Maximally Recoverable Codes: the Bounded Case

Clay Thomas
Joint work with
Venkata Gandikota, Elena Grigorescu, Minshen Zhu

October 26, 2017

PURDUE
UNIVERSITY

Motivation - Distributed Storage

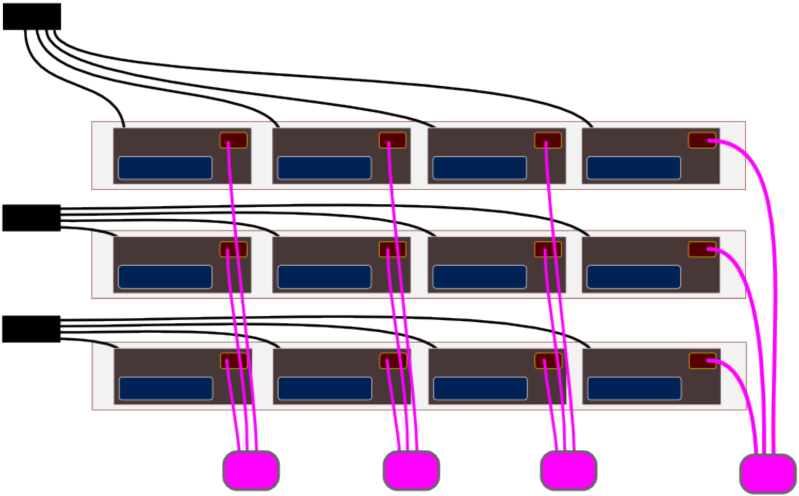

Row failures Column failures Arbitrary failures

Motivation - Distributed Storage

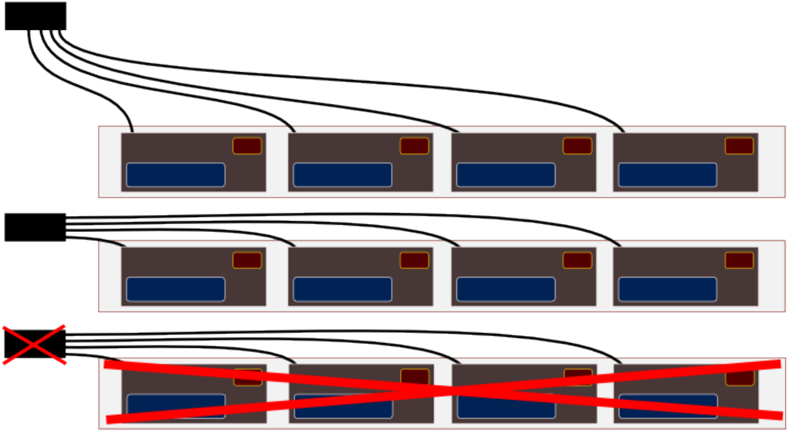

Row failures   Column failures   Arbitrary failures

Motivation - Distributed Storage

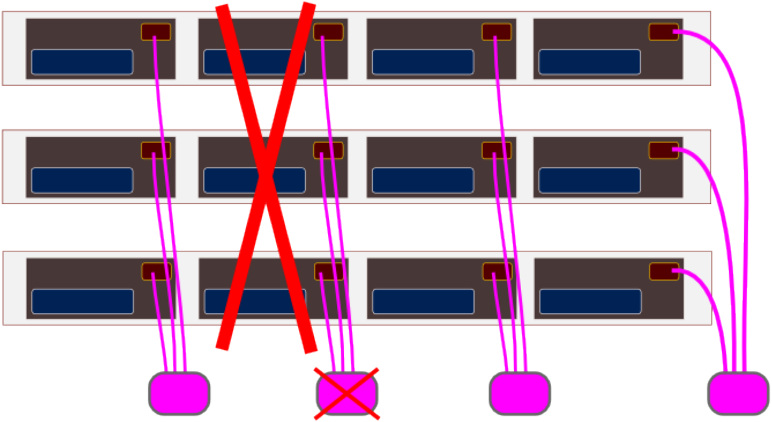

Row failures

Column failures

Arbitrary failures

Motivation - Distributed Storage

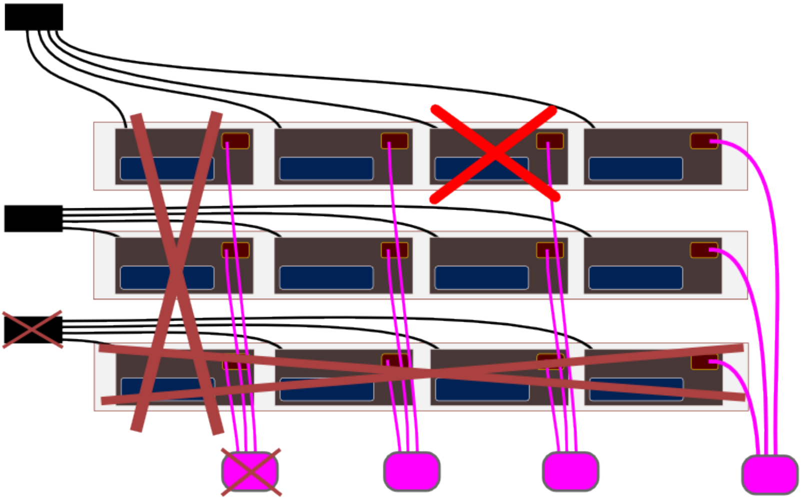

Row failures Column failures Arbitrary failures

Designing a Code

Typical failures are either small OR large but structured (Correlated Erasures)

[CHL07] [Gopalan,Hu,Kopparty,Saraf,Wang,Yekhanin (SODA 17)]

1. Design the topology
    - Known failure patterns
    - Heuristics; hardware
    - Set limits of recoverability
2. Set the coefficients
    - Maximize recoverability
    - Pure math

The Code Topology $T(1,1,1)$ [GHK+17]

Design the topology

Row parities

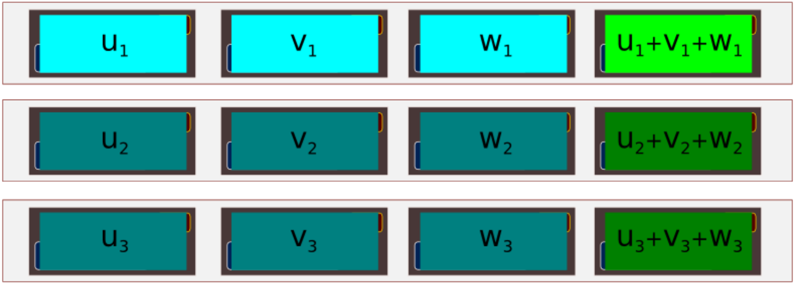

The Code Topology $T(1,1,1)$ [GHK+17]

Design the topology

Column parities

u_1	v_1	w_1	$u_1+v_1+w_1$
u_2	v_2	w_2	$u_2+v_2+w_2$
u_1+u_2	v_1+v_2	w_1+w_2	

The Code Topology $T(1,1,1)$ [GHK+17]

Maximize Recoverability

One extra, global redundancy

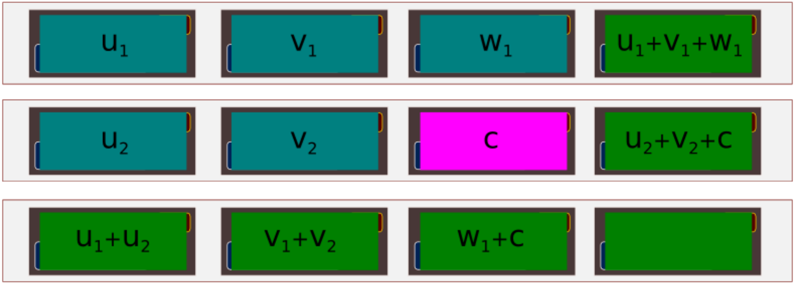

Where c is some linear function of the data symbols

T(a,b,h) [GHK+17]

$T(a, b, h)$:

- a column parities
- b row parities
- h global parities

Facebook: $T(1, 4, 0)$ [SLR+14]

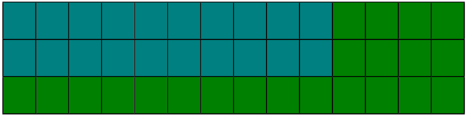

Microsoft: $T(0, 1, 2)$ [HSX+12]

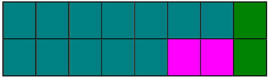

Local Codes: $T(1, 0, h)$ [GHSY12]

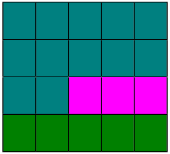

The Code Topology $T(1,1,1)$ [GHK+17]

View codewords as an $m \times n$ grid of symbols $x_{i,j}$

- Rows are codewords in C_{Row}
- Columns are codewords in C_{Col}
- Satisfy global constraint that

$$0 = \sum_{i\in[m],j\in[n]} \gamma_{i,j} x_{i,j}$$

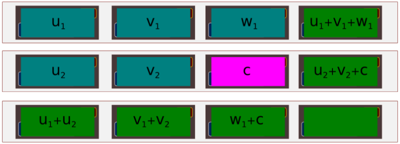

- WLOG C_{Row} and C_{Col} are parity checks
    - entries sum to 0
- A code instantiating $T(1,1,1)$ is defined by the constants $\gamma_{i,j}$.

Decoding Examples

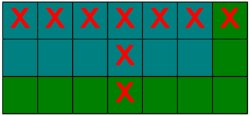

Decoding Examples

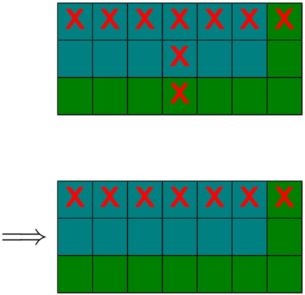

Decoding Examples

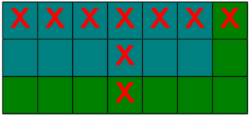

$\Longrightarrow$

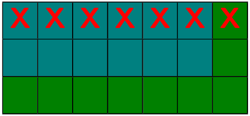

$\Longrightarrow$

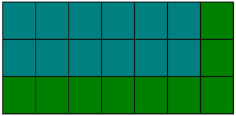

Decoding Examples

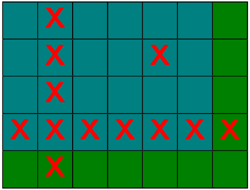

Decoding Examples

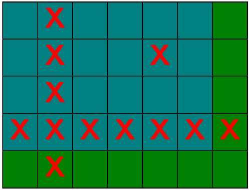

$\Longrightarrow$

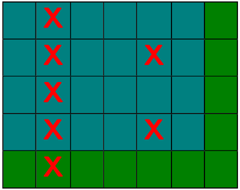

Decoding Examples

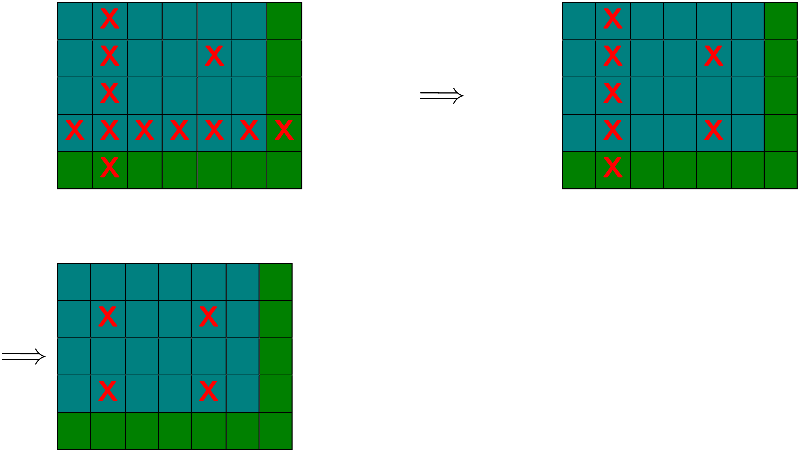

Decoding Examples

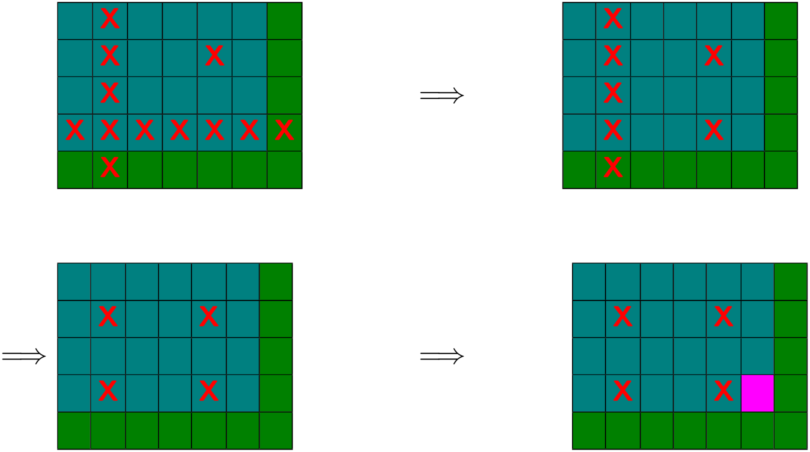

Maximal Recoverability [GHK+17]

- Recall C instantiates $T(1,1,1)$ by setting the coefficients $\gamma_{i,j}$
- As code C instantiating $T(1,1,1)$ corrects an erasure pattern $E \subseteq [n] \times [n]$ if the symbols $x_{i,j}$ for $(i,j) \in E$ can be recovered from those in $[n] \times [n] - E$.
- $E \subseteq [n] \times [n]$ is correctable for $T(1,1,1)$ if there is some code instantiating $T(1,1,1)$ which corrects E
- A code for $T(1,1,1)$ is maximally recoverable (MR) if it can correct every correctable erasure pattern
  - Good news: they exist [CHL07]
  - Bad news: require $\gamma_{i,j} \in \mathbb{F}_2^d$ for d linear in n
    - [Kane, Lovett, Rao (FOCS 17)] [GHK+17]

e-Maximal Recoverability

- An e-MR code for a topology corrects all correctable erasure patterns of size $\leq e$
  - For $T(1,1,1)$ and constant e, attain field size polynomial in n

Reducing to labeling problem [GHK+17]

Subsets of code symbols	$\longleftrightarrow$	Sets of vertices in $K_{n,n}$
Symbol (i,j) erased	$\longleftrightarrow$	Edge $(i,j) \in K_{n,n}$
Parity check weights $\gamma_{i,j}$	$\longleftrightarrow$	Edge weight $\gamma(i,j) \in \mathbb{F}_2^d$
(Irreducible) correctable pattern	$\longleftrightarrow$	Simple cycle with nonzero weight

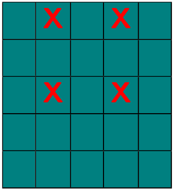

$\longleftrightarrow$

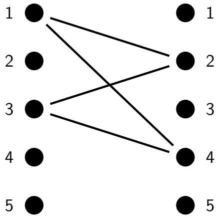

Reducing to labeling problem [GHK$^+$17]

Subsets of code symbols $\longleftrightarrow$ Sets of vertices in $K_{n,n}$

Symbol (i, j) erased $\longleftrightarrow$ Edge $(i, j) \in K_{n,n}$

Parity check weights $\gamma_{i,j}$ $\longleftrightarrow$ Edge weight $\gamma(i, j) \in \mathbb{F}_2^d$

(Irreducible) correctable pattern $\longleftrightarrow$ Simple cycle with nonzero weight

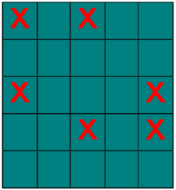

$\longleftrightarrow$

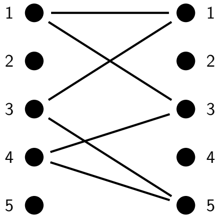

Reducing to labeling problem [GHK+17]

Subsets of code symbols	$\longleftrightarrow$	Sets of vertices in $K_{n,n}$
Symbol (i,j) erased	$\longleftrightarrow$	Edge $(i,j) \in K_{n,n}$
Parity check weights $\gamma_{i,j}$	$\longleftrightarrow$	Edge weight $\gamma(i,j) \in \mathbb{F}_2^d$
(Irreducible) correctable pattern	$\longleftrightarrow$	Simple cycle with nonzero weight

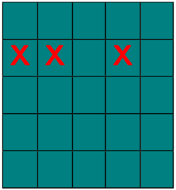

$\longleftrightarrow$

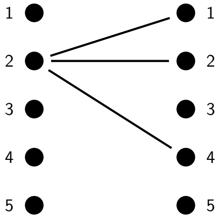

Reducing to labeling problem [GHK+17]

Subsets of code symbols	$\longleftrightarrow$	Sets of vertices in $K_{n,n}$
Symbol (i,j) erased	$\longleftrightarrow$	Edge $(i,j) \in K_{n,n}$
Parity check weights $\gamma_{i,j}$	$\longleftrightarrow$	Edge weight $\gamma(i,j) \in \mathbb{F}_2^d$
(Irreducible) correctable pattern	$\longleftrightarrow$	Simple cycle with nonzero weight

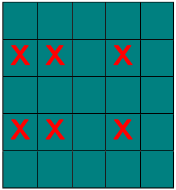

$\longleftrightarrow$

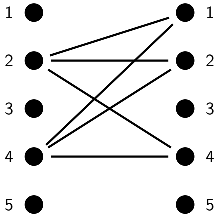

Reducing to labeling problem

Theorem [GHK$^+$17]: A code with $\gamma(i,j) \in \mathbb{F}_2^d$ corrects an (irreducible) error $E \subseteq [n] \times [n]$ iff $E \subseteq K_{n,n}$ is a simply cycle with

$$0 \neq \gamma(E) := \sum_{(i,j)\in E} \gamma(i,j)$$

- Correcting unbounded-length cycles requires $d = \Theta(n)$ [KLR17]
- Our observation: $|\mathbb{F}_2^d|$ polynomial in n if you only correct cycles of bounded length

Our Problem

Given n, e.
Find a labeling $\gamma : [n] \times [n] \to \mathbb{F}_2^d$ such that for all simple cycles E in $K_{n,n}$ of length at most e, we have $\gamma(E) \neq 0$
Goal: minimize d.

- Handle small (constant) number of arbitrary erasures
- Implies existance of e-MR code

Our Results

(Asymptotic) bounds on $|\mathbb{F}_2^d|$ for e-MR codes on $n \times n$ codewords

e	u.b.	l.b.
4	n	n
6	n^2	n^2
8	n^3	n^2
10	n^4	n^3
12	n^5	n^3

Previous results:

- $\leq n^e$ (implied in [GHJY14])
- $\geq \Omega((n/e)^{\log(e/2)})$ for $e \leq \sqrt{n}$ (implied in [GHK$^+$17])

Edge labeling for $e = 4$

Take $\{P_i\}$ and $\{Q_j\}$ each distinct and let

$$\gamma(i,j) = P_i Q_j \in \mathbb{F}_{2^d} \cong \mathbb{F}_2^d$$

Only need n distinct $\{P_i\}, \{Q_j\} \subseteq \mathbb{F}_2^d$ so $|\mathbb{F}_2^d| = O(n)$.

$$\begin{aligned}
\gamma(E) &= \gamma(i_1, j_1) + \gamma(i_1, j_2) \\
&\quad + \gamma(i_2, j_1) + \gamma(i_2, j_2) \\
&= P_{i_1} Q_{j_1} + P_{i_1} Q_{j_2} \\
&\quad + P_{i_2} Q_{j_1} + P_{i_2} Q_{j_2} \\
&= (P_{i_1} + P_{i_2})(Q_{j_1} + Q_{j_2}) \\
&\neq 0
\end{aligned}$$

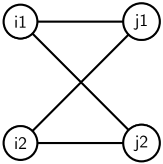

Edge labeling for $e = 8$

Take $\{P_i\}$ and $\{Q_j\}$ each distinct and let

$$\gamma(i,j) = \left(P_i Q_j, P_i^2 Q_j, P_i^4 Q_j\right) \in \mathbb{F}_{2^{\log n}}^3 \cong \mathbb{F}_2^{3 \log n}$$

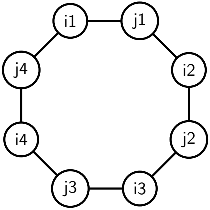

$$\gamma(E) = \begin{pmatrix} a_1 & a_2 & a_3 \\ a_1^2 & a_2^2 & a_3^2 \\ a_1^4 & a_2^4 & a_3^4 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

Where $a_m = P_{i_1} + P_{i_{m+1}}$ and $b_m = Q_{j_m} + Q_{j_{m+1}}$.

Edge labeling for $e = 8$

Take $\{P_i\}$ and $\{Q_j\}$ each distinct and let

$$\gamma(i,j) = \left(P_i Q_j, P_i^2 Q_j, P_i^4 Q_j\right) \in \mathbb{F}_{2^{\log n}}^3 \cong \mathbb{F}_2^{3 \log n}$$

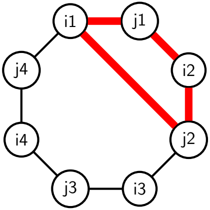

$$\gamma(E) = \begin{pmatrix} a_1 & a_2 & a_3 \\ a_1^2 & a_2^2 & a_3^2 \\ a_1^4 & a_2^4 & a_3^4 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

Where $a_m = P_{i_1} + P_{i_{m+1}}$ and $b_m = Q_{j_m} + Q_{j_{m+1}}$.

Edge labeling for $e = 8$

Take $\{P_i\}$ and $\{Q_j\}$ each distinct and let

$$\gamma(i,j) = \left(P_i Q_j, {\color{red} P_i^2 Q_j}, P_i^4 Q_j\right) \in \mathbb{F}^3_{2^{\log n}} \cong \mathbb{F}_2^{3 \log n}$$

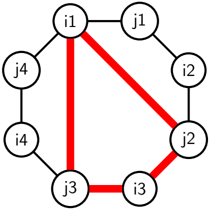

$$\gamma(E) = \begin{pmatrix} a_1 & a_2 & a_3 \\ {\color{red} a_1^2} & {\color{red} a_2^2} & {\color{red} a_3^2} \\ a_1^4 & a_2^4 & a_3^4 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

Where $a_m = P_{i_1} + P_{i_{m+1}}$ and $b_m = Q_{j_m} + Q_{j_{m+1}}$.

Edge labeling for $e = 8$

Take $\{P_i\}$ and $\{Q_j\}$ each distinct and let

$$\gamma(i,j) = \left(P_iQ_j, P_i^2Q_j, P_i^4Q_j\right) \in \mathbb{F}_{2^{\log n}}^3 \cong \mathbb{F}_2^{3\log n}$$

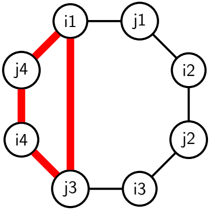

$$\gamma(E) = \begin{pmatrix} a_1 & a_2 & a_3 \\ a_1^2 & a_2^2 & a_3^2 \\ a_1^4 & a_2^4 & a_3^4 \end{pmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix}$$

Where $a_m = P_{i_1} + P_{i_{m+1}}$ and $b_m = Q_{j_m} + Q_{j_{m+1}}$.

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

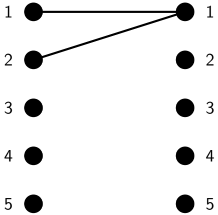

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

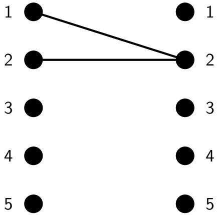

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

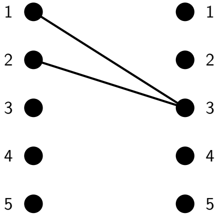

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

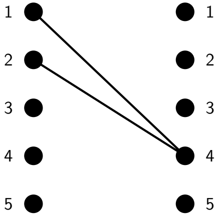

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

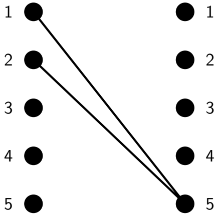

Lower Bound for $e = 4$

Theorem: at least n different labels in $\mathbb{F}_2^d$ are required.

Proof: If 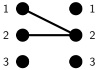 and 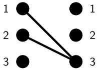 have the same weight, then 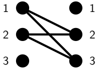 is a simple cycle with zero weight.

Lower Bound for $e \leq 12$

Create a graph.

- Vertices: some collection of paths in $K_{n,n}$
- Edges: $p_1 - p_2$ if p_1 and p_2 together form a simple cycle of length $\leq e$
  - Connected paths must have different γ weight
  - $|\mathbb{F}_2^d|$ is at least the chromatic number
  - Observation: a valid labeling induces a valid coloring

For $e = 4$ we build a clique of size n

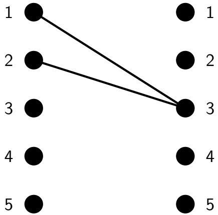

Lower Bound for $e \leq 12$

Create a graph.

- Vertices: some collection of paths in $K_{n,n}$
- Edges: $p_1 - p_2$ if p_1 and p_2 together form a simple cycle of length $\leq e$
  - Connected paths must have different γ weight
  - $|\mathbb{F}_2^d|$ is at least the chromatic number
  - Observation: a valid labeling induces a valid coloring

For $e = 6$ we build a clique of size $(n-1)^2$

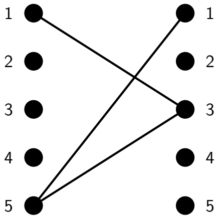

Lower Bound for $e \leq 12$

Create a graph.

- Vertices: some collection of paths in $K_{n,n}$
- Edges: $p_1 - p_2$ if p_1 and p_2 together form a simple cycle of length $\leq e$
  - Connected paths must have different γ weight
  - $|\mathbb{F}_2^d|$ is at least the chromatic number
  - Observation: a valid labeling induces a valid coloring

For $e = 10$ we get chromatic number $\Omega(n^3)$

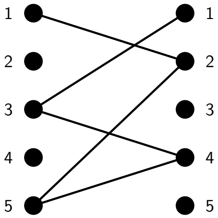

Conclusions and Open Questions

We showed

- $|\mathbb{F}_2^d| \leq n^{e/2-1}$ for $e \leq 12$
  - Tight for $e = 4, 6$
- $|\mathbb{F}_2^d| \geq n^3$ for $e = 10$

Open questions

- Tight bounds for all e
- Construct e-MR codes for other topologies