_W, \ i = 1, 2, \ldots, n. \\
\text{(Datacollection)} \quad & d_i \quad : \quad \mathcal{A}_W^k \longrightarrow \mathcal{A}_B, \ i = 1, 2, \ldots, \binom{n}{k}.
\end{aligned}
$$

$$
B := \log(\mathcal{A}_B), \quad \alpha := \log(\mathcal{A}_W), \quad \beta := \log(\mathcal{A}_R)
$$

Setting Up the Random Variables

- M picked uniformly at random from $\mathcal{A}_B$.
- M induces distributions on:

$$
\begin{aligned}
W_i &= e_i(M),\ i = 1, 2, \ldots, n. \\
S_i^j &= h_i^j(W_i),\ i \neq j,\ i, j \in [n]
\end{aligned}
$$

$$
\begin{aligned}
W_X &= \{W_x \mid x \in X\} \\
S_Z^Y &= \{S_z^y \mid z \in Z,\ y \in Y\}
\end{aligned}
$$

$(n + n(n-1) = n^2)$ random variables in all.

Constraints

$H(W_i) \leq \alpha,$	entropy of ith node
$H(S_i^j) \leq \beta,$	entropy of repair data

$H(M \mid W_A) = 0, \mid A \mid = k,$	data collection property
$H(W_i \mid M) = 0,$	node contents a function of file data
$H(S_i^j \mid W_i) = 0,$	repair data draws from node contents
$H(W_i \mid S_A^i) = 0,$ if $i \notin A, \mid A \mid = d$	repair property

Domain Constraints Viewed on Entropic Vector Space

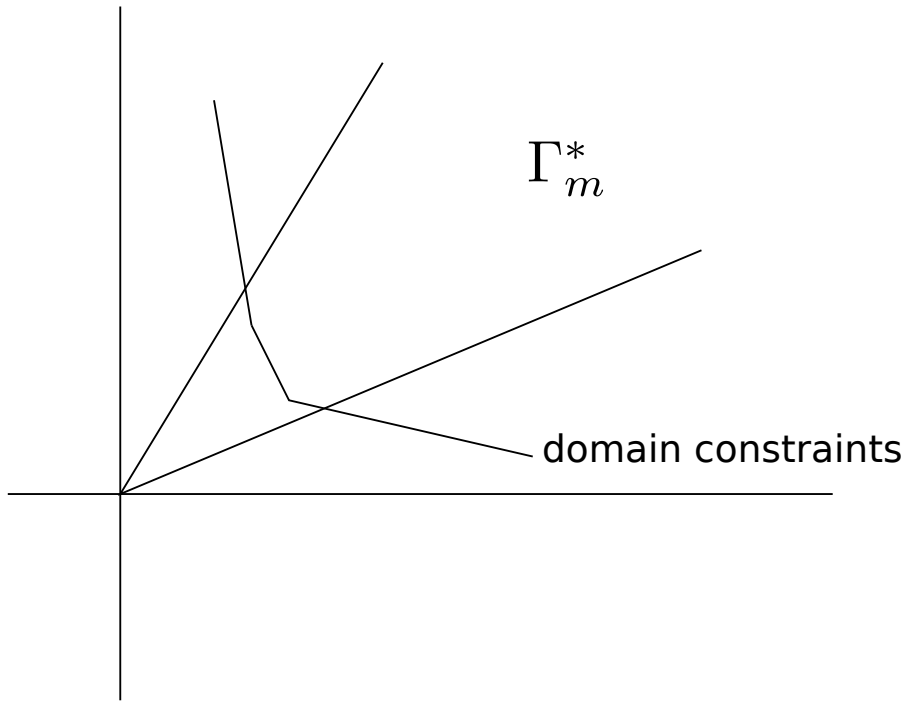

The domain constraints limit the permissible entropic vectors, and therefore limit $H(W_{[k]})$ to much less a value than $k\alpha$.

Exact-Repair File Size Bound

Consider ER code with $\alpha = (d-\mu)\beta,\ 1 \le \mu \le k-2$.

$$\begin{aligned} [d+1] &= X \,\dot\cup\, Y \,\dot\cup\, Z \\ |\,X\,| &= \mu + 1 \\ |\,Y\,| &= k - (\mu+1) \\ |\,Z\,| &= (d+1-k) \end{aligned}$$

Then

$$\begin{aligned} B &= H(W_X, S_Y, S_Z^Y) \\ S_Y &= \{S_i^j \mid i,j \in Y, i > j\} \\ S_Z^Y &= \{S_z^y \mid z \in Z, y \in Y\} \end{aligned}$$

Exact-Repair File Size Bound

Then we have

$$
\begin{aligned}
B &= H(W_X, S_Y, S_Z^Y) \\
&= H(W_X) + H(S_Y \mid W_X) + H(S_Z^Y \mid W_X, S_Y) \\
&\leq H(W_X) + H(S_Y) + H(S_Z^Y) && (1) \\
&\leq |X|\alpha + |S_Y|\beta + |S_Z^Y|\beta && (2) \\
&= (\mu+1)\alpha + \binom{k-\mu-1}{2}\beta + (d-k+1)(k-\mu-1)\beta.
\end{aligned}
$$

Exact-Repair File Size Bound

The optimal FR file size when $\alpha = (d-\mu)\beta$:

$$
\begin{aligned}
B_{\text{opt-fr}} &= \sum_{i=0}^{k-1} \min\{\alpha, (d-i)\beta\} \\
&= (\mu+1)\alpha \; + \sum_{i=\mu+1}^{k-1} (d-i)\beta \\
&= (\mu+1)\alpha \; + \sum_{j=0}^{k-\mu-2} (d-k+1+j)\beta \\
&= (\mu+1)\alpha \; + \binom{k-\mu-1}{2}\beta \; + \; (d-k+1)(k-\mu-1)\beta.
\end{aligned}
$$

So we must have equalities in (1), (2), if ER code were to achieve optimal FR file size. i.e.,

$$
B = |X|\alpha + |S_Y|\beta + |S_Z^Y|\beta.
$$

Non-Existence via Properties of the Repair Matrix $\mathcal{R}$

Assuming the existence of an optimal exact-repair code, we must have:

$$B = H(W_X, S_Y, S_Z^Y) \;\; = \;\; |\, X \,|\, \alpha + \;\; |\, S_Y \,|\, \beta + \;\; |\, S_Z^Y \,|\, \beta.$$

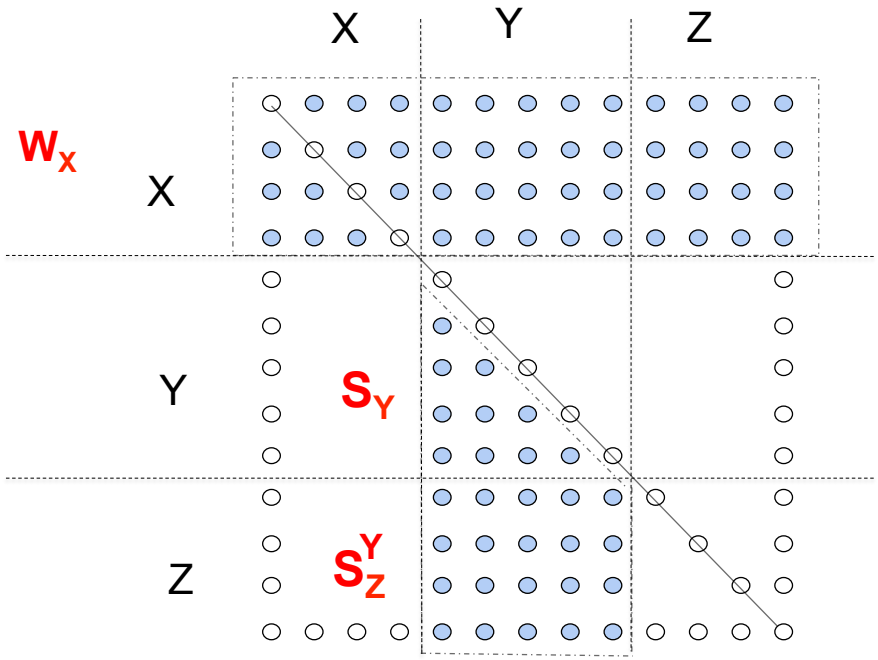

- Turns out however, every row of $\mathcal{R}$ has entropy at most β - contradiction!

Shah, Rashmi, K, Ramchandran, T-IT, 2012.

Explaining Why Rows Have Small Entropy

Goal: Explain why every row of $\mathcal{R}$ has entropy at most β. In figure below, $p = (\mu + 1)$.

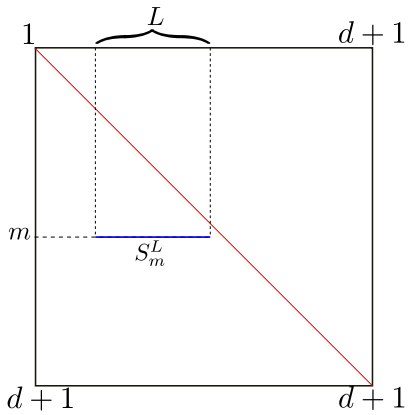

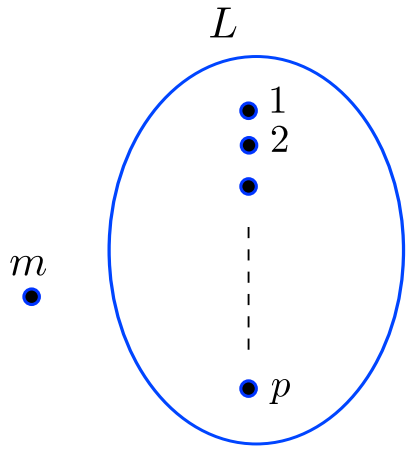

$$H(S_m^L) = \underbrace{H(S_m^L \mid W_L)}_{\leq(\mu+1)H(S_m^{\ell_0}|W_L)=0} + \underbrace{I(S_m^L : W_L)}_{\leq I(W_m;W_L)\leq\beta} \leq \beta.$$

Because $\mid L \mid = (\mu + 1)$ is

- large enough to permit interference cancellation to take place while passing repair information
- small enough that the mutual information is limited by β

The Computation

Step 1: If ER code achieves the optimal FR file size, then

$$
\begin{aligned}
H(W_{i+1} \mid W_{[i]}) &= \min\{\alpha, (d-i)\beta\} \\
&= \begin{cases} \alpha & 0 \le i \le \mu \\ \alpha - (i-\mu)\beta & \mu < i \le k-1 \\ 0 & k \le i < n \end{cases}
\end{aligned}
$$

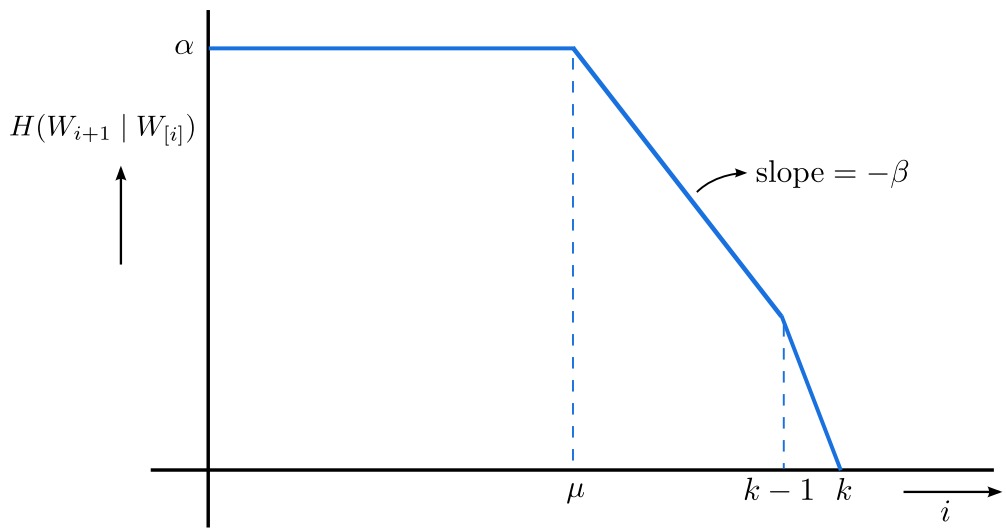

The Computation

Step 2: Let m, L be chosen such that $|L| = (\mu + 1)$, $m \notin L$, $\ell_0 \in L$, $\bar{L} = L \setminus \{\ell_0\}$. Then

$$
\begin{aligned}
H(S_m^L) &= H(S_m^L \mid W_L) + I(S_m^L : W_L) \\
&\leq \ell H(S_m^{\ell_0} \mid W_L) + I(W_m : W_L) \\
&\leq \ell \left\{ H(S_m^{\ell_0} \mid W_{\bar{L}}) + H(W_{\ell_0} \mid S_m^{\ell_0}, W_{\bar{L}}) - H(W_{\ell_0} \mid W_{\bar{L}}) \right\} + \{H(W_m) - H(W_m \mid W_L)\} \\
&\leq \ell \underbrace{\{\beta + (\alpha - \beta) - \alpha\}}_{=0} + \underbrace{\{\alpha - (\alpha - \beta)\}}_{=\beta} \\
&= \beta.
\end{aligned}
$$

CAN AN INTERIOR POINT BE APPROACHED ?

No! From Characterization of the $(4, 3, 3)$ Tradeoff

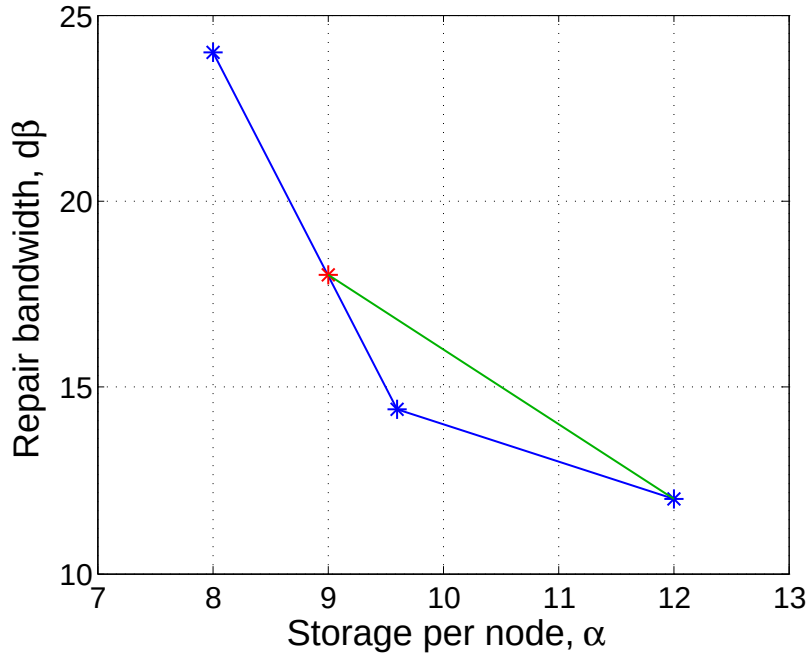

- FR Tradeoff = Blue
- ER Tradeoff = Max{Blue, Green}
- Chao Tian provides an explicit proof by using Raymond Yeung's ITIP framework to extract an additional inequality for the $(4, 3, 3)$ case.

Chao Tian, ISIT 2013

Our Subsequent Results (2014)

$(4, 3, 3)$ Case

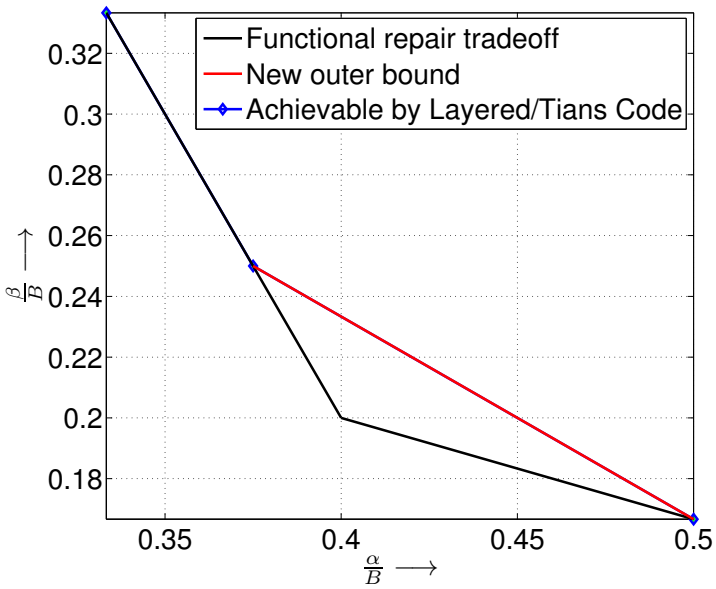

$(5, 4, 4)$ Case

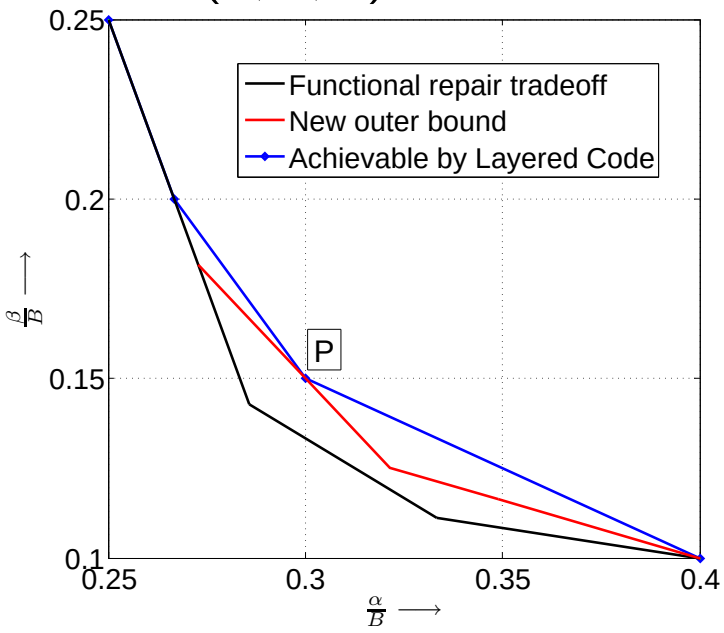

- First outer bound on the ER tradeoff that improves upon the FR tradeoff for all $[n, k, d]$
- Coincides with the ER tradeoff characterized by Tian for the $[4, 3, 3]$ case
- Shown alongside is the outer bound in the $[5, 4, 4]$ case
- In the $[5, 4, 4]$ case, bound coincides at one point P with performance of a layered code.
- First instance of an optimal code operating off of the FR tradeoff.

Layered code: (Tian, Sasidharan, Aggarwal, Vaishampayan, K, Apr 2015)

Our Approach

Let $\mathcal{T}$ denote the 'trapezium-shaped' region of the repair matrix:

$$\mathcal{T} \;=\; S_Y \,\dot\cup\, S_Z^Y \;\subseteq\; \mathcal{R}$$

Assuming the existence of an optimal exact-repair code, we must have:

$$H(\mathcal{T}) \;=\; |\,\mathcal{T}\,|\,\beta$$

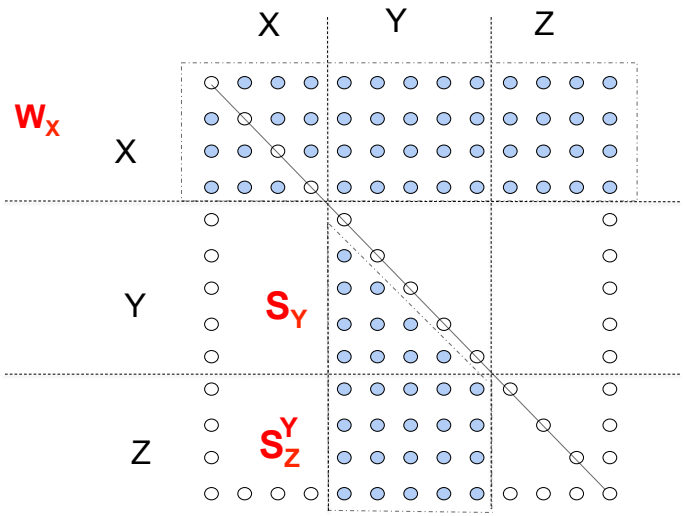

- On the other hand, every row of $\mathcal{T}$ has entropy at most β, this is a large gap which we exploit!

Birenjith, Senthoor, Kumar, ISIT 2014.

Approach to Deriving the New Bound

Decreasing file size to B-ε moves the bounds close together until contradiction is resolved

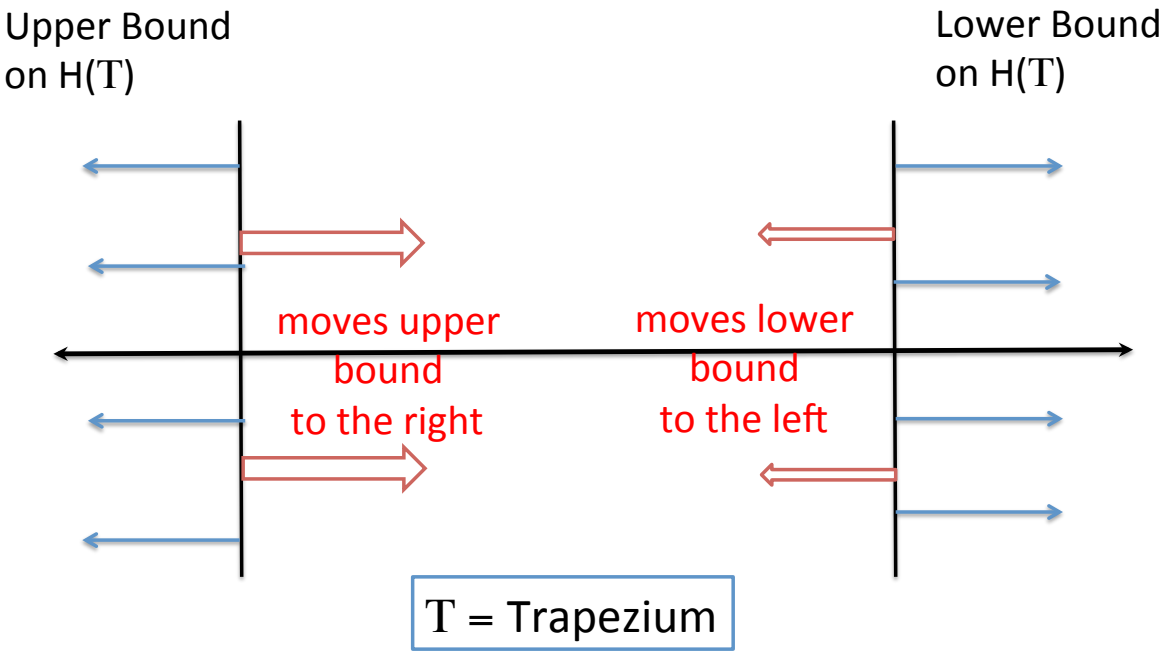

Upper bound on $H(\mathcal{T})$

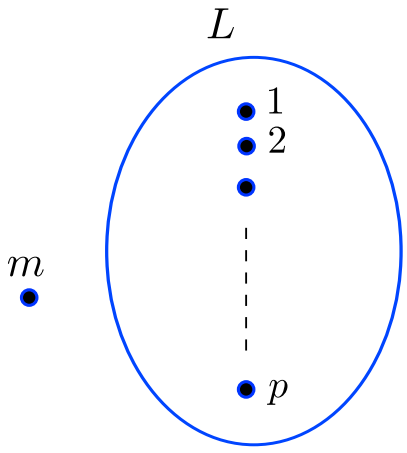

We have chosen m, L such that $|L| = (\mu + 1)$, $m \notin L$, $\ell_0 \in L$, $\bar{L} = L \setminus \{\ell_0\}$. Then

$$
\begin{aligned}
H(S_m^L) &= H(S_m^L \mid W_L) + I(S_m^L : W_L) \\
&\leq \ell H(S_m^{\ell_0} \mid W_L) + I(W_m : W_L) \\
&\leq \ell \left\{ H(S_m^{\ell_0} \mid W_{\bar{L}}) + H(W_{\ell_0} \mid S_m^{\ell_0}, W_{\bar{L}}) - H(W_{\ell_0} \mid W_{\bar{L}}) \right\} + \{H(W_m) - H(W_m \mid W_L)\} \\
&\leq \ell \underbrace{\{\beta + (\alpha - \beta) - \alpha + \epsilon_1\}}_{=\epsilon_1} + \underbrace{\{\alpha - (\alpha - \beta) + \epsilon_2\}}_{=\beta+\epsilon_2} \\
&= \beta + \ell\epsilon_1 + \epsilon_2
\end{aligned}
$$

Lower bound on $H(\mathcal{T})$

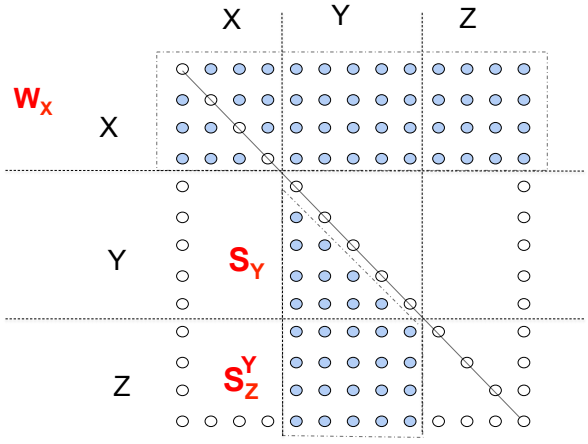

$$
\begin{aligned}
H(\mathcal{T}) &\geq H(\mathcal{T} \mid W_X) \\
&\geq H(W_Y \mid W_X) \\
&= B_{\text{opt-fr}} - \epsilon - H(W_X) \\
&\geq B_{\text{opt-fr}} - \epsilon - (\mu + 1)\alpha
\end{aligned}
$$

Matching the lower and upper bounds leads to an new tradeoff as shown earlier.

The New Outer Bound

The case of $(4, 3, 3)$

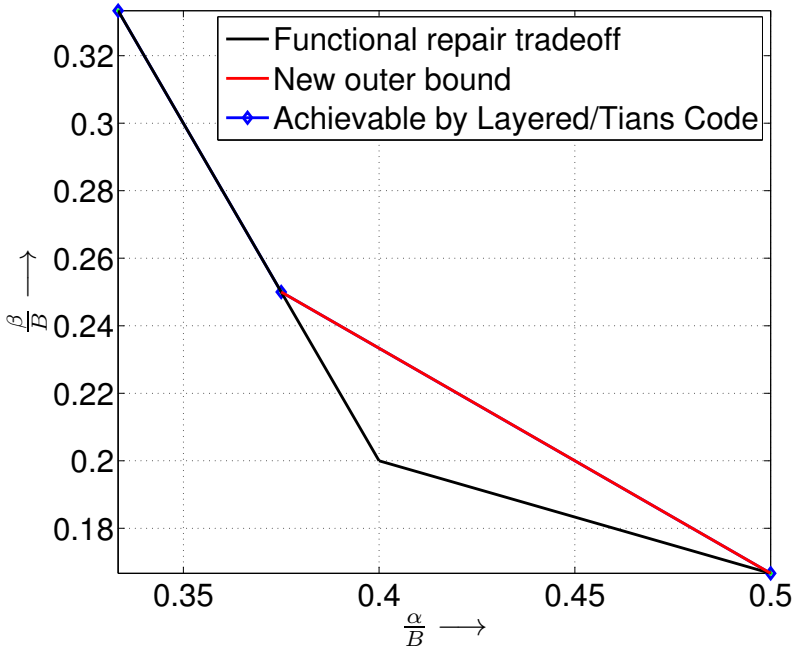

The case of $(5, 4, 4)$

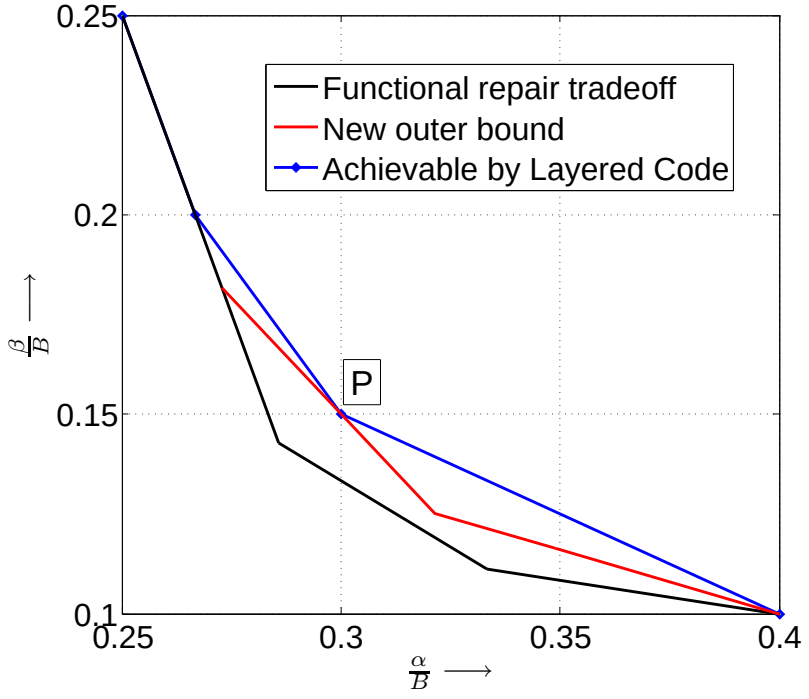

- Provides a new outer bound on ER tradeoff for all $[n, k, d]$
- Bound coincides with the tradeoff characterized by Tian in $[4, 3, 3]$ case.
- The bound in $[5, 4, 4]$ case coincides at one point P with an achievable region by layered codes.
- First instance of an optimal code operating off of the FR tradeoff.

MOHAJER-TANDON OUTER BOUND

AND AN IMPROVEMENT

Exact-Repair File Size Bound

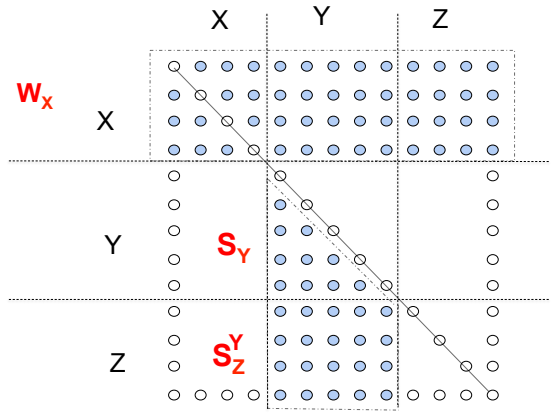

$$
\begin{aligned}
B &= H(W_X, S_Y, S_Z^Y) \\
&= H(W_X) + H(S_Y \mid W_X) + H(S_Z^Y \mid W_X, S_Y) \\
&\leq H(W_X) + H(S_Y) + H(S_Z^Y) \\
&\leq |X|\alpha + |S_Y|\beta + |S_Z^Y|\beta
\end{aligned}
$$

- Earlier, we fixed the size of X ($|Y|$,$|Z|$ determined by $|X|$).
- The bound holds for any value of $|X| = q$, and $|Y| = k - q = p$.

Exact-Repair File Size Bound

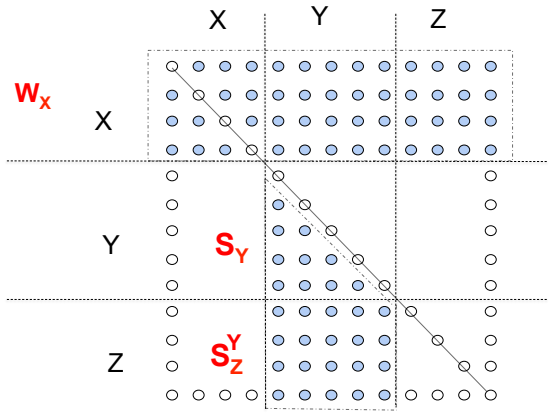

$$
\begin{aligned}
B &= H(W_X, S_Y, S_Z^Y) \\
&= H(W_X) + H(S_Y \mid W_X) + H(S_Z^Y \mid W_X, S_Y) \\
&\leq H(W_X) + H(S_Y) + H(S_Z^Y) \\
&\leq |X|\alpha + |S_Y|\beta + |S_Z^Y|\beta
\end{aligned}
$$

- Earlier, we fixed the size of X ($|Y|$,$|Z|$ determined by $|X|$).
- The bound holds for any value of $|X| = q$, and $|Y| = k - q = p$.

Exact-Repair File Size Bound

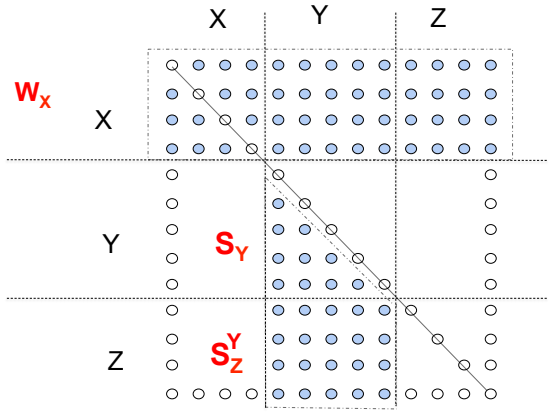

$$
\begin{aligned}
B &= H(W_X, S_Y, S_Z^Y) \\
&= H(W_X) + H(S_Y \mid W_X) + H(S_Z^Y \mid W_X, S_Y) \\
&\leq H(W_X) + H(S_Y) + H(S_Z^Y) \quad \text{(Can we cancel out terms?)} \\
&\leq |X|\alpha + |S_Y|\beta + |S_Z^Y|\beta
\end{aligned}
$$

Yes. By Finding Other Inequalities and Gaussian Elimination

We have:

$$B \leq H(W_X) + \underbrace{\sum_{i=1}^{p} H\left(S_i^{[i-1]} \mid W_X\right)}_{\mathcal{R}(p)} + H(S_Z^Y)$$

We will find two additional inequalities that involve $\mathcal{R}(p)$.

Two Additional Inequalities: Inequality 1

$$
\begin{aligned}
B &= H(W_X) + \sum_{i=1}^{p} H(W_i \mid W_X) - \sum_{i=1}^{p} I(W_i; W_{[i-1]} \mid W_X) \\
&\leq H(W_X) + \sum_{i=1}^{p} H(W_i \mid W_X) - \sum_{i=1}^{p} I(S_i^{[i-1]}; S_{[i-1]}^{i} \mid W_X) \\
&= k\alpha - \underbrace{\sum_{i=1}^{p} H\left(S_i^{[i-1]} \mid W_X\right)}_{\mathcal{R}(p)} - \underbrace{\sum_{i=1}^{p} H\left(S_{[i-1]}^{i} \mid W_X\right)}_{\mathcal{C}(p)} + \underbrace{\sum_{i=1}^{p} H\left(S_{[i-1]}^{i}, S_i^{[i-1]} \mid W_X\right)}_{\mathcal{J}(p)}.
\end{aligned}
$$

Two Additional Inequalities: Inequality 2

$$
\begin{aligned}
B &\leq H(W_X, S^Y_{[d+1]}) \\
&\leq q\alpha + \sum_{i=1}^{p} H(S^i_{[d+1]}|W_X) - \sum_{i=1}^{p} I\left(S^i_{[d+1]}; S^{[i-1]}_{[d+1]} \mid W_Q\right) \\
&\leq q\alpha + \sum_{i=1}^{p} H\left(S^i_{[i-1]} \mid W_X\right) + \sum_{i=1}^{p} H\left(S^i_{[i+1\ d+1]} \mid W_Q\right) - \sum_{i=1}^{p} I\left(S^i_{[d+1]}; S^{[i-1]}_{[d+1]} \mid W_X\right) \\
&\leq q\alpha + \underbrace{\sum_{i=1}^{p} H\left(S^i_{[i-1]} \mid W_Q\right)}_{\mathcal{C}(p)} + \sum_{i=1}^{p} H\left(S^i_{[i+1\ d+1]} \mid W_Q\right) - \underbrace{\sum_{i=1}^{p} H\left(S^i_{[i-1]}, S^{[i-1]}_i \mid W_Q\right)}_{\mathcal{J}(p)}.
\end{aligned}
$$

Adding them leads to:

$$
3B \leq (3k-2p)\alpha + \sum_{i=1}^{p} H\left(S^i_{[i+1\ d+1]} \mid W_X\right) + H\left(S^Y_Z \mid W_X\right).
$$

We further improve upon this bound avoiding union bound on $H\left(S^Y_Z \mid W_X\right)$ by identifying certain symmetries.

The Improved Mohajer-Tandon Bound

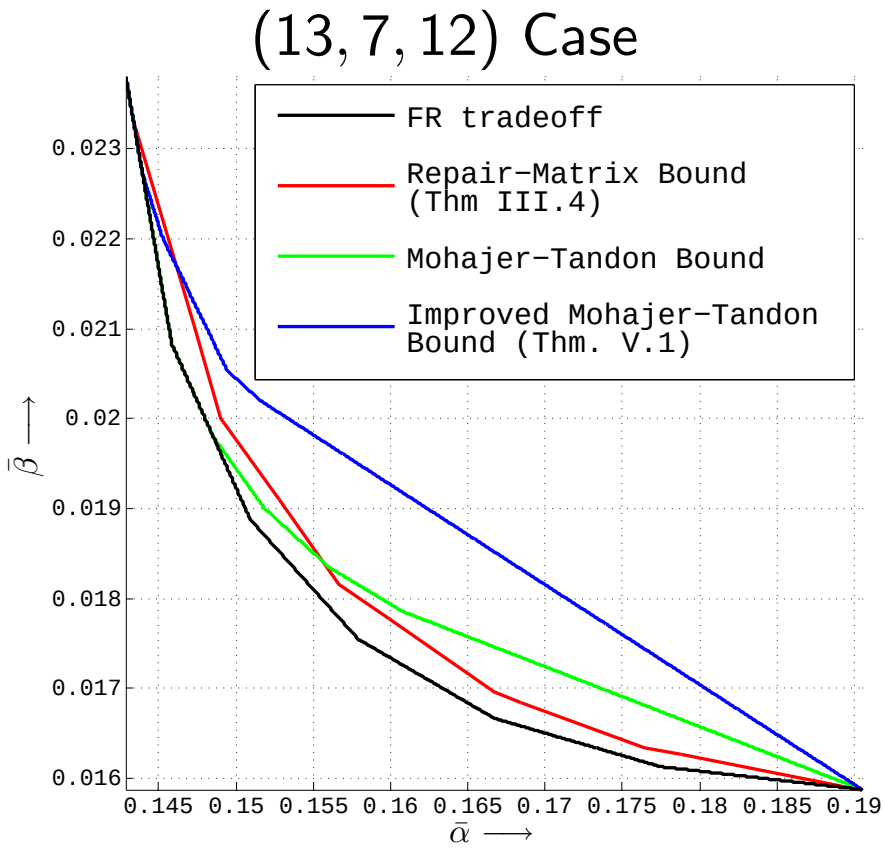

Other Works

- Iwan Duursma, "Outer bounds for exact repair codes," 2014.
- Iwan Duursma, "Shortened regenerating codes," 2015.
- Chao Tian, A Note on the Rate Region of Exact-Repair Regenerating Codes, 2015
- N. Prakash, M. Nikhil Krishnan, "The Storage-Repair-Bandwidth Trade-off of Exact Repair Linear Regenerating Codes for the Case $d = k = (n-1)$", 2015.

(not described here for lack of time)

Thanks!