Efficient Universal Recovery in Broadcast Networks

Thomas Courtade and Rick Wesel

UCLA

September 30, 2010

System Model and Problem Statement

Universal recovery in a broadcast network

- Consider a network of N nodes.

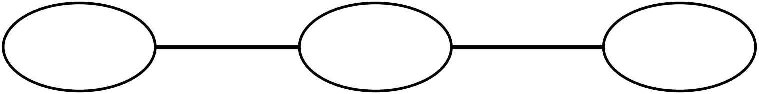

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$.

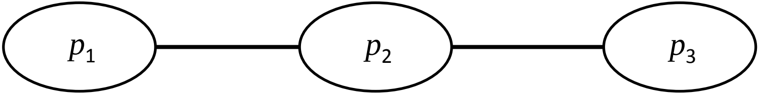

System Model and Problem Statement

Universal recovery in a broadcast network

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$.
- Each node broadcasts packets to its neighbors. One transmission consists of one packet.

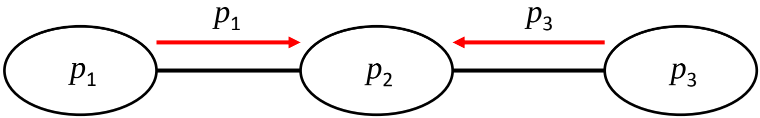

System Model and Problem Statement

Universal recovery in a broadcast network

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$.
- Each node broadcasts packets to its neighbors. One transmission consists of one packet.
- Links are noiseless and interference-free. Each node knows the network topology and which packets each node already knows.

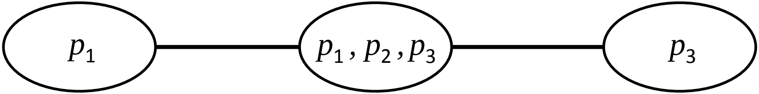

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$.
- Each node broadcasts packets to its neighbors. One transmission consists of one packet.
- Links are noiseless and interference-free. Each node knows the network topology and which packets each node already knows.

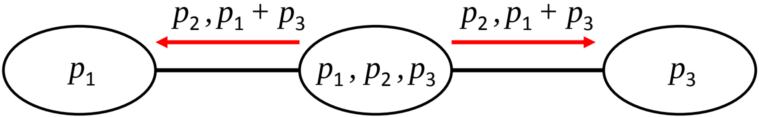

System Model and Problem Statement

Universal recovery in a broadcast network

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \ldots, p_k\}$.
- Each node broadcasts packets to its neighbors. One transmission consists of one packet.
- Links are noiseless and interference-free. Each node knows the network topology and which packets each node already knows.
- "Universal Recovery" is achieved if some transmission scheme permits each node to recover all k packets.

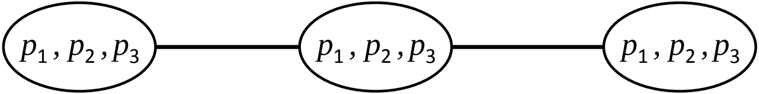

System Model and Problem Statement

Universal recovery in a broadcast network

- Consider a network of N nodes.
- Node i begins with a set of packets $\mathcal{P}_i \subseteq \{p_1, \dots, p_k\}$.
- Each node broadcasts packets to its neighbors. One transmission consists of one packet.
- Links are noiseless and interference-free. Each node knows the network topology and which packets each node already knows.
- "Universal Recovery" is achieved if some transmission scheme permits each node to recover all k packets.

Question

What is the minimum number of transmissions required for universal recovery?

Transmission Strategies

The scheduling of transmissions

Definition (Transmission Strategy)

Let b_i^j be the number of transmissions by node i during round j. The set of values $\{b_i^j\}_{i,j}$ is called a transmission strategy.

Transmission Strategies

The scheduling of transmissions

Definition (Transmission Strategy)

Let b_i^j be the number of transmissions by node i during round j. The set of values $\{b_i^j\}_{i,j}$ is called a transmission strategy.

- For the previous example, $b_1^1 = 1, b_3^1 = 1$, and $b_2^2 = 2$ is a transmission strategy that permits universal recovery.

Transmission Strategies

The scheduling of transmissions

Definition (Transmission Strategy)

Let b_i^j be the number of transmissions by node i during round j. The set of values $\{b_i^j\}_{i,j}$ is called a transmission strategy.

- For the previous example, $b_1^1 = 1, b_3^1 = 1$, and $b_2^2 = 2$ is a transmission strategy that permits universal recovery.
- Given a transmission strategy that permits universal recovery, the encoding operations can be efficiently computed using the algorithm developed by Jaggi et al.

Transmission Strategies in $\mathcal{R}_r$ Permit Universal Recovery

Theorem

For a fixed number of communication rounds r, a transmission strategy in which node i makes at most b_i^j transmissions during round j permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Transmission Strategies in $\mathcal{R}_r$ Permit Universal Recovery

Theorem

For a fixed number of communication rounds r, a transmission strategy in which node i makes at most b_i^j transmissions during round j permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Proof.

1. Create an equivalent single-source network coding problem.

Transmission Strategies in $\mathcal{R}_r$ Permit Universal Recovery

Theorem

For a fixed number of communication rounds r, a transmission strategy in which node i makes at most b_i^j transmissions during round j permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Proof.

1. Create an equivalent single-source network coding problem.
2. Take any source-sink cut in the network coding graph and reduce to a minimal cut.

Transmission Strategies in $\mathcal{R}_r$ Permit Universal Recovery

Theorem

For a fixed number of communication rounds r, a transmission strategy in which node i makes at most b_i^j transmissions during round j permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Proof.

1. Create an equivalent single-source network coding problem.
2. Take any source-sink cut in the network coding graph and reduce to a minimal cut.
3. Each minimal cut corresponds to a constraint defining $\mathcal{R}_r$ (and vice versa).

Transmission Strategies in $\mathcal{R}_r$ Permit Universal Recovery

Theorem

For a fixed number of communication rounds r, a transmission strategy in which node i makes at most b_i^j transmissions during round j permits universal recovery if and only if $\{b_i^j\} \in \mathcal{R}_r$.

Proof.

1. Create an equivalent single-source network coding problem.
2. Take any source-sink cut in the network coding graph and reduce to a minimal cut.
3. Each minimal cut corresponds to a constraint defining $\mathcal{R}_r$ (and vice versa).
4. Thus, a transmission strategy permits universal recovery iff $\{b_i^j\} \in \mathcal{R}_r$.

□

Neighborhood of a set $\mathcal{S}$

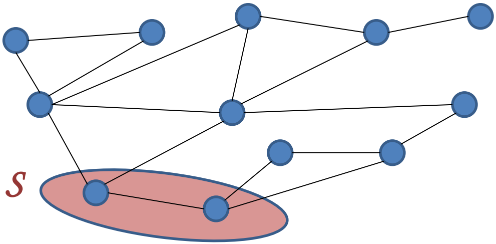

Neighborhood of a set $\mathcal{S}$

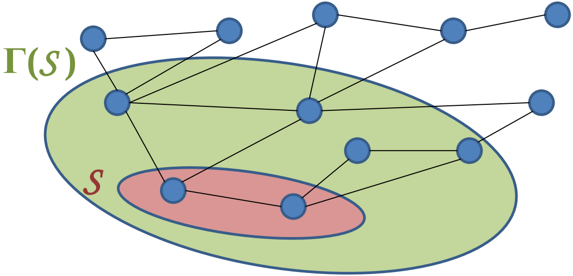

The Region $\mathcal{R}_r$

A Set of Transmission Strategies

Definition (The region $\mathcal{R}_r \subset \mathbb{Z}_+^{N\times r}$)

$\{b_i^j\} \in \mathcal{R}_r$ if and only if:

$\forall \emptyset \subsetneq \mathcal{S}_0 \subseteq \cdots \subseteq \mathcal{S}_r \subsetneq [N]$ satisfying $\mathcal{S}_j \subseteq \Gamma(\mathcal{S}_{j-1})$

for each $j \in [r]$, the following inequalities hold :

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|.$$

An Example
3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

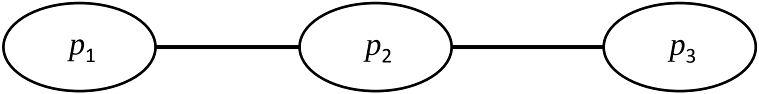

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

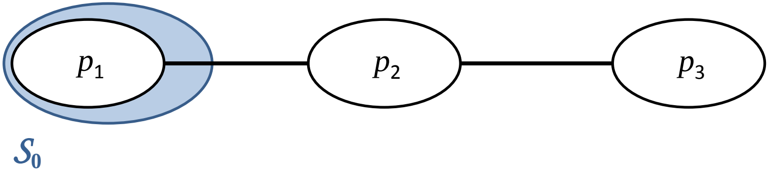

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

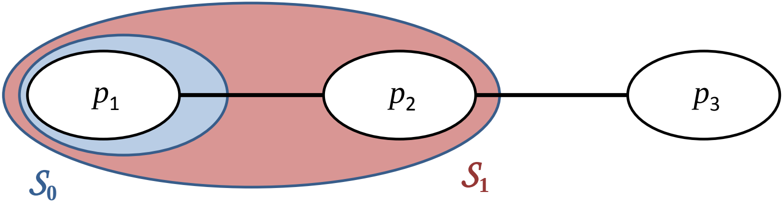

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

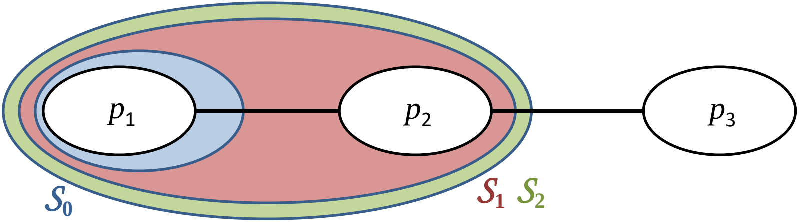

- $\mathcal{S}_0 = \{1\}$, $\mathcal{S}_1 = \mathcal{S}_2 = \{1, 2\} \Rightarrow b_3^1 \geq 1$

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

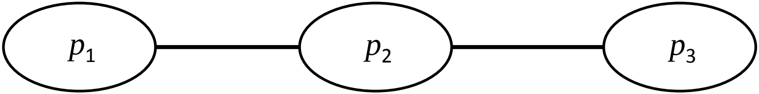

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

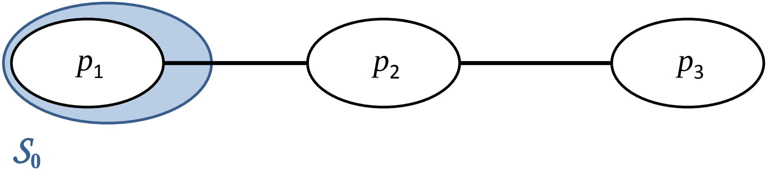

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

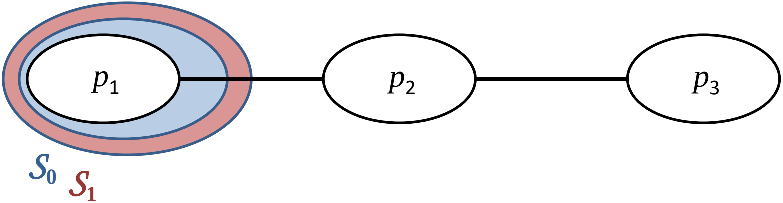

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

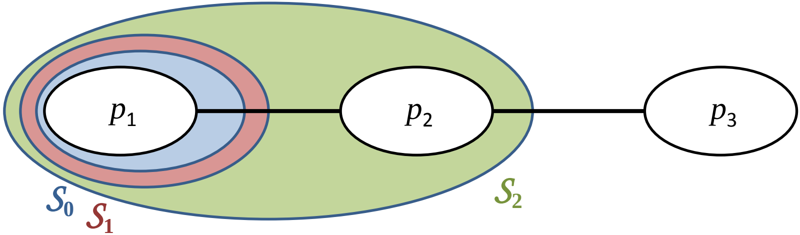

- $\mathcal{S}_0 = \mathcal{S}_1 = \{1\}$, $\mathcal{S}_2 = \{1,2\} \Rightarrow b_2^2 \geq 1$

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

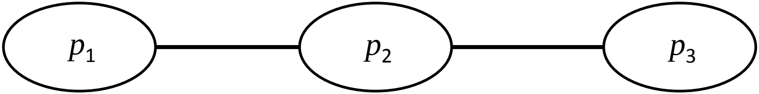

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

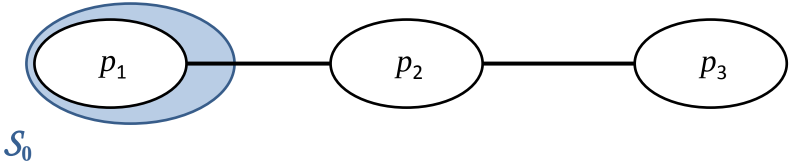

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

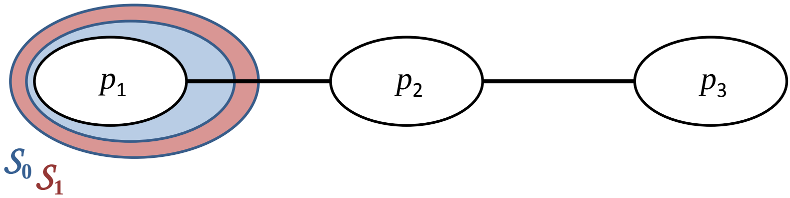

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

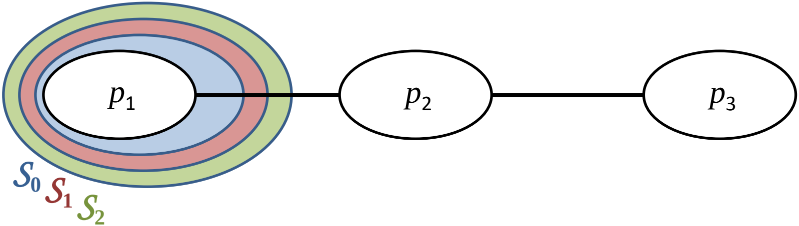

- $\mathcal{S}_0 = \mathcal{S}_1 = \mathcal{S}_2 = \{1\} \Rightarrow b_2^1 + b_2^2 \geq 2$

An Example

3-Node Line Network Revisited

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

- $\mathcal{S}_0 = \{1\}$, $\mathcal{S}_1 = \mathcal{S}_2 = \{1,2\} \Rightarrow b_3^1 \geq 1$
- $\mathcal{S}_0 = \mathcal{S}_1 = \{1\}$, $\mathcal{S}_2 = \{1,2\} \Rightarrow b_2^2 \geq 1$
- $\mathcal{S}_0 = \mathcal{S}_1 = \mathcal{S}_2 = \{1\} \Rightarrow b_2^1 + b_2^2 \geq 2$

Constraints defining $\mathcal{R}_r$

$$\sum_{j=1}^{r} \sum_{i \in \mathcal{S}_j^c \cap \Gamma(\mathcal{S}_{j-1})} b_i^{(r+1-j)} \geq \left| \bigcap_{i \in \mathcal{S}_r} \mathcal{P}_i^c \right|$$

- $\mathcal{S}_0 = \{1\}$, $\mathcal{S}_1 = \mathcal{S}_2 = \{1, 2\} \Rightarrow b_3^1 \geq 1$
- $\mathcal{S}_0 = \mathcal{S}_1 = \{1\}$, $\mathcal{S}_2 = \{1, 2\} \Rightarrow b_2^2 \geq 1$
- $\mathcal{S}_0 = \mathcal{S}_1 = \mathcal{S}_2 = \{1\} \Rightarrow b_2^1 + b_2^2 \geq 2$
- By symmetry: $b_1^1 \geq 1$. All other sequences of sets give redundant constraints.

Sequences of Networks

- We can consider sequences of networks indexed by k.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^{N}$.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^N$.
- Naturally, we have $|\bigcup_{i=1}^N \mathcal{P}_i(k)| = k$.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^N$.
- Naturally, we have $|\bigcup_{i=1}^N \mathcal{P}_i(k)| = k$.

Definition (Well-behaved sequence of packet distributions)

A sequence of packet distributions is well-behaved if

$$\mathbb{P}_{\mathcal{S}} \triangleq \lim_{k\to\infty} \frac{1}{k}\left|\bigcup_{i\in\mathcal{S}} \mathcal{P}_i(k)\right|$$

exists for all subsets $\mathcal{S} \subseteq \{1,\ldots,N\}$.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^N$.
- Naturally, we have $|\bigcup_{i=1}^N \mathcal{P}_i(k)| = k$.

Definition (Well-behaved sequence of packet distributions)

A sequence of packet distributions is well-behaved if

$$\mathbb{P}_{\mathcal{S}} \triangleq \lim_{k\to\infty} \frac{1}{k} \left| \bigcup_{i\in\mathcal{S}} \mathcal{P}_i(k) \right|$$

exists for all subsets $\mathcal{S} \subseteq \{1, \ldots, N\}$.

- $\mathbb{P}_{\mathcal{S}}$ is the limit of the empirical probability that any node in $\mathcal{S}$ receives a particular packet.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^N$.
- Naturally, we have $|\bigcup_{i=1}^N \mathcal{P}_i(k)| = k$.

Definition (Well-behaved sequence of packet distributions)

A sequence of packet distributions is well-behaved if

$$\mathbb{P}_{\mathcal{S}} \triangleq \lim_{k\to\infty} \frac{1}{k} \left| \bigcup_{i\in\mathcal{S}} \mathcal{P}_i(k) \right|$$

exists for all subsets $\mathcal{S} \subseteq \{1, \ldots, N\}$.

- $\mathbb{P}_{\mathcal{S}}$ is the limit of the empirical probability that any node in $\mathcal{S}$ receives a particular packet.
- Let $\mathbb{P}_{\mathcal{S}}^c = 1 - \mathbb{P}_{\mathcal{S}}$.

Sequences of Networks

- We can consider sequences of networks indexed by k.
- The packet distributions are denoted $\{\mathcal{P}_i(k)\}_{i=1}^N$.
- Naturally, we have $|\bigcup_{i=1}^N \mathcal{P}_i(k)| = k$.

Definition (Well-behaved sequence of packet distributions)

A sequence of packet distributions is well-behaved if

$$\mathbb{P}_{\mathcal{S}} \triangleq \lim_{k\to\infty} \frac{1}{k} \left| \bigcup_{i\in\mathcal{S}} \mathcal{P}_i(k) \right|$$

exists for all subsets $\mathcal{S} \subseteq \{1, \ldots, N\}$.

- $\mathbb{P}_{\mathcal{S}}$ is the limit of the empirical probability that any node in $\mathcal{S}$ receives a particular packet.
- Let $\mathbb{P}_{\mathcal{S}}^c = 1 - \mathbb{P}_{\mathcal{S}}$.
- Convergence can take place in probability.

Border Nodes

Characterized by simple cuts

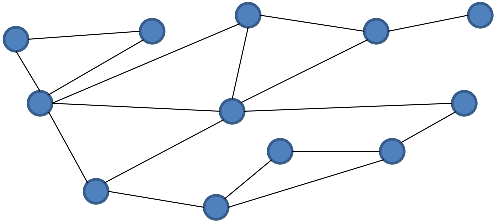

Border Nodes

Characterized by simple cuts

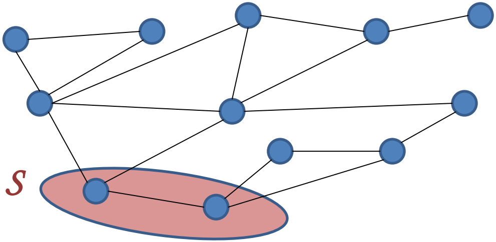

Border Nodes

Characterized by simple cuts

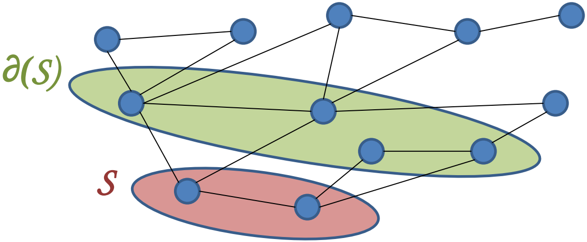

Per-round Transmission Constraints

Simple Cuts Suffice

Theorem

For a fixed network topology, $\tau \triangleq k/r$ fixed, and a well-behaved sequence of packet distributions $\{\mathcal{P}_i(k)\}_{i=1}^N$, if node i is allowed to transmit b_i times per round and

$$\sum_{i \in \partial(\mathcal{S})} b_i > \tau \mathbb{P}_{\mathcal{S}}^c, \quad \forall\, \emptyset \subsetneq \mathcal{S} \subsetneq [N]$$

then universal recovery is possible for all sufficiently large k.

Per-round Transmission Constraints

Simple Cuts Suffice

Theorem

For a fixed network topology, $\tau \triangleq k/r$ fixed, and a well-behaved sequence of packet distributions $\{\mathcal{P}_i(k)\}_{i=1}^N$, if node i is allowed to transmit b_i times per round and

$$\sum_{i \in \partial(\mathcal{S})} b_i > \tau \mathbb{P}_{\mathcal{S}}^c, \qquad \forall\, \emptyset \subsetneq \mathcal{S} \subsetneq [N]$$

then universal recovery is possible for all sufficiently large k.

- Interpretation: If border nodes have the ability to relay the packets $\mathcal{S}$ is missing, then universal recovery is possible.

Per-round Transmission Constraints

A corresponding converse

Theorem

Universal recovery is not possible if node i is allowed to make at most b_i transmissions per communication round and there is some set $\emptyset \subsetneq \mathcal{S} \subsetneq [N]$ for which

$$\sum_{i \in \partial(\mathcal{S})} b_i < \frac{1}{r} \left| \bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c \right|.$$

Per-round Transmission Constraints

A corresponding converse

Theorem

Universal recovery is not possible if node i is allowed to make at most b_i transmissions per communication round and there is some set $\emptyset \subsetneq \mathcal{S} \subsetneq [N]$ for which

$$\sum_{i \in \partial(\mathcal{S})} b_i < \frac{1}{r} \left| \bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c \right|.$$

- Interpretation: The nodes in $\partial(\mathcal{S})$ form a bottleneck which is "too tight".

Per-round Transmission Constraints

A corresponding converse

Theorem

Universal recovery is not possible if node i is allowed to make at most b_i transmissions per communication round and there is some set $\emptyset \subsetneq \mathcal{S} \subsetneq [N]$ for which

$$\sum_{i \in \partial(\mathcal{S})} b_i < \frac{1}{r} \left| \bigcap_{i \in \mathcal{S}} \mathcal{P}_i^c \right|.$$

- Interpretation: The nodes in $\partial(\mathcal{S})$ form a bottleneck which is "too tight".

Question

What is the minimum number of transmissions required for universal recovery?

Nonsingular, d-regular, d-connected Networks

Definitions and Remarks

Definition (d-regular Networks)

A network is d-regular if each node is adjacent to d other nodes.

Nonsingular, d-regular, d-connected Networks

Definition (d-regular Networks)

A network is d-regular if each node is adjacent to d other nodes.

Definition (d-connected Networks)

A network is d-connected if at least d nodes must be removed in order to disconnect the network.

Nonsingular, d-regular, d-connected Networks

Definitions and Remarks

Definition (d-regular Networks)

A network is d-regular if each node is adjacent to d other nodes.

Definition (d-connected Networks)

A network is d-connected if at least d nodes must be removed in order to disconnect the network.

- Theorem (Bollobás): Almost every d-regular network is d-connected.

Nonsingular, d-regular, d-connected Networks

Definitions and Remarks

Definition (d-regular Networks)

A network is d-regular if each node is adjacent to d other nodes.

Definition (d-connected Networks)

A network is d-connected if at least d nodes must be removed in order to disconnect the network.

- Theorem (Bollobás): Almost every d-regular network is d-connected.

Definition (Nonsingular Networks)

A network is nonsingular if its adjacency matrix is nonsingular.

Definition (d-regular Networks)

A network is d-regular if each node is adjacent to d other nodes.

Definition (d-connected Networks)

A network is d-connected if at least d nodes must be removed in order to disconnect the network.

- Theorem (Bollobás): Almost every d-regular network is d-connected.

Definition (Nonsingular Networks)

A network is nonsingular if its adjacency matrix is nonsingular.

- Conjecture (Costello and Vu): For $d \geq 3$, almost every d-regular network is nonsingular.

Bounding Excess Transmissions by the Number of Nodes

A theorem for structured networks

Theorem

For all nonsingular d-regular, d-connected networks with $\rho > 0$ fixed,

$$\mathbb{P}_i = \rho < \mathbb{P}_{\mathcal{S}}, \quad \forall i \in [N], \ \forall \mathcal{S} : \{i\} \subsetneq \mathcal{S},$$

and

$$(N - |\mathcal{S}|) \cdot \mathbb{P}_i^c > d \cdot \mathbb{P}_{\mathcal{S}}^c, \quad \forall \mathcal{S} : |\mathcal{S}| > N - d$$

we can explicitly construct a vector $\{b_i^j\}$ that is within N transmissions of optimum. Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i \in [N] \\ j \in [r]}} b_i^j \leq \frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| + N.$$

Packets Available to d Nodes

Corollary

For all nonsingular d-regular, d-connected networks with each packet available to at least d nodes, and

$$\mathbb{P}_i = \rho < \mathbb{P}_{\mathcal{S}}, \quad \forall i \in [N], \; \forall \mathcal{S} : \{i\} \subsetneq \mathcal{S},$$

we can explicitly construct a vector $\{b_i^j\}$ that is within N transmissions of optimum. Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i \in [N] \\ j \in [r]}} b_i^j \leq \frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| + N.$$

- If each packet is originally available to d nodes, then almost all transmissions are useful.

Sparsely Distributed Packets

Corollary

For all nonsingular d-regular, d-connected networks with $0 < \rho < \frac{1}{d+1}$ fixed, and

$$\mathbb{P}_i = \rho < \mathbb{P}_{\mathcal{S}}, \quad \forall i \in [N], \; \forall \mathcal{S} : \{i\} \subsetneq \mathcal{S},$$

we can explicitly construct a vector $\{b_i^j\}$ that is within N transmissions of optimum. Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i \in [N] \\ j \in [r]}} b_i^j \leq \frac{1}{d} \sum_{i \in [N]} |\mathcal{P}_i^c(k)| + N.$$

- If packets are sparsely distributed and each node has $\sim \rho \times k$ packets, then almost every transmission is useful.

Independent and Identically Distributed Packets

Corollary

For all nonsingular d-regular, d-connected networks with

$$\mathbb{P}_{\mathcal{S}} = \frac{1-(1-q)^{|\mathcal{S}|}}{1-(1-q)^N}, \ \forall \mathcal{S} \subseteq [N]$$

for some $0 < q < 1$, we can explicitly construct a vector $\{b_i^j\}$ that is within N transmissions of optimum. Moreover, for this $\{b_i^j\}$, we have that:

$$\frac{1}{d}\sum_{i\in[N]} |\mathcal{P}_i^c(k)| \leq \sum_{\substack{i\in[N]\\ j\in[r]}} b_i^j \leq \frac{1}{d}\sum_{i\in[N]} |\mathcal{P}_i^c(k)| + N.$$

- If packets are available independently at each node with probability q, then almost every transmission is useful in recovering the set of packets available to the network.

Examples

Two hypothetical scenarios

$N = 10$, Ring Network

Ten servers connected in a ring collectively have 1000 files. Each server has a unique 200-file subset. Between 4000 and 4010 file transfers are required for each server to recover all files.

$N = 1000$, Fully Connected Network

1000 peers in a P2P network collectively have a movie consisting of 1,000,000 packets. Each peer has a random half of the movie (i.e. 500,000 randomly selected packets). Between 500,500 and 501,500 packet transmissions are required for universal recovery.

Conclusions and Extensions

- Main Results
  - Defined a region $\mathcal{R}_r$ which characterizes all transmission strategies that permit universal recovery.
  - Proved that it suffices to consider simple cuts in networks with a per-round transmission constraint.
  - Bounded excess transmissions by the number of nodes for many structured networks.
- Extensions
  - Can include scenarios where helper nodes are present, etc.
  - Convergence in probability can replace "Well-behaved" and all results hold with arbitrarily high probability as the total number of packets grows.

Thank You!