Rank-Metric Codes with Local Recoverability

Swanand Kadhe

Texas A&M University

Joint work with

Salim El Rouayheb

Iwan Duursma

Alex Sprintson

Allerton '16

Sept 29, 2016

Cloud Storage: Very Large Scale Storage!

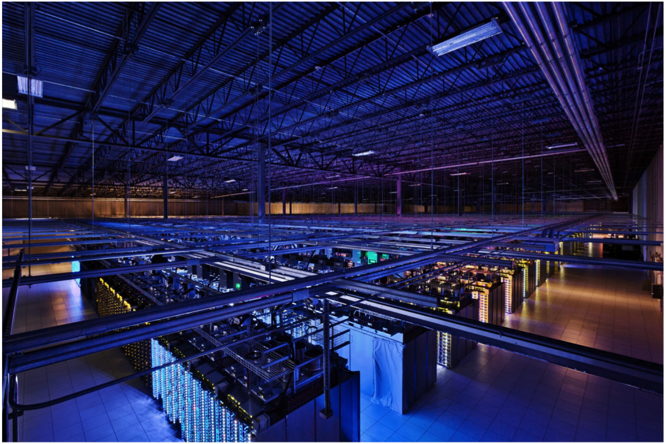

Google data center at Council Bluffs, Iowa

We want cloud systems to be reliable, efficient, and available

Coding for Distributed Storage

Two metrics have received primary research attention

- Repair bandwidth

  Dimakis *et al.* '10,

  Suh-Ramachandran '10,

  Cadambe *et al.* '10,

  Rashmi *et al.* '11,

  Tamo *et al.* '13,

  Ye-Barg '16, ..., ..., ...

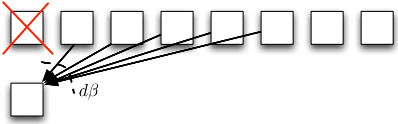

Regenerating Codes

- Locality

  Huang *et al.* 07,

  Oggier-Datta '11,

  Gopalan *et al.* '12,

  Papailiopoulos-Dimakis '14,

  Goparaju-Calderbank '14,

  Tamo-Barg '14, ..., ..., ...

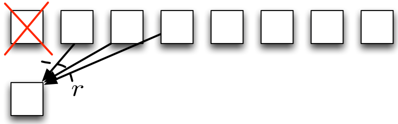

Locally Repairable Codes

Mixed and Correlated Failure Patterns

- Coding has predominantly focused on following type of failures
  - The unit of failure is entire disk
  - Failures occur independently

Mixed and Correlated Failure Patterns

- Coding has predominantly focused on following type of failures
  - The unit of failure is entire disk
  - Failures occur independently
- Storage systems suffer from a large number of mixed and correlated failures
  - Mixed failures: entire drive (node) plus a few blocks fail
  - Correlated failures: a bunch of nodes fail simultaneously

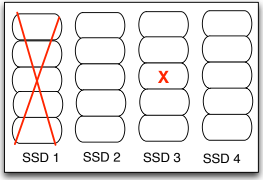

SSD Array

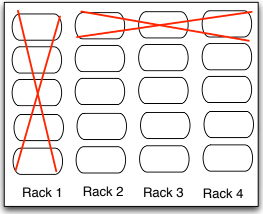

Data center

Example: Mixed failure in a solid state drive (SSD) array, and a correlated failure in a data center

Mixed and Correlated Failure Patterns: Related Work

- Cooperative or centralized regeneration, cooperative local recovery [Shum-Hu '13, Rawat-Mazumdar-Vishwanath '14, Wang-Tamo-Bruck '16]
- Local error correction [Prakash-Kamath-Lalitha-Kumar '12, Song-Dau-Yuen-Li '14]
- Maximally recoverable codes [Gopalan-Huang-Jenkins-Yekhanin '14, Gopalan-Hu-Saraf-Wang-Yekhanin '16]
- Sector-Disk codes, partial MDS codes [Blaum-Hafner-Hetzler '13, Blaum-Plank-Schwartz-Yaakobi '14, Plank-Blaum '14]

We are interested in codes that allow local recoverability from mixed and/or correlated erasures and errors

Crisscross Failure Patterns

- We focus on crisscross failures that form a subclass of mixed and correlated failures
- A crisscross failure pattern affects a limited number of number of rows or columns (or both)

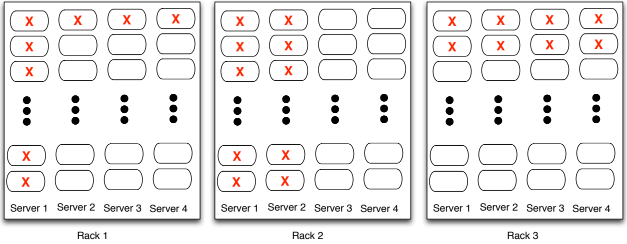

- Codes for crisscross errors (with no locality) have been studied previously [Roth '91, Blaum-Bruck '00]
- We construct codes that allow local recovery from *small weight* crisscross failures. We take a rank-metric approach for code design.

Our Contributions

1. We consider the notion of rank-locality
2. We establish a Singleton-like upper bound on the minimum rank-distance for codes with rank-locality
3. We present an optimal code construction

Rank-Metric Codes

- A rank-metric code $\mathcal{C}$ is a non-empty subset of $\mathbb{F}_q^{m \times n}$ of size q^{mk} endowed with rank-distance metric
  $d_R(A, B) = \mathrm{rank}(A - B)$ [Delsarte '78, Gabidulin '85, Roth '91]

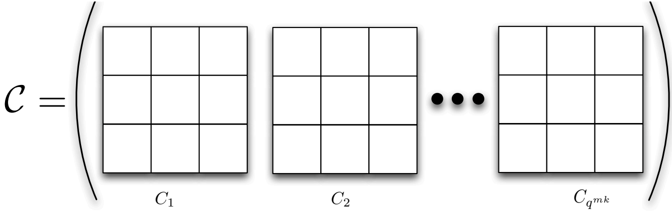

- Maximum rank-distance (MRD) codes are analogues of the maximum distance separable (MDS) codes in the Hamming metric
  - MRD codes achieve the Singleton bound for the rank-metric codes

$$|\mathcal{C}| \leqslant q^{\max\{n,m\}(\min\{n,m\}-d+1)}$$

Gabidulin Codes

Rank-metric analogues of Reed-Solomon codes

- Let $P = \{p_1, \cdots, p_n\}$ be a set of n elements in $\mathbb{F}_{q^m}$ that are linearly independent over $\mathbb{F}_q$ $(m \geqslant n)$
- Let $G_{\mathbf{m}}(x) \in \mathbb{F}_{q^m}[x]$ denote the linearized polynomial of q-degree at most $k-1$ with coefficients $\mathbf{m}$ as follows.

$$G_{\mathbf{m}}(x) = \sum_{j=0}^{k-1} m_j x^{q^j}, \qquad G = \begin{bmatrix} p_1 & p_2 & \cdots & p_n \\ p_1^q & p_2^q & \cdots & p_n^q \\ p_1^{q^2} & p_2^{q^2} & \cdots & p_n^{q^2} \\ \vdots & \vdots & \ddots & \vdots \\ p_1^{q^{k-1}} & p_2^{q^{k-1}} & \cdots & p_n^{q^{k-1}} \end{bmatrix}$$

- Gabidulin code is obtained by the following evaluation map

$$\begin{aligned} \mathrm{Enc} : \mathbb{F}_{q^m}^k &\to \mathbb{F}_{q^m}^n \\ \mathbf{m} &\mapsto \{G_{\mathbf{m}}(p_i), p_i \in P\} \end{aligned}$$

(r, δ)-Locality [Prakash-Lalitha-Kumar '12]

- An (n, k) code $\mathcal{C}$ is said to have (r, δ) locality, if for each symbol c_i, $i \in [n]$, of a codeword $\mathbf{c} = [c_1\ c_2\ \cdots c_n] \in \mathcal{C}$, there exists a set of indices $\Gamma(i)$ such that
  1. $i \in \Gamma(i)$,
  2. $|\Gamma(i)| \leqslant r + \delta - 1$, and
  3. $d_{min}\left(\mathcal{C}|_{\Gamma(i)}\right) \geqslant \delta$,

  where $\mathcal{C}|_{\Gamma(i)}$ is the restriction of $\mathcal{C}$ on the coordinates $\Gamma(i)$
- Any $\delta - 1$ erasures can be repaired from at most r symbols

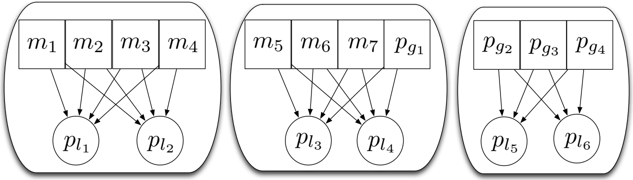

Example: An $(17, 7)$ code with $(4, 3)$-locality containing three local codes

(r, δ)-Locality [Prakash-Lalitha-Kumar '12]

- An (n, k) code $\mathcal{C}$ is said to have (r, δ) locality, if for each symbol c_i, $i \in [n]$, of a codeword $\mathbf{c} = [c_1\ c_2\ \cdots c_n] \in \mathcal{C}$, there exists a set of indices $\Gamma(i)$ such that
  1. $i \in \Gamma(i)$,
  2. $|\Gamma(i)| \leqslant r + \delta - 1$, and
  3. $d_{min}\left(\mathcal{C}|_{\Gamma(i)}\right) \geqslant \delta$,

  where $\mathcal{C}|_{\Gamma(i)}$ is the restriction of $\mathcal{C}$ on the coordinates $\Gamma(i)$
- Any $\delta - 1$ erasures can be repaired from at most r symbols

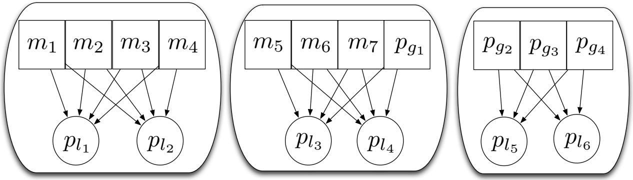

Example: An $(17, 7)$ code with $(4, 3)$-locality containing three local codes

We are interested in locality with respect to rank-metric

(r, δ) Rank-Locality

- An $(m \times n, k)$ rank-metric code $\mathcal{C}$ is said to have (r, δ) rank-locality if for each column $i \in [n]$ of the codeword matrix, there exists a set of columns $\Gamma(i) \subset [n]$ such that
  1. $i \in \Gamma(i)$,
  2. $|\Gamma(i)| \leqslant r + \delta - 1$, and
  3. $d_R\left(\mathcal{C}|_{\Gamma(i)}\right) \geqslant \delta$,

  where $\mathcal{C}|_{\Gamma(i)}$ is the restriction of $\mathcal{C}$ on the columns indexed by $\Gamma(i)$
- The code $\mathcal{C}|_{\Gamma(i)}$ is said to be the local code associated with the i-th column

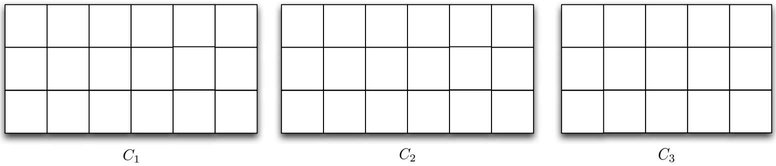

Rank-metric code with $(4, 3)$ rank-locality: local codes C_1, C_2, and C_3 are rank-metric codes with rank-distance at least 2

Rank-Locality: Minimum Distance Bound

Theorem: For a rank-metric code $\mathcal{C} \subseteq \mathbb{F}_q^{m \times n}$ of cardinality q^{mk} with (r, δ) rank-locality, it holds that

$$d_R(\mathcal{C}) \leqslant n - k + 1 - \left(\left\lceil \frac{k}{r} \right\rceil - 1 \right) (\delta - 1).$$

Rank-Locality: Minimum Distance Bound

Theorem: For a rank-metric code $\mathcal{C} \subseteq \mathbb{F}_q^{m \times n}$ of cardinality q^{mk} with (r, δ) rank-locality, it holds that

$$d_R(\mathcal{C}) \leqslant n - k + 1 - \left(\left\lceil \frac{k}{r} \right\rceil - 1 \right)(\delta - 1).$$

Remarks:

- Above Singleton-like bound for the rank-metric coincides with the Singleton-like bound for the Hamming metric by [Prakash *et al.* '13, Rawat *et al.* '14]
- Singleton-optimal code constructions exist for the Hamming metric [Silberstein *et al.* '13, Tamo-Barg '14]

Rank-Locality: Minimum Distance Bound

Theorem: $d_R(\mathcal{C}) \leqslant n - k + 1 - \left(\left\lceil \frac{k}{r} \right\rceil - 1\right)(\delta - 1).$

Proof sketch:

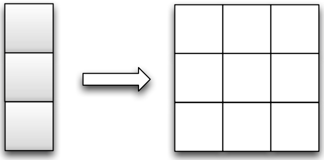

$\mathbf{c} \in \mathbb{F}_{q^m}^n$ $\qquad$ $C \in \mathbb{F}_q^{m \times n}$

By fixing a basis for $\mathbb{F}_{q^m}$, we get a bijection

$\phi : \mathbb{F}_{q^m}^n \to \mathbb{F}_q^{m \times n}$

- Let $C = \phi(\mathbf{c})$. Then, we have $\text{rank}(C) \leqslant \text{weight}(\mathbf{c})$
- An $(m \times n, k, d)$ rank-metric code $\mathcal{C}$ over $\mathbb{F}_q$ can be considered as a block code $\mathcal{C}'$ of length n over $\mathbb{F}_{q^m}$
  - Hence, we have $d_R(\mathcal{C}) \leqslant d_{min}(\mathcal{C}')$
- The result follows from an upper bound on the minimum Hamming distance of an (n, k, d')-LRC

Rank-Locality: Code Construction

We build upon the construction of [Tamo-Barg '14]

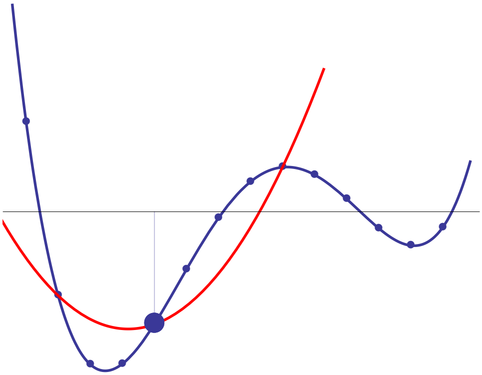

- Intuition: What if we can interpolate low degree polynomials to recover an erased symbol?
- For the rank-locality, we need to use linearized polynomials

Rank-Locality: Code Construction

Assume: $r \mid k$, $(r+\delta-1) \mid n$, $n \mid m$, $\mu := n/(r+\delta-1)$, $q \geqslant 2$

- Encoding Linearized Polynomial:
  - Given k information symbols m_{ij}, $i = 0, \ldots, r-1; j = 0, \ldots, \frac{k}{r}-1$, define the encoding polynomial as

$$G_{\mathbf{m}}(x) = \sum_{i=0}^{r-1} \sum_{j=0}^{\frac{k}{r}-1} m_{ij} x^{q^{(r+\delta-1)j+i}}.$$

Rank-Locality: Code Construction

Assume: $r \mid k$, $(r+\delta-1) \mid n$, $n \mid m$, $\mu := n/(r+\delta-1)$, $q \geqslant 2$

- Encoding Linearized Polynomial:
  - Given k information symbols m_{ij}, $i = 0, \ldots, r-1; j = 0, \ldots, \frac{k}{r} - 1$, define the encoding polynomial as

$$G_{\mathbf{m}}(x) = \sum_{i=0}^{r-1} \sum_{j=0}^{\frac{k}{r}-1} m_{ij} x^{q^{(r+\delta-1)j+i}}.$$

- Evaluation Points:
  - $\{\alpha_1, \ldots, \alpha_{r+\delta-1}\}$: basis of $\mathbb{F}_{q^{r+\delta-1}}$ as a vector space over $\mathbb{F}_q$
  - $\{\beta_1, \ldots, \beta_\mu\}$: basis of $\mathbb{F}_{q^n}$ as a vector space over $\mathbb{F}_{q^{r+\delta-1}}$
  - Evaluation points are $P_1, P_2, \cdots, P_\mu$, where $P_j = \{\alpha_i \beta_j, 1 \leqslant i \leqslant r+\delta-1\}$

Rank-Locality: Code Construction

Assume: $r \mid k$, $(r+\delta-1) \mid n$, $n \mid m$, $\mu := n/(r+\delta-1)$, $q \geqslant 2$

- Encoding Linearized Polynomial:
  - Given k information symbols m_{ij}, $i = 0, \ldots, r-1; j = 0, \ldots, \frac{k}{r}-1$, define the encoding polynomial as

$$G_{\mathbf{m}}(x) = \sum_{i=0}^{r-1} \sum_{j=0}^{\frac{k}{r}-1} m_{ij} x^{q^{(r+\delta-1)j+i}}.$$

- Evaluation Points:
  - $\{\alpha_1, \ldots, \alpha_{r+\delta-1}\}$: basis of $\mathbb{F}_{q^{r+\delta-1}}$ as a vector space over $\mathbb{F}_q$
  - $\{\beta_1, \ldots, \beta_\mu\}$: basis of $\mathbb{F}_{q^n}$ as a vector space over $\mathbb{F}_{q^{r+\delta-1}}$
  - Evaluation points P and their partition $(P_1, P_2, \cdots, P_\mu)$ is given as $P_j = \{\alpha_i \beta_j, 1 \leqslant i \leqslant r+\delta-1\}$
- Codeword is the evaluations of $G_{\mathbf{m}}(x)$ on points in P, *i.e.*, $\mathbf{c} = (G_{\mathbf{m}}(\gamma), \gamma \in P)$

Proposed Construction: Example

$n = 9, k = 4, r = 2, \delta = 2$. Set $q = 2$ and $m = n$

ω: primitive element of $\mathbb{F}_{2^9}$

- Define the encoding polynomial as

$$G_{\mathbf{m}}(x) = m_{00}x^{2^0} + m_{01}x^{2^3} + m_{10}x^{2^1} + m_{11}x^{2^4}.$$

- The evaluation points P are obtained as:
  - $\{1, \omega^{73}, \omega^{146}\}$ as a basis for $\mathbb{F}_{2^3}$ over $\mathbb{F}_2$
  - $\{1, \omega^{309}, \omega^{107}\}$ forms a basis of $\mathbb{F}_{2^9}$ over $\mathbb{F}_{2^3}$

$$P = \{\{1, \omega^{73}, \omega^{146}\}, \{\omega^{309}, \omega^{382}, \omega^{455}\}, \{\omega^{107}, \omega^{180}, \omega^{253}\}\}.$$

- $\mathcal{C}_{\mathrm{Loc}} = \left\{(G_{\mathbf{m}}(\gamma), \gamma \in P) \mid \mathbf{m} \in \mathbb{F}_{2^9}^4\right\}$, and the local codes are $\mathcal{C}_j = \left\{(G_{\mathbf{m}}(\gamma), \gamma \in P_j) \mid \mathbf{m} \in \mathbb{F}_{2^9}^4\right\}$ for $1 \leqslant j \leqslant 3$

Proposed Construction: Example

$n = 9, k = 4, r = 2, \delta = 2$. Set $q = 2$ and $m = n$

ω: primitive element of $\mathbb{F}_{2^9}$

- Define the encoding polynomial as

$$G_{\mathbf{m}}(x) = m_{00}x^{2^0} + m_{01}x^{2^3} + m_{10}x^{2^1} + m_{11}x^{2^4}.$$

- The evaluation points P are obtained as:
  - $\{1, \omega^{73}, \omega^{146}\}$ as a basis for $\mathbb{F}_{2^3}$ over $\mathbb{F}_2$
  - $\{1, \omega^{309}, \omega^{107}\}$ forms a basis of $\mathbb{F}_{2^9}$ over $\mathbb{F}_{2^3}$

$$P = \{\{1, \omega^{73}, \omega^{146}\}, \{\omega^{309}, \omega^{382}, \omega^{455}\}, \{\omega^{107}, \omega^{180}, \omega^{253}\}\}.$$

- $\mathcal{C}_{\mathrm{Loc}} = \left\{ (G_{\mathbf{m}}(\gamma), \gamma \in P) \mid \mathbf{m} \in \mathbb{F}_{2^9}^4 \right\}$, and the local codes are $\mathcal{C}_j = \left\{ (G_{\mathbf{m}}(\gamma), \gamma \in P_j) \mid \mathbf{m} \in \mathbb{F}_{2^9}^4 \right\}$ for $1 \leqslant j \leqslant 3$

- $\mathcal{C}_j$ can be obtained by evaluating the repair polynomial $R_j(x)$ on P_j

$$R_1(x) = (m_{00} + m_{01})x^{2^0} + (m_{10} + m_{11})x^{2^1},$$

$$R_2(x) = (m_{00} + \omega^{119}m_{01})x^{2^0} + (m_{10} + \omega^{238}m_{11})x^{2^1},$$

$$R_3(x) = (m_{00} + \omega^{238}m_{01})x^{2^0} + (m_{10} + \omega^{476}m_{11})x^{2^1}$$

Rank-Distance Optimality of the Proposed Construction

Theorem: The proposed construction is Singleton-optimal, *i.e.*, $d_R\left(\mathcal{C}_{Loc}\right) = n - k + 1 - \left(\left\lceil \frac{k}{r} \right\rceil - 1\right)(\delta - 1)$.

Proof Idea:

The proposed code $\mathcal{C}_{Loc}$ is a subcode of an $\left(n, k + \left(\frac{k}{r} - 1\right)(\delta - 1)\right)$ Gabidulin code

- Example:
  - Recall our example, $n = 9, k = 4, r = 2, \delta = 2$
  - $G_m(x) = m_0 x^{2^0} + m_1 x^{2^1} + m_3 x^{2^3} + m_4 x^{2^4}$
  - This is a subcode of a $(9, 5)$ Gabidulin code, $d_R\left(\mathcal{C}_{Loc}\right) = 5$

Rank-Locality of the Proposed Construction

Theorem: The proposed construction has (r, δ) rank-locality.

Proof Sketch:

- We write the encoding polynomial $G_m(x)$ in terms of a good polynomial $H(x) := x^{q^{r+\delta-1}-1}$ as $G_m(x) = \sum_{i=0}^{r-1} G_i(x) x^{q^i}$, where $G_i(x) = m_{i0} + \sum_{j=1}^{\frac{k}{r}-1} m_{ij} [H(x)]^{\sum_{l=0}^{j-1} q^{(r+\delta-1)l+i}}$.
- Define the repair polynomial for a $\gamma \in P_j$ as

$$R_j(x) = \sum_{i=0}^{r-1} G_i(\gamma) x^{q^i}.$$

- We show that $H(x)$ is constant on P_j, and thus, the evaluations of the encoding polynomial $G_m(x)$ and the repair polynomial $R_j(x)$ on points in P_j are identical

Proposed Construction: Example

$n = 9, k = 4, r = 2, \delta = 2$. Set $q = 2$ and $m = n$

ω: primitive element of $\mathbb{F}_{2^9}$

- Define the encoding polynomial as

$$G_{\mathbf{m}}(x) = m_{00}x^{2^0} + m_{01}x^{2^3} + m_{10}x^{2^1} + m_{11}x^{2^4}.$$

- The evaluation points P are:

$$P = \{\{1, \omega^{73}, \omega^{146}\}, \{\omega^{309}, \omega^{382}, \omega^{455}\}, \{\omega^{107}, \omega^{180}, \omega^{253}\}\}.$$

- $\mathcal{C}_j$ can be obtained by evaluating the repair polynomial $R_j(x)$ on P_j

$$R_1(x) = (m_{00} + m_{01})x^{2^0} + (m_{10} + m_{11})x^{2^1},$$
$$R_2(x) = (m_{00} + \omega^{119}m_{01})x^{2^0} + (m_{10} + \omega^{238}m_{11})x^{2^1},$$
$$R_3(x) = (m_{00} + \omega^{238}m_{01})x^{2^0} + (m_{10} + \omega^{476}m_{11})x^{2^1}$$

Erasure Correction Capability

Proposition: A rank-metric code with (r, δ) rank-locality can locally recover from a crisscross failure that affects at most $\delta - 1$ rows and/or columns.

- Follows from the rank-distance guarantee of a local code

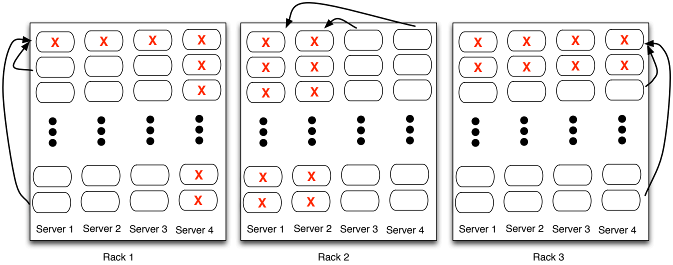

Rank-metric code with $(2, 3)$ rank-locality can locally recover from crisscross erasures affecting any two rows and/or columns

Conclusion and Future Directions

- Rank-locality: Local codes possess good rank-distance.
  We computed tight upper bound on the rank-distance of codes with rank-locality and constructed optimal codes
- Crisscross erasures: Rank-locality ensures local recovery from small weight crisscross failure patterns

Future Directions

- Can we construct rank-metric codes such that every column as well as row is associated with a local code?
- Can we improve the recovery performance by combining rank-metric decoding and Hamming-metric decoding for individual node failures?
- Recovering from a broader class of erasures?