ON THROUGHPUT BENEFITS OF NETWORK CODING

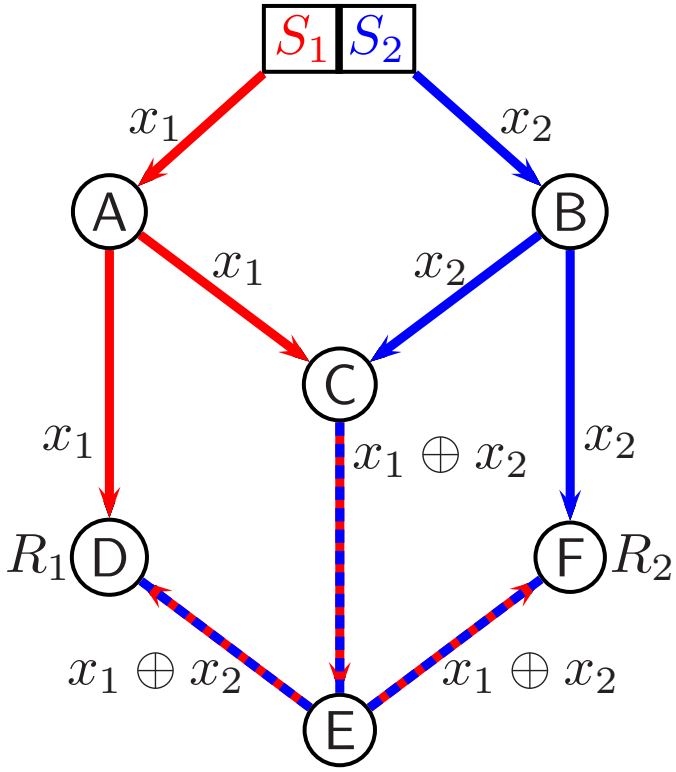
Emina Soljanin
Bell Labs

Joint work with C. Chekuri, Bell Labs, and C. Fragouli, EPFL

September 2, 2005
ACORN/CUBIN Information Theory Workshop

NETWORK MULTICAST

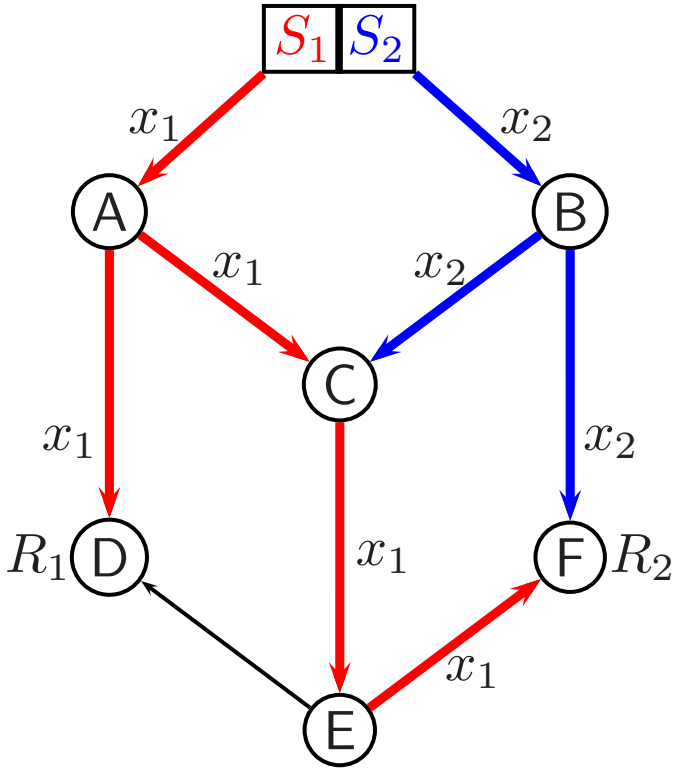
Example 1



Each receiver gets 2 bits. The average throughput is 2.

NETWORK MULTICAST

Example 1



With time sharing, the average throughput is $\frac{3}{2} = \frac{3}{4} \cdot 2$.

OUTLINE OF THE TALK

- Network coding basics
- Examples of networks with small coding benefits
- Examples of networks with large coding benefits
- Some general results on coding benefits
- On alphabet size

COMMUNICATION NETWORK
A Mathematical Model

- A network is represented as a directed acyclic graph $G = (V, E)$.
- G has unit-capacity edges and parallel edges are allowed.
- There are h unit-rate information sources $S_1, \ldots, S_h$.
- There are N receivers $R_1, \ldots, R_N$ located at N distinct nodes.
- Can all sources simultaneously transmit information to all receivers? (We are interested in the transmission at full rate.)

NETWORK MULTICAST
Throughput

- Can all sources simultaneously transmit information to receiver j?

  Yes, if between the sources and the j-th receiver node

  – the number of edges in the min-cut is h (or equivalently)
  – there are h edge-disjoint paths (S_i, R_j) for $1 \leq i \leq h$.

  [Ford, Fulkerson], [Elias, Feinstein, Shannon] $\sim$ 50s

- Can all sources simultaneously transmit information to all receivers?

  Yes, if in addition each node of G can re-encode information.

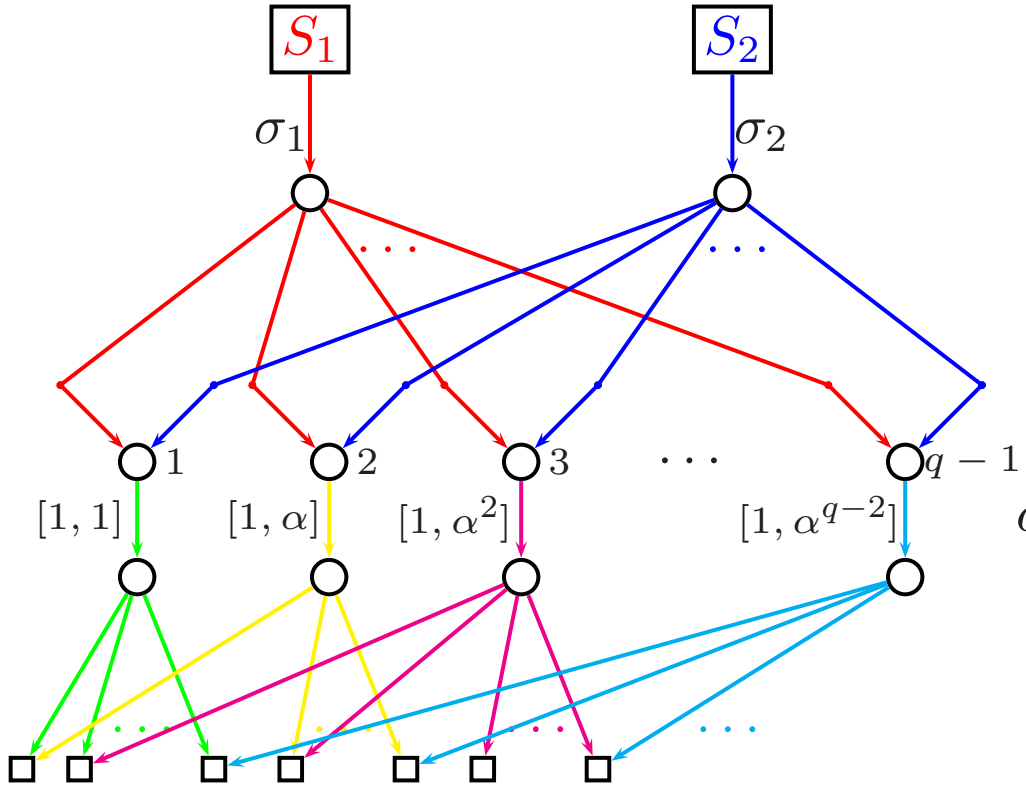
  [Alshwede, Cai, Li, Yeung] $\sim$ 2000

NETWORK MULTICAST
Linear Combining

- Source S_i emits σ_i which is an element of some finite field $\mathbb{F}_q$.
- Each edge carries a linear combination of its parent node inputs.
- Consequently,
  each edge carries a linear combination of source symbols.
- The h edges a receiver observes
  should carry independent linear combination of source symbols.
- Network code design problem: How should nodes combine their inputs to ensure this independence?

NETWORK MULTICAST

An Example



α is a primitive element of $\mathbb{F}_q$.

$\leftarrow \binom{q-1}{2}$ receivers.

A receiver observes $\sigma_1 + \alpha^i \sigma_2$ and $\sigma_1 + \alpha^j \sigma_2$.

AVERAGE THROUGHPUT WITH ROUTING

- T_i^r is the throughput to receiver i when only routing is used.
- Compare the average routing throughput with coding throughput:

$$T^r = \frac{1}{N}\sum_{i=1}^{N} T_i^r \quad \text{vs} \quad T^c = \frac{1}{N}\sum_{i=1}^{N} T_i^c = 2.$$

- Route σ_1 through one half of the $q-1$ intermediate nodes, and σ_2 through the other half.
- The routing throughput (for even $q-1$):

$$T^r = \frac{1}{\binom{q-1}{2}}\left[\frac{q-1}{2}\left(\frac{q-1}{2}-1\right)\cdot 1 + \left(\frac{q-1}{2}\right)^2 \cdot 2\right] > \frac{3}{4}\cdot 2$$

CODING v.s. ROUTING

Throughput Measures

- $T_j^c = h$ is the throughput to receiver j when coding is used.
- T_j^r is the throughput to receiver j when only routing is used.
- Compare the average routing throughput with coding throughput:

$$T^r = \frac{1}{N}\sum_{i=1}^{N} T_i^r \quad \text{vs} \quad T^c = \frac{1}{N}\sum_{i=1}^{N} T_i^c = h.$$

(We are interested in the maximal routing throughput.)

- There are other throughput measures, and coding benefits depend on the measure.

CODING v.s. ROUTING
Measures and Scenarios

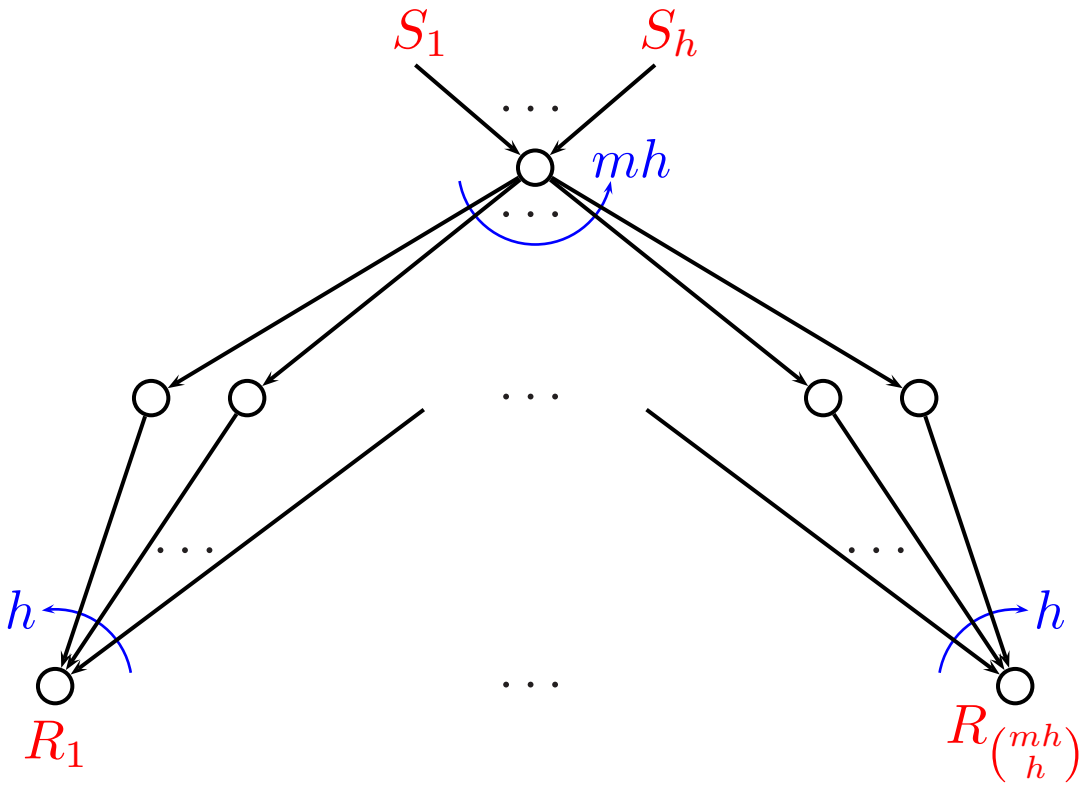
- Individual receiver requirements: symmetric or sum.
- Type of routing: integral or fractional.
- Communications schemes: time sharing or channel coding.
- Type of network: directed or undirected.

General results for symmetric throughput and fractional routing:

1. In undirected networks, coding can at most double the throughput. (Li&Li '04)
2. In both undirected and directed networks, coding benefits are equal to the integrality gap of the standard Steiner tree problem. (Agrawal&Charikar '04).

NETWORK MULTICAST

An Example



Encode the h source symbols by an (mh, h) RS code.

AVERAGE THROUGHPUT WITH ROUTING

An Occupancy Model

- There are h urns (sources) and each receiver has h balls.
- $\mu_0(h)$ is the number of unobserved sources (empty urns):

$$\Pr\{\mu_0(h) = k\} = \binom{h}{k}\left(1 - \frac{k}{h}\right)^h \Pr\{\mu_0(h-k) = 0\}$$

$$\Pr\{\mu_0(n) = 0\} = \sum_{l=0}^{n} \binom{n}{l}(-1)^l\left(1 - \frac{l}{h}\right)^h$$

- A receiver observes

$$h[1 - (1 - 1/h)^h] \gtrsim h(1 - e^{-1}) > h/2$$

sources on the average.

SYMMETRIC THROUGHPUT

A Coding Theorem

- The probability that a receiver does not observe source S_i is

$$\epsilon = (1 - 1/h)^h.$$

- The expected value of the throughput is

$$h[1 - (1 - 1/h)^h] = h(1 - \epsilon).$$

- Theorem: There exist a sequence of channel codes of rates

$$k/n \to 1 - \epsilon$$

and a routing strategy such that the symmetric throughput

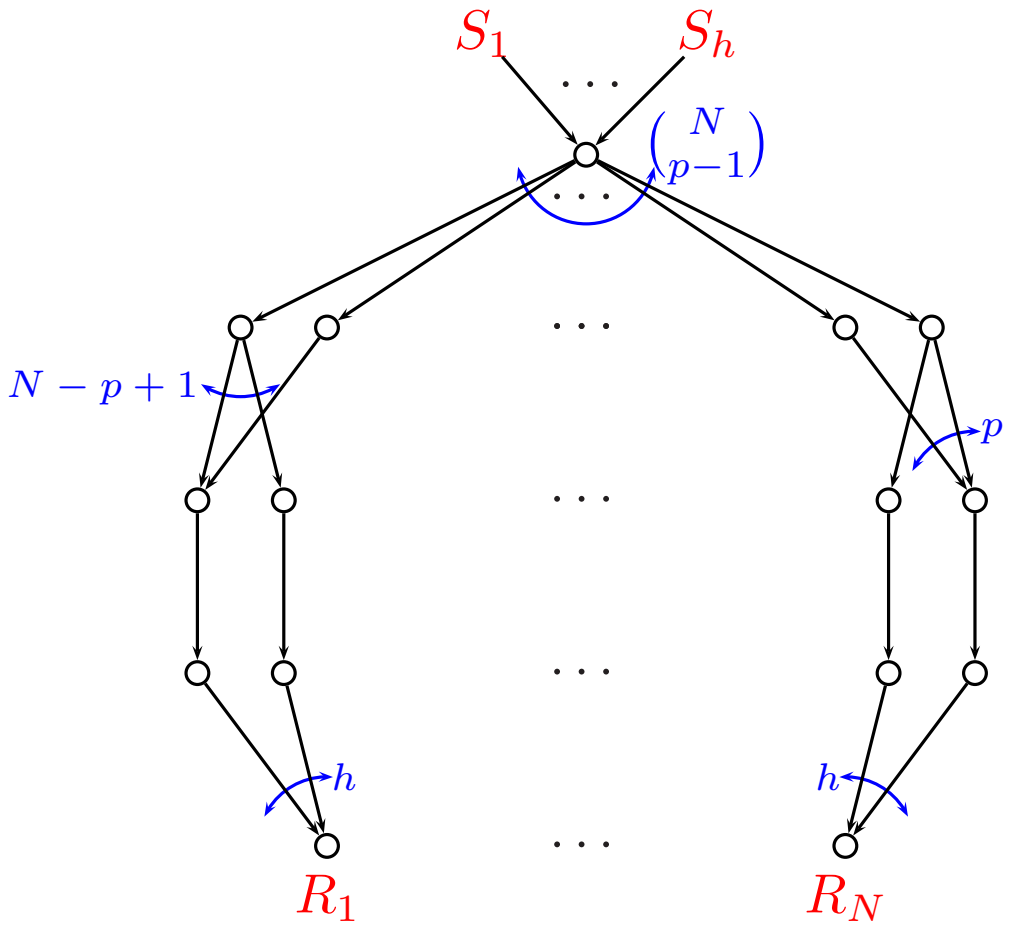
$$T_i(n) \to hk/n \to h(1 - \epsilon) \text{ as } n \to \infty.$$

AVERAGE THROUGHPUT BENEFITS OF NETWORK CODING

- In our two-source network example, $T^r \geq 3T^c/4$.
- For general networks with two sources, $T^r \geq T^c/2 + 1/N$.
- For networks with two receivers, $T^r \geq T^c/2$.
- For several other classes of networks, $T^r \geq T^c/2$.
- There are networks for which

$$T^r \simeq \frac{1}{\sqrt{N}} \cdot T^c.$$

AVERAGE THROUGHPUT WITH ROUTING



$\leftarrow$ A-nodes

$\leftarrow$ B-nodes

$\leftarrow$ C-nodes

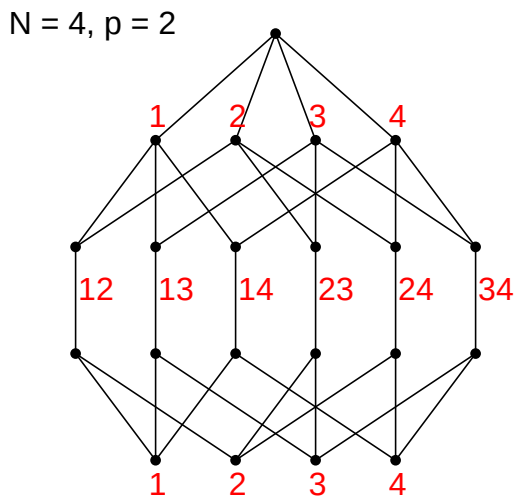
$h = \binom{N-1}{p-1}$

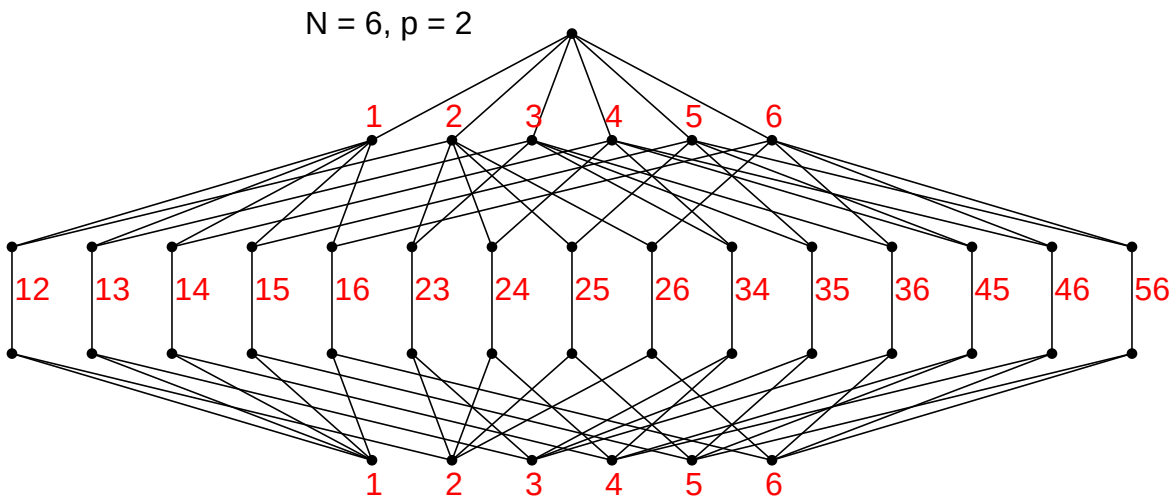
We have $\frac{T^r}{T^c} = \frac{p-1}{N-p+1} + \frac{1}{p}$.

For $p = \frac{N+1}{\sqrt{N}+1}$, we get $\frac{T^r}{T^c} = 2\,\frac{\sqrt{N}}{1+N} \lesssim \frac{2}{\sqrt{N}}$.

A NETWORK CODE FOR THE EXAMPLE

$p = 2$

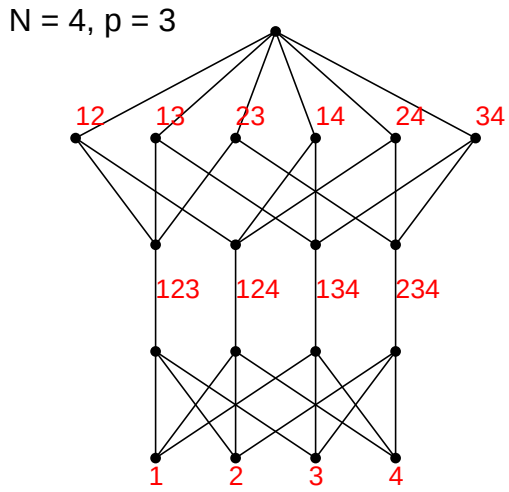


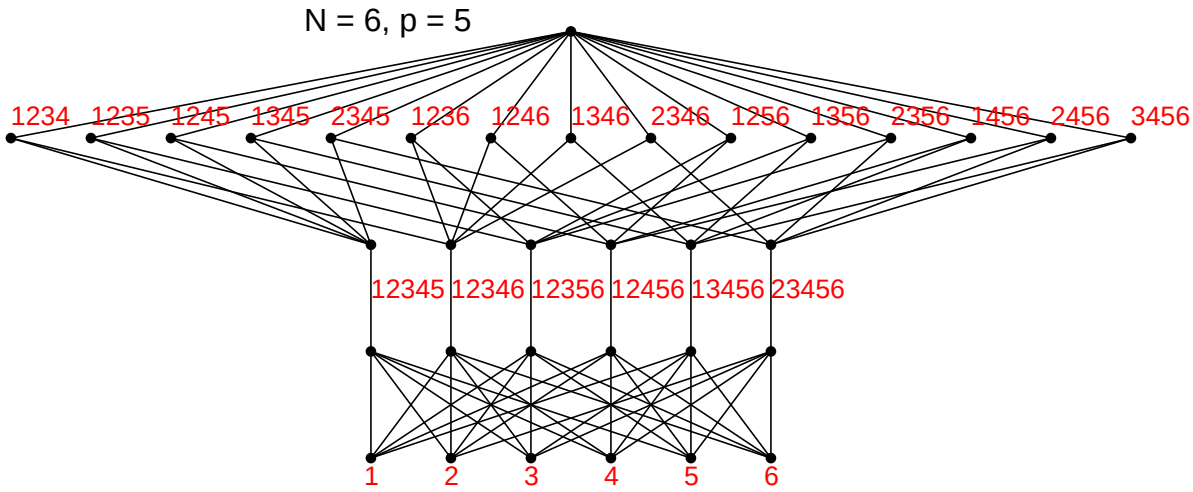


- The number of information sources is $h = N - 1$.
- Remove the A-node with label N. B-nodes add their inputs.
- Receiver i gets σ_i and $\sigma_j + \sigma_i$ for $i, j = 1, \ldots, N-1$ and $j \neq i$.

A NETWORK CODE FOR THE EXAMPLE

$$p = N - 1$$

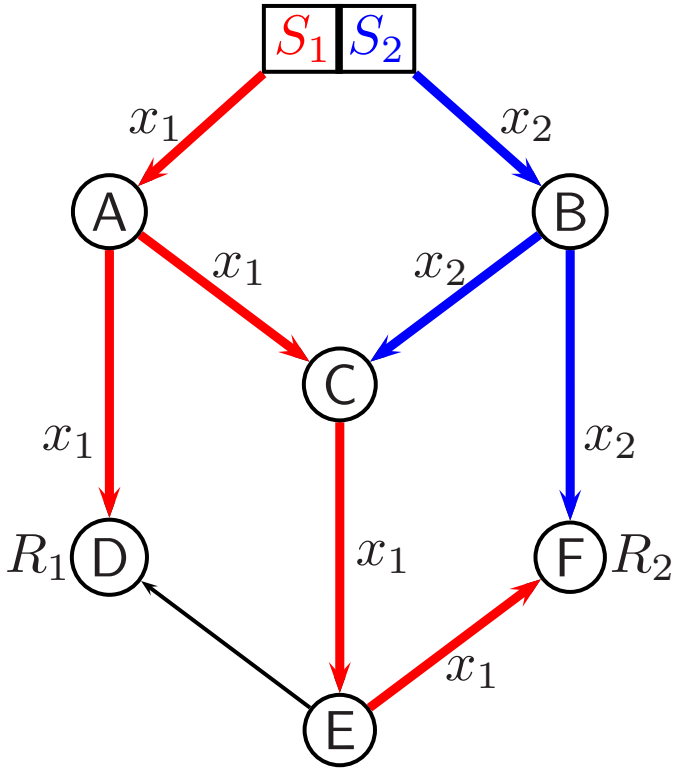




- The number of information sources is $h = N - 1$.
- Remove the A-node whose label contains N.
- The $B-C$ edges carry the following: $\sum_j \sigma_j$, σ_i for $i = 1, \ldots, N-1$.

ROUTING THROUGHPUT

Packing Partial Steiner Trees



$t_1, t_2 \in \tau \qquad n_{t_1} = 2 \quad n_{t_2} = 1$

STEINER TREE PROBLEMS

Graph Instance $\{G, S, \mathcal{R}, w, c\}$

Optimal packing $T^r(G, S, \mathcal{R}, c)$:

$$\max \sum_{t \in \tau} \frac{n_t}{N} y_t$$

$$\sum_{t \in \tau : e \in t} y_t \leq c_e, \quad \forall \ e \in E$$

$$y_t \geq 0, \quad \forall \ t \in \tau$$

Minimum weight $\mathrm{OPT}(G, w, S, \mathcal{R})$:

$$\min \sum_{e \in E} w_e x_e$$

$$\sum_{e \in \delta(\mathcal{D})} x_e \geq 1, \quad \forall\ \mathcal{D}\text{: } \mathcal{D} \text{ is separating}$$

$$x_e \in \{0, 1\}, \quad \forall\ e \in E$$

$$\beta(G, S, \mathcal{R}) = \max_c \frac{T^c(G, c, S, \mathcal{R})}{T^r(G, c, S, \mathcal{R})} \qquad \alpha(G, S, \mathcal{R}) = \max_w \frac{\mathrm{OPT}(G, w, S, \mathcal{R})}{\mathrm{LP}(G, w, S, \mathcal{R})}$$

w^*, x^*_{LP}

Claim: $\max_{\mathcal{R}' \subseteq \mathcal{R}} \beta(G, S, \mathcal{R}') \geq \max\{1, \alpha(G, S, \mathcal{R}) / \sum_{j=1}^{N} 1/j\}$

Note: $$\beta(G, S, \mathcal{R}) \geq \frac{T^c(G, c = x^*, S, \mathcal{R})}{T^r(G, c = x^*, S, \mathcal{R})} \geq \frac{1}{T^r(G, c = x^*, S, \mathcal{R})} = \frac{1}{\sum \frac{n_t}{N} y^*_t}$$

STEINER TREE PROBLEMS

Graph Instance $\{G, S, \mathcal{R}, w, c\}$

- There exists a partial tree t_1 of weight $w_{t_1} = \sum_{e \in t_1} w_e^* x_e^*$ such that

$$w_{t_1} \leq \frac{1}{\sum_{t \in \tau} \frac{n_t}{N} y_t^*} \cdot \frac{n_{t_1}}{N} \sum_{t \in \tau} w_t y_t^* \leq \boxed{\frac{1}{\sum_{t \in \tau} \frac{n_t}{N} y_t^*}}_{\leq \beta} \cdot \frac{n_{t_1}}{N} \sum_{e \in E} w_e^* x_e^*$$

- Therefore,

$$w_{t_1} \leq \beta(G, S, \mathcal{R}) \cdot \frac{n_{t_1}}{N} \sum_{e \in E} w_e^* x_e^*.$$

- If $n_1 = N$, we have

$$\beta(G, S, \mathcal{R}) \geq \frac{w_{t_1}}{\sum_{e \in E} w_e^* x_e^*} \geq \alpha(G, S, \mathcal{R}).$$

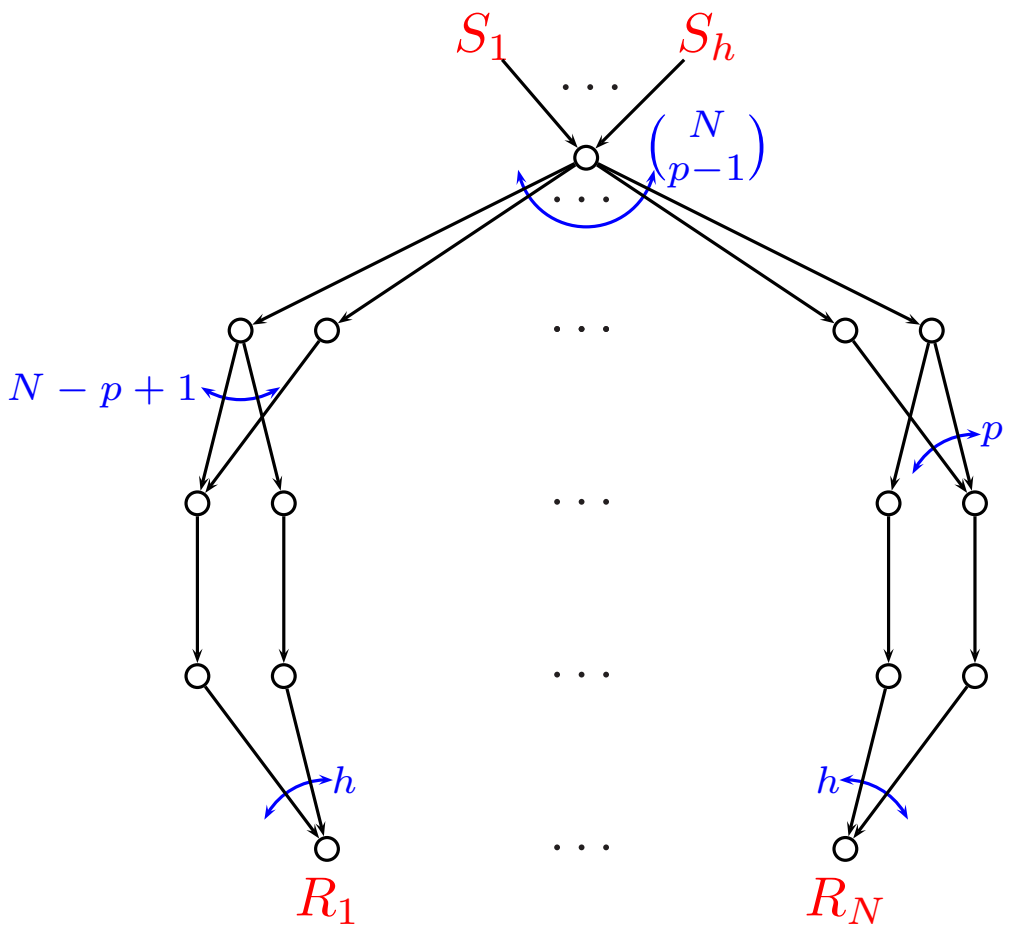
CODING ALPHABET SIZE

Theorem (Fragouli&Soljanin '04): For networks with two sources and N receivers, the code alphabet $\mathbb{F}_q$ of size

$$q = \lfloor \sqrt{2N - 7/4} + 1/2 \rfloor$$

is sufficient. There exist configurations for which it is necessary.

NETWORK CODES FOR THE EXAMPLE



NETWORK CODES FOR THE EXAMPLE

Alphabet Size

- The code design generalizes to arbitrary p.
- The binary alphabet is sufficient.
- What about coding randomly over the alphabet $\mathbb{F}_q$?
- A bound on the probability that all receivers can decode

$$P_N^d \geq \left(1 - \frac{N}{q}\right)^{\binom{N}{p}} \approxeq e^{-\frac{N\binom{N}{p}}{q}}.$$

CONCLUSIONS

- Network coding benefits
  - are related to certain combinatorial optimization problems
  - depend on the throughput measure
- Symmetric throughput
  - benefits from network coding more than the average, but
  - can approach the average in symmetric networks by channel coding and routing.
- Average throughput
  - gets at most doubled by coding for large classes of networks.