Transmission and Scheduling Aspects of Distributed Storage and Their Connections with Index Coding

Parastoo Sadeghi

Research School of Engineering, The Australian National University

Background and motivations

- In practical storage systems, rarely are all messages available at a single source.
- They are distributed at different sources across the network.
- **Majority** of recent works on distributed storage focus on **repair performance**.
- We do not have a good understanding of how index coding works in distributed storage systems.

Centralized index coding example and notation

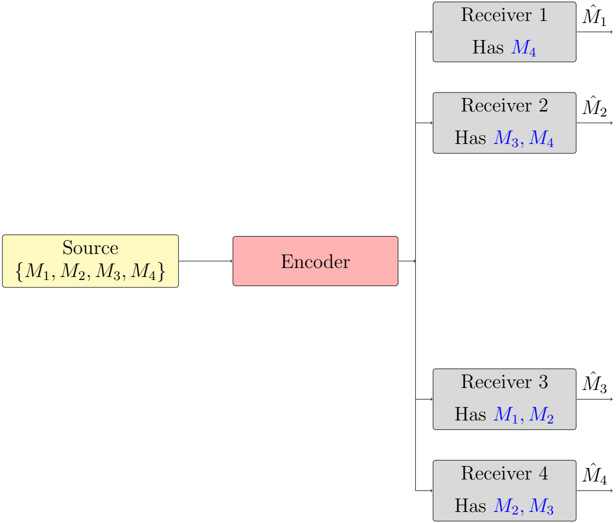

Notation:

Demands	Has
1	4
2	3, 4
3	1, 2
4	2, 3

Distributed example

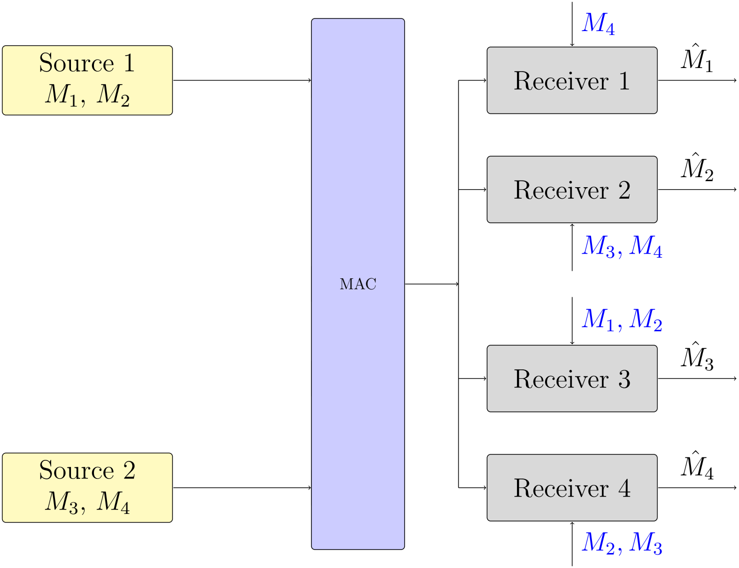

Contributions

- This work looks at **distributed index coding**.
- And studies the impact of message distribution across the network on index coding **achievable rates (and not repair properties)**.

Road map

- Brief review of existing work.
- **Base** example of distributed index coding and establishing an achievable rate region.
- **Different** message distribution.
- Some messages **repeated**.
- **Optimal** message distribution?
- (only if time allows) A simple **MDS code**.

Centralized index coding solution

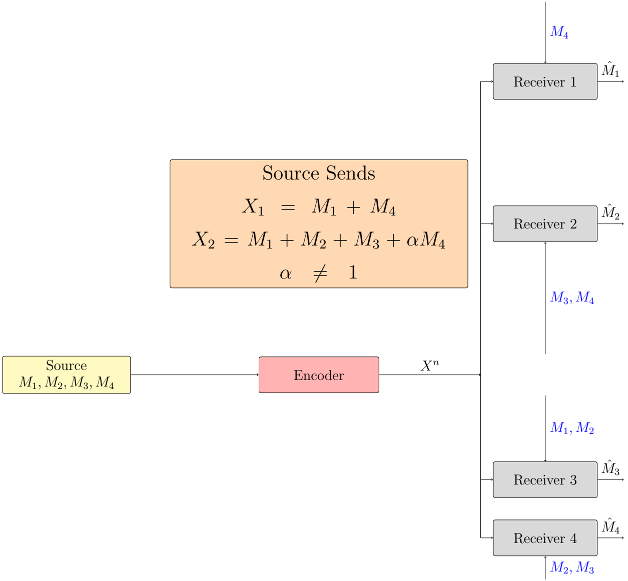

Notation:
Demands | Has

Demands	Has
1	4
2	3, 4
3	1, 2
4	2, 3

Source Sends

$$X_1 = M_1 + M_4$$

$$X_2 = M_1 + M_2 + M_3 + \alpha M_4$$

$$\alpha \neq 1$$

Receivers Decode

$$M_1 = X_1 - M_4$$

$$M_2 = X_2 - \underbrace{(X_1 - M_4)}_{M_1} - M_3 - \alpha M_4$$

$$\vdots$$

Index coding rate region

- It has been shown in [1][1] that the following rate region is achievable for this example:

$$
\begin{aligned}
R_1 + R_2 &\leq 1, \\
R_1 + R_3 &\leq 1, \\
R_1 + R_4 &\leq 1, \\
R_3 + R_4 &\leq 1.
\end{aligned}
$$

- Resulting in sum rate $R_1 + R_2 + R_3 + R_4 \leq 2$.

- The method uses the concept of **virtual composite message encoders**.

[1] [1] F. Arbabjolfaei, B. Bandemer, Y.-H. Kim, E. Sasoglu, and L. Wang, "On the capacity region for index coding," in IEEE Int. Symp. on Information Theory (ISIT), July 2013, pp. 962–966.

Virtual composite message encoder

Original Messages:

$$M_1 \in [1, 2, 3, \cdots, 2^{nR_1}]$$
$$M_4 \in [1, 2, 3, \cdots, 2^{nR_4}]$$

$\xrightarrow{(M_1, M_4)}$

Composite Message:

$$W_{1,4} \in [1, 2, 3, \cdots, 2^{nS_{1,4}}]$$

Example: Arbitrary Mapping

$$(M_1, M_4) = (1, 1) \rightarrow W_{1,4} = 3$$
$$(M_1, M_4) = (1, 2) \rightarrow W_{1,4} = 2$$
$$\vdots$$
$$(M_1, M_4) = (x, y) \rightarrow W_{1,4} = 2^{nS_{1,4}}$$
$$\vdots$$
$$(M_1, M_4) = (2^{nR_1}, 2^{nR_4}) \rightarrow W_{1,4} = 1$$

Example: Linear Mapping

$$(M_1, M_4) \rightarrow W_{1,4} = M_1 + M_4$$

Virtual composite message encoder

Original Messages:

$$M_1 \in [1, 2, 3, \cdots, 2^{nR_1}]$$
$$M_2 \in [1, 2, 3, \cdots, 2^{nR_2}]$$
$$M_3 \in [1, 2, 3, \cdots, 2^{nR_3}]$$
$$M_4 \in [1, 2, 3, \cdots, 2^{nR_4}]$$

$$(M_1, M_2, M_3, M_4) \longrightarrow$$

Composite Message:

$$W_{1,2,3,4} \in [1, 2, 3, \cdots, 2^{nS_{1,2,3,4}}]$$

Example: Linear Mapping

$$(M_1, M_2, M_3, M_4) \rightarrow W_{1,2,3,4} = M_1 + M_2 + M_3 + \alpha M_4$$

where

$$M_1, \cdots, W_{1,2,3,4}, \alpha \in F_q, \quad q = 2^m$$

System block diagram

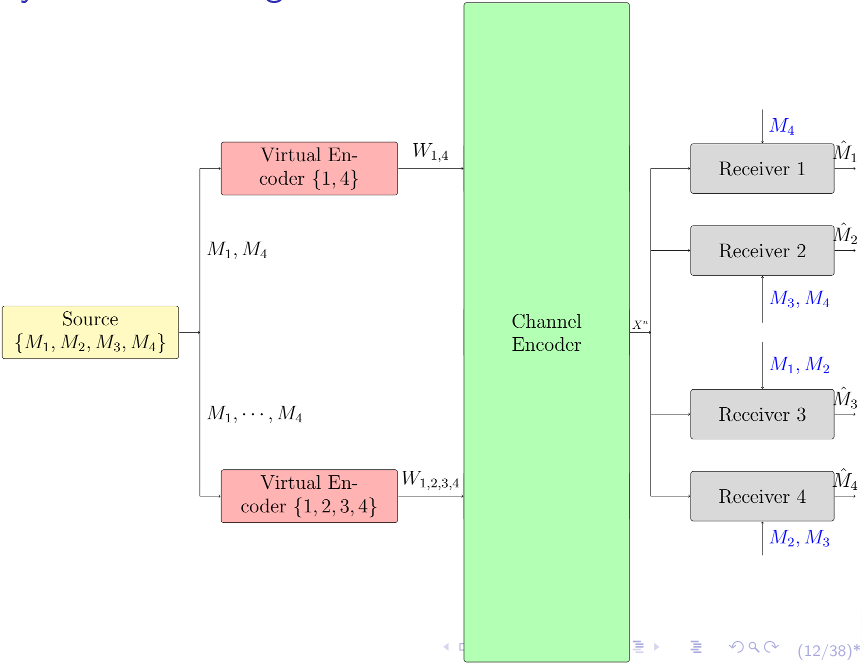

Existing theory

- Let $\mathcal{K}_j$ be index of messages receiver j decodes ($j \in \mathcal{K}_j$) and $\mathcal{A}_j$ be its side information. Then the index coding rates for j $\mathcal{R}(\mathcal{K}_j|\mathcal{A}_j)$ follow

$$\sum_{j \in \mathcal{J}} R_j < \sum_{\mathcal{J}' \subseteq \mathcal{K}_j \cup \mathcal{A}_j : \mathcal{J}' \cap \mathcal{J} \neq \emptyset} S_{\mathcal{J}'}$$

  for all $\mathcal{J} \subseteq \mathcal{K}_j \setminus \mathcal{A}_j$. Any composite message in $\mathcal{K}_j \cup \mathcal{A}_j$ common with $\mathcal{K}_j \setminus \mathcal{A}_j$ is relevant.

- Achievable rate region is given by

$$(R_1, R_2, \cdots, R_N) \in \bigcap_{j \in [1:N]} \bigcup_{\mathcal{K}_j \subseteq [1:N] : j \in \mathcal{K}_j} \mathcal{R}(\mathcal{K}_j|\mathcal{A}_j)$$

Existing theory - 2

- Constraints on composite message rates $S_{\mathcal{J}}$ come from the unit channel capacity (but are somewhat relaxed by receivers' side information):

$$\sum_{\mathcal{J}:\mathcal{J}\subsetneq\mathcal{A}_j} S_{\mathcal{J}} \leq 1$$

for all $j \in [1:N]$.

- Any composite message that is fully embedded in $\mathcal{A}_j$ does not constrain the composite rates.

Example

$$\mathcal{K}_1 = \{1\}, \mathcal{A}_1 = \{4\} \rightarrow \quad R_1 \leq S_{1,4},$$

$$\mathcal{K}_2 = \{1,2\}, \mathcal{A}_2 = \{3,4\} \rightarrow \quad \begin{cases} R_2 \leq S_{1,2,3,4} \\ R_1 + R_2 \leq S_{1,4} + S_{1,2,3,4}, \end{cases}$$

$$\mathcal{K}_3 = \{3,4\}, \mathcal{A}_3 = \{1,2\} \rightarrow \quad \begin{cases} R_3 \leq S_{1,2,3,4} \\ R_3 + R_4 \leq S_{1,4} + S_{1,2,3,4} \end{cases}$$

$$\mathcal{K}_4 = \{1,4\}, \mathcal{A}_4 = \{2,3\} \rightarrow \quad R_1 + R_4 \leq S_{1,4} + S_{1,2,3,4},$$

$$S_{1,4} + S_{1,2,3,4} \leq 1$$

Extension to distributed index coding - base example

- The key difference is that only a subset of composite messages may be computable in the network that are available at distributed sources.

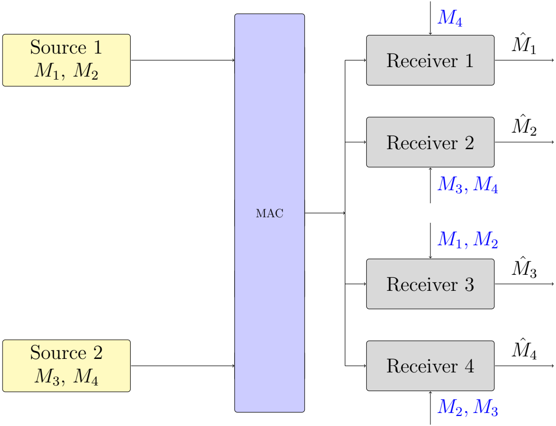

Characterizing the rate region for this example

- The set of computable composite indices is

$$\mathcal{P}' = \{\{1\},\{2\},\{1,2\},\{3\},\{4\},\{3,4\}\}.$$

$$\mathcal{K}_1 = \{1\}, \mathcal{A}_1 = \{4\} \rightarrow \quad R_1 \leq S_1,$$

$$\mathcal{K}_2 = \{1,2\}, \mathcal{A}_2 = \{3,4\} \rightarrow \quad \begin{cases} R_2 & \leq S_{1,2} \\ R_1 + R_2 & \leq S_1 + S_{1,2} \end{cases}$$

$$\mathcal{K}_3 = \{3\}, \mathcal{A}_3 = \{1,2\} \rightarrow \quad R_3 \leq S_3$$

$$\mathcal{K}_2 = \{1,4\}, \mathcal{A}_2 = \{2,3\} \rightarrow \quad \begin{cases} R_4 & \leq S_4 \\ R_1 + R_4 & \leq S_1 + S_{1,2} + S_4 \end{cases}$$

Characterizing the rate region -2

- We consider a binary erasure MAC without noise

$$Y = X_1 + X_2$$

- Effective MAC constraints on composite rates:

$$\begin{aligned} S_1 + S_{1,2} &\leq 1 \\ S_1 + S_{1,2} + S_3 &\leq 1.5 \\ S_3 + S_4 &\leq 1 \\ S_1 + S_{1,2} + S_4 &\leq 1.5 \end{aligned}$$

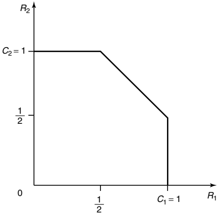

Achievable rate region

- Rate region is specified by

$$R_1 + R_2 \leq 1, \qquad R_1 + R_2 + R_3 \leq 1.5,$$
$$R_3 + R_4 \leq 1, \qquad R_1 + R_2 + R_4 \leq 1.5.$$

- The same sum rate of

$$R_1 + R_2 + R_3 + R_4 \leq 2$$

with

$$R_1 = R_2 = R_3 = R_4 = 0.5$$

is achievable as shown in the next slide.

- Despite distributed storage constraints, MAC transmissions helped to create **key "channel" composite messages**.

Achievable Scheme

$$Y_1 = M_1 + M_4$$

$$Y_2 = (M_1 \oplus M_2) + M_3$$

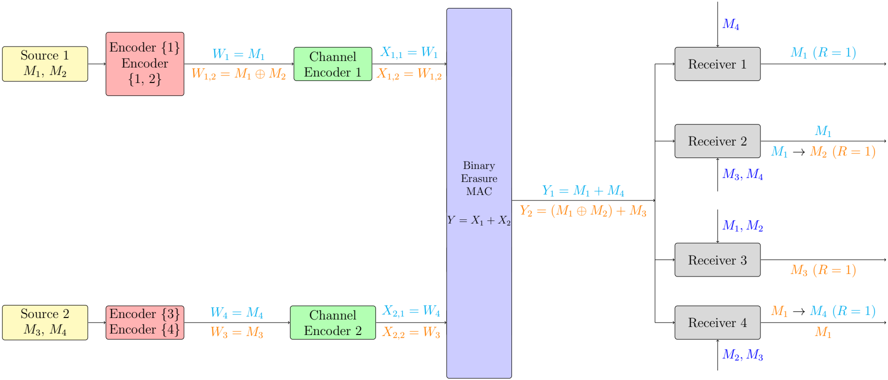

Developed theory

- Let $\mathcal{K}_j$ be index of messages receiver j decodes ($j \in \mathcal{K}_j$) and $\mathcal{A}_j$ be its side information. Then the index coding rates for j $\mathcal{R}(\mathcal{K}_j|\mathcal{A}_j)$ follow

$$\sum_{j \in \mathcal{J}} R_j < \sum_{\mathcal{J}' \in (\mathcal{P}(\mathcal{K}_j \cup \mathcal{A}_j) \cap \mathcal{P}') : \mathcal{J}' \cap \mathcal{J} \neq \emptyset} S_{\mathcal{J}}$$

  for all $\mathcal{J} \subseteq \mathcal{K}_j \setminus \mathcal{A}_j$.

  Composite messages in the power set of $\mathcal{K}_j \cup \mathcal{A}_j$ that are computable in the network (belong to $\mathcal{P}'$) are relevant.

Developed theory 2

- As before achievable rate region is given by

$$(R_1, R_2, \cdots, R_N) \in \bigcap_{j \in [1:N]} \bigcup_{\mathcal{K}_j \subseteq [1:N] : j \in \mathcal{K}_j} \mathcal{R}(\mathcal{K}_j | \mathcal{A}_j)$$

Developed theory - 3

- Constraints on composite message rates $S_{\mathcal{J}}$ come from the MAC capacity (but are somewhat relaxed by receivers' side information):
- The rate of every selected composite message that is overlapping with $\mathcal{K}_j$ and not fully embedded in $\mathcal{A}_j$ must belong to MAC capacity region. More mathematically:
- Find a suitable subset of composite messages computable in the network $\mathcal{J}^* \subseteq \mathcal{P}'$
  - such that for all $j \in [1:N]$ and for all $\tilde{\mathcal{J}} \subseteq \mathcal{J}^* : \exists \mathcal{J} \in \tilde{\mathcal{J}} : \mathcal{K}_j \cap \tilde{\mathcal{J}} \neq \emptyset$ we have

$$\sum_{\mathcal{J} \in \tilde{\mathcal{J}} : \mathcal{J} \subsetneq \mathcal{A}_j} S_{\mathcal{J}}$$

belong to the MAC capacity region $\mathcal{M}$.

Different message distribution.

- How does message distribution affect performance?

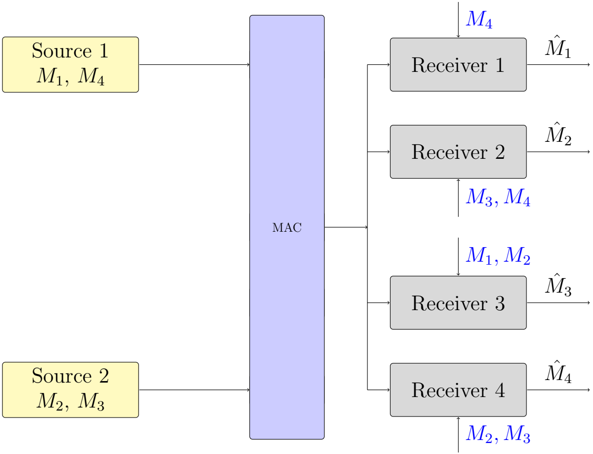

Achievable rate region

- Rate region is specified by more relaxed conditions

$$
\begin{aligned}
R_1 + R_2 + R_3 &\le 1.5, \\
R_1 + R_2 + R_4 &\le 1.5 \\
R_1 + R_4 &\le 1 \\
R_2 &\le 1 \\
R_3 &\le 1
\end{aligned}
$$

- 25% higher same sum rate of

$$R_1 + R_2 + R_3 + R_4 \le 2.5$$

with

$$R_1 = R_4 = 0.25$$

and

$$R_2 = R_3 = 1$$

is achievable as shown next.

Achievable Scheme

[2]

$$Y_1 = LT(M_1) + (M_2 \oplus M_3)$$

$$Y_2 = LT(M_4) + (M_2 \oplus M_3)$$

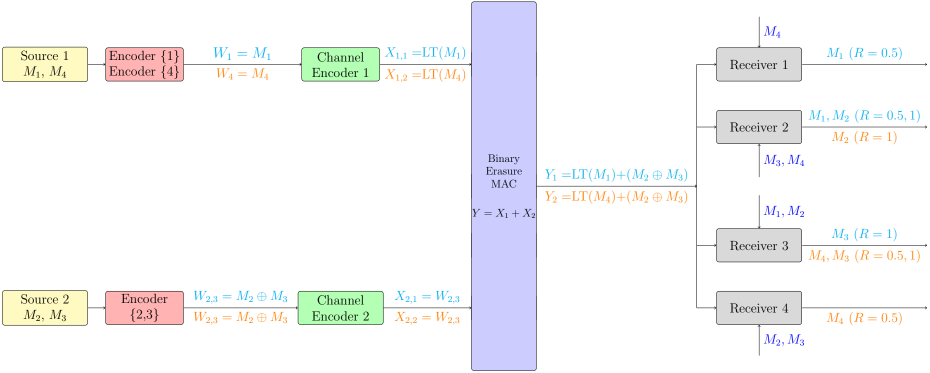

[2] Or any suitable block erasure code.

M_1 repeated.

- How does message repetition across the network affect performance?
- Sources can cooperate for transmission of M_1 to achieve higher rates.

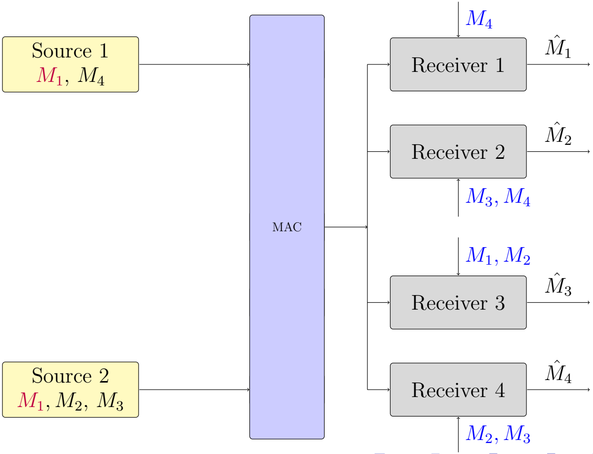

Achievable rate region

Fix $P(x_1, x_2) = \frac{1}{4}$

$$R_2 \leq 1, \qquad R_3 \leq 1, \qquad R_4 \leq 1,$$

$$R_1 + R_2 \leq 1.5, \quad R_1 + R_3 \leq 1.5,$$

$$R_3 + R_4 \leq 1.5, \quad R_1 + R_4 \leq 1.5$$

50% higher same sum rate of

$$R_1 + R_2 + R_3 + R_4 \leq 3$$

with

$$R_1 = R_3 = 0.5$$

and

$$R_2 = R_4 = 1$$

is achievable as shown next.

Achievable scheme using non-unique decoding

$$Y = (M_1 \oplus M_4) + (M_1 \oplus M_2 \oplus M_3)$$

$$Y = 0/2 \rightarrow (M_1 \oplus M_4) = (M_1 \oplus M_2 \oplus M_3) = 0/1$$

$$Y = 1 \rightarrow X_1 \oplus X_2 = 1 \rightarrow (M_1 \oplus M_4) \oplus (M_1 \oplus M_2 \oplus M_3) = M_2 \oplus M_3 \oplus M_4 = 1$$

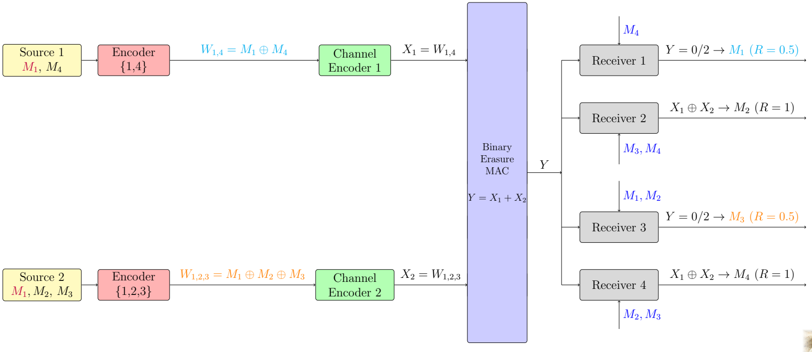

All messages **repeated**.

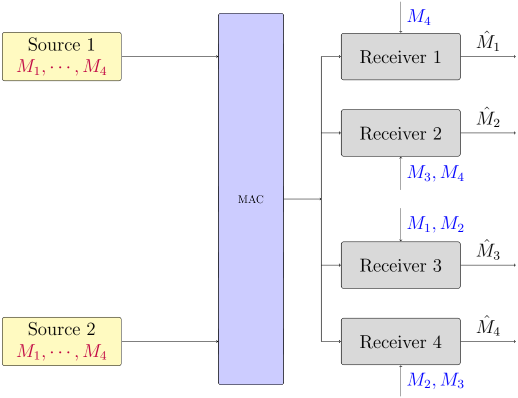

$$R_1 + R_2 + R_3 + R_4 \leq 2 \times \log_2 3$$

is achievable, which is only marginally better than previous case which needed only 62.5% of storage.

Optimal (min storage - max rate) solution?

- How can we optimally use distributed storage and MAC capacity?
- All sources should be able to compute all composite messages
- **Stripe** each message in two parts and store each part on one source

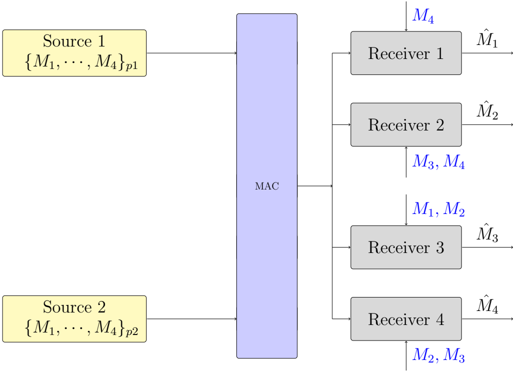

Achievable rate region-symmetric case

- Each half can achieve the rate as if it was **centralized** index coding:

$$R_{1_{pk}} + R_{2_{pk}} \leq 1, \quad R_{1_{pk}} + R_{3_{pk}} \leq 1$$
$$R_{1_{pk}} + R_{4_{pk}} \leq 1, \quad R_{3_{pk}} + R_{4_{pk}} \leq 1$$

for $k = 1, 2$.

- Moreover, due to MAC constraints, we can symmetrically achieve

$$R_1 + R_2 \leq 1.5, \quad R_1 + R_3 \leq 1.5$$
$$R_1 + R_4 \leq 1.5, \quad R_3 + R_4 \leq 1.5$$

As shown next, $R_1 = R_2 = R_3 = R_4 = 0.75$ is achievable.

Achievable scheme

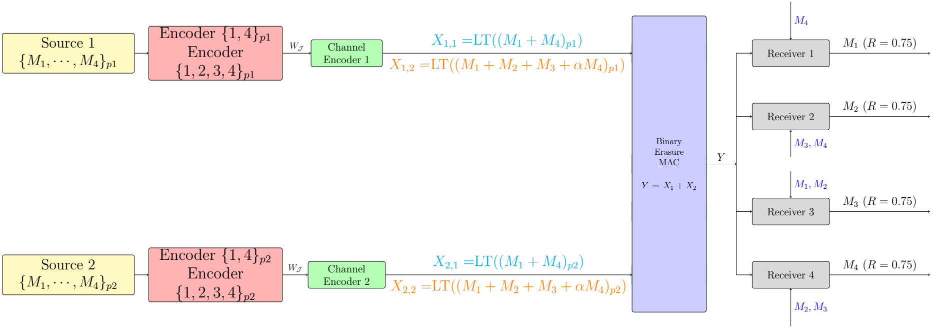

Take-home messages

- The distribution of messages across the network can greatly affect index coding solutions and rates.
- **Striping** seems to be the optimal thing to do in **symmetric** networks, but the effect of **heterogeneous** conditions is unknown.
- Research is needed to better understand the interactions between **storage, repair bandwidth, data availability, and index coding transmission rates**.
- Research is needed to develop **practical scheduling and high rate transmission schemes** for distributed index coding.

Simple (3,2) MDS code

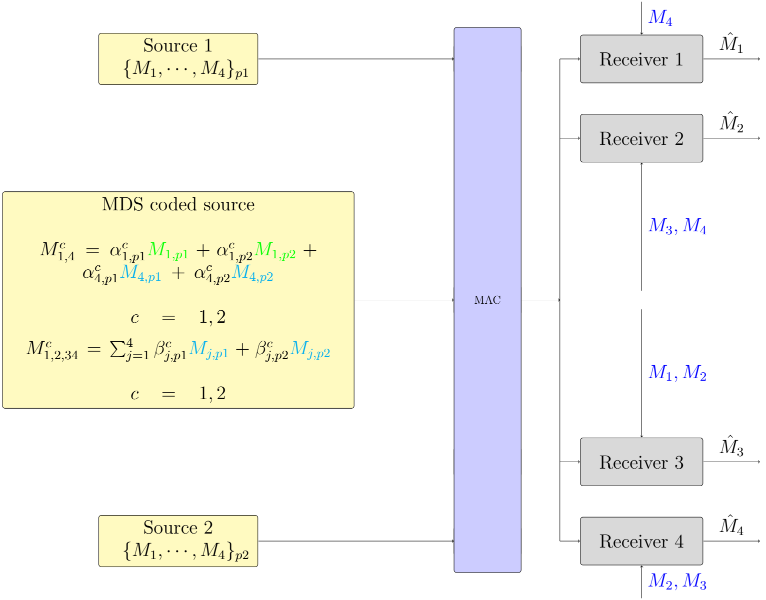

Achievable rate region

$$R_1 + R_2 \leq 1.81, \quad R_1 + R_3 \leq 1.81$$
$$R_1 + R_4 \leq 1.81, \quad R_3 + R_4 \leq 1.81$$

- Symmetric rates

$$R_1 = R_2 = R_3 = R_4 = 0.905$$

are achievable.