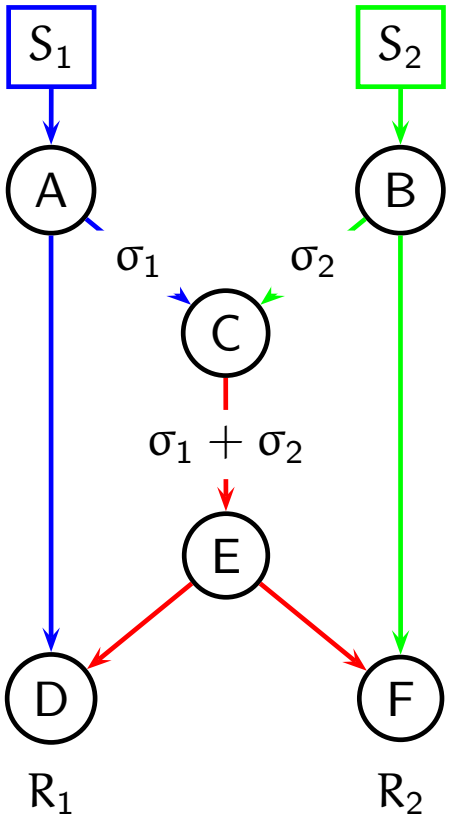
NETWORK CODING TUTORIAL

Emina Soljanin

Comm. & Stat. Research, Bell Labs

Murray Hill, June 2007

NETWORK MULTICAST – The Butterfly



- Sources S_1 and S_2 produce bits σ_1 and σ_2.
- Each receiver needs bits from both sources.
- The edges have unit capacity.

NETWORK MULTICAST MODEL

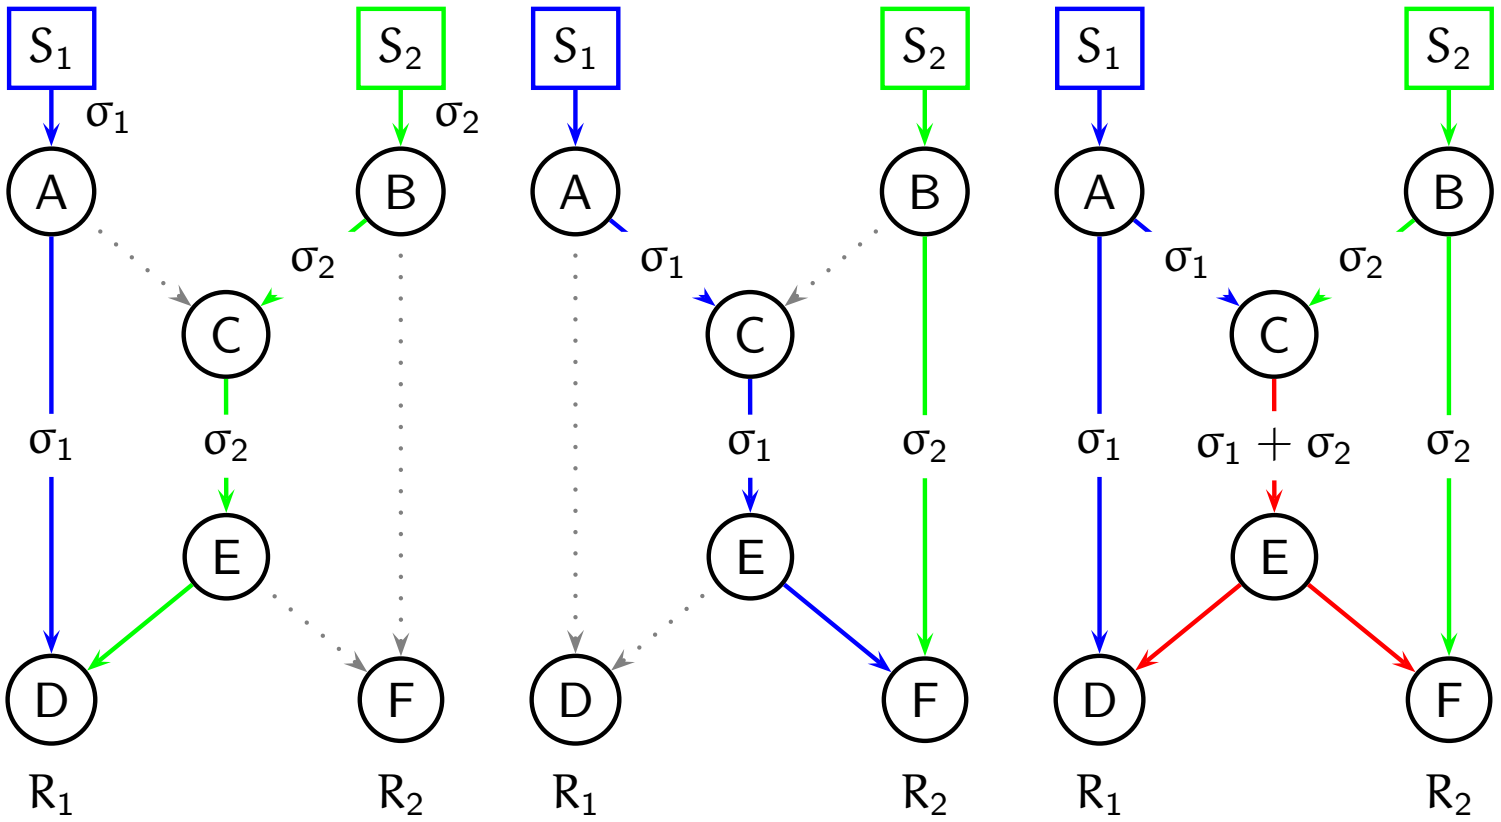
- Network is represented as a directed, acyclic graph.
- Edges have unit-capacity and parallel edges are allowed.
- There are h unit-rate information sources $S_1, \ldots, S_h$.
- There are N receivers $R_1, \ldots, R_N$ located at N distinct nodes.
- Can all sources simultaneously transmit information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

NETWORK MULTICAST – Throughput

- Can all sources simultaneously transmit to receiver R_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 $1 \leqslant i \leqslant h$.

  [Ford, Fulkerson], [Elias, Feinstein, Shannon] $\sim$ 50s
- Can all sources simultaneously transmit 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



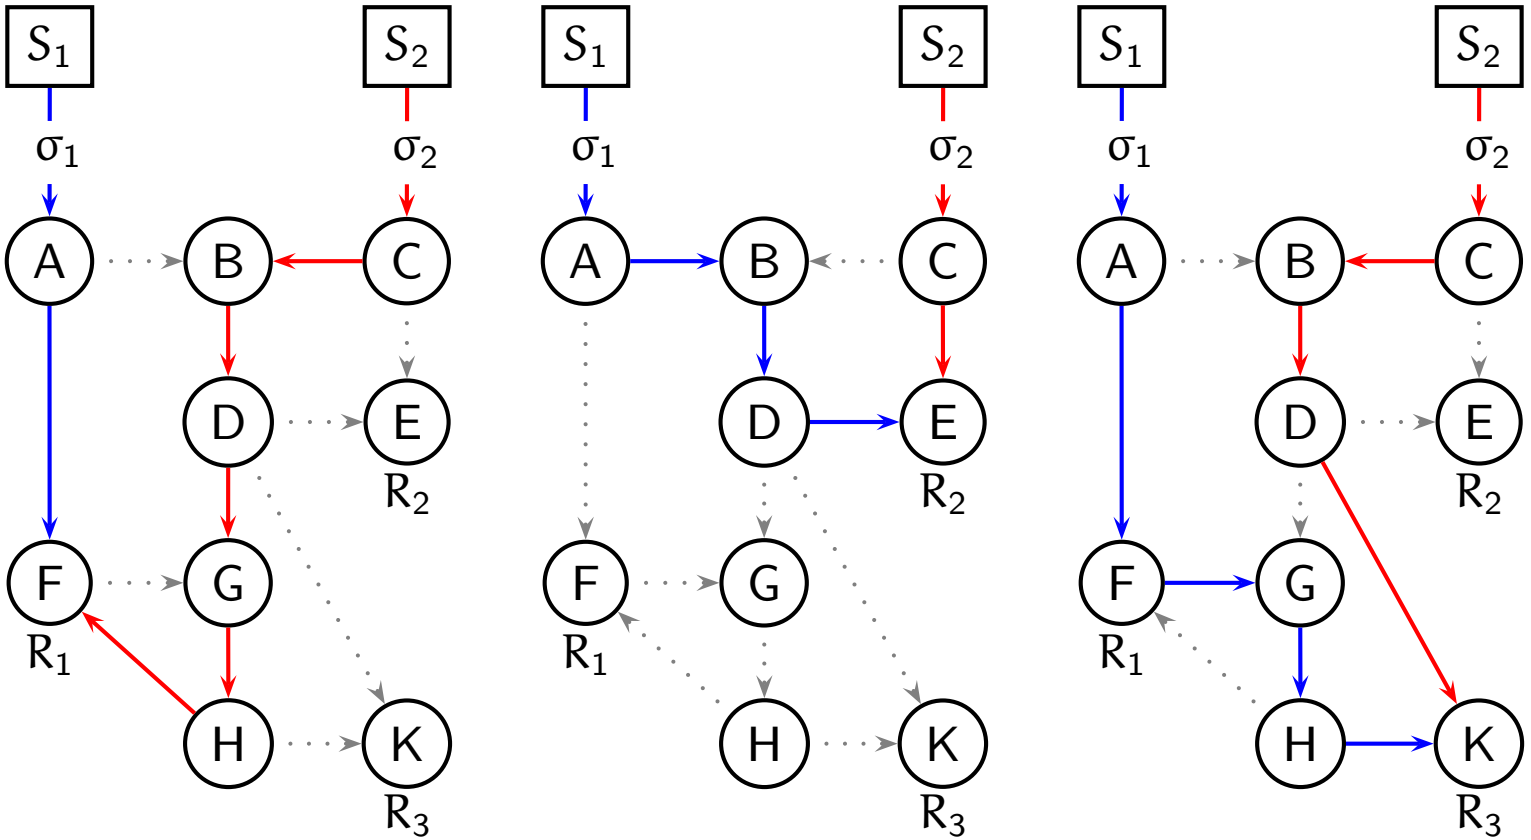
(a) Routing to R_1 (b) Routing to R_2 (c) Network coding

NETWORK MULTICAST – Linear Combining

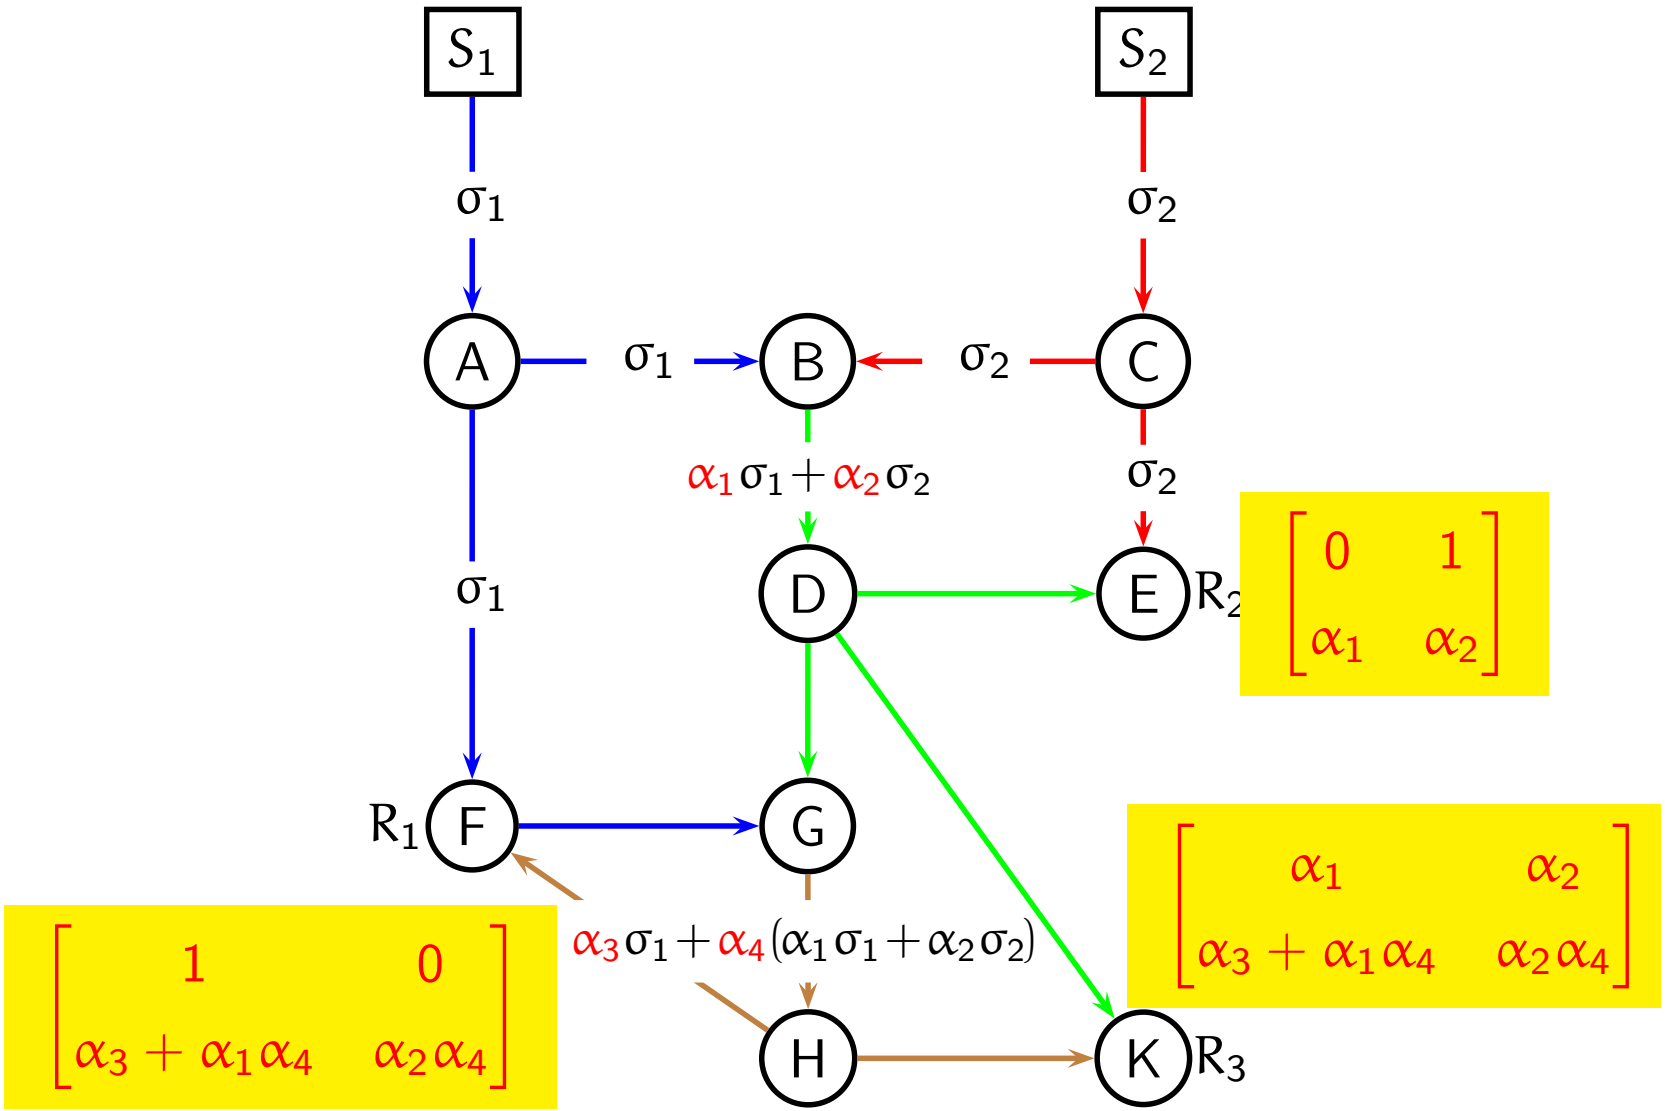
- Source S_i emits σ_i which is an element of some finite field.
- Edges carry linear combinations of their parent node inputs.
- Consequently,
  edges carry linear combinations of source symbols σ_i.
- Network Multicast Problem:

  How should nodes combine their inputs to ensure that any h edges observed by a receiver carry h independent linear combinations of source symbols?

NETWORK MULTICAST – Example



NETWORK MULTICAST – Example



NETWORK MULTICAST – Code Design

- Edges carry linear combinations of their parent node inputs.
- $\{\alpha_k\}$ are the coefficients used in these linear combinations.
- ρ_i^j is the symbol on the last edge of the path (S_i, R_j).
- Receiver j has to solve the following system of equations:

$$\begin{bmatrix} \rho_1^j \\ \vdots \\ \rho_h^j \end{bmatrix} = \mathbf{C}_j \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_h \end{bmatrix}$$

  where the elements of matrix $\mathbf{C}_j$ are polynomials in $\{\alpha_k\}$.

- The Code Design Problem:

  Select $\{\alpha_k\}$ so that all matrices $\mathbf{C}_1 \ldots \mathbf{C}_N$ are full rank.

NETWORK MULTICAST – Code Design

- The goal is to select $\{\alpha_k\}$ so that $\mathbf{C}_1 \ldots \mathbf{C}_N$ are full rank.
- Equivalently, the goal is to select $\{\alpha_k\}$ so that

$$f(\{\alpha_k\}) \triangleq \det(\mathbf{C}_1) \cdots \det(\mathbf{C}_N) \neq 0.$$

- Can such $\{\alpha_k\}$ be found? How? Over how large field?

NETWORK MULTICAST – Code Existence

Sparse Zeros Lemma:

Let

- $f(\alpha_1, \ldots, \alpha_\eta)$ be a multivariate polynomial, and
- $\mathbb{F}_q$ be the field with q elements.

such that

1. $f(\alpha_1, \ldots, \alpha_\eta)$ is not identically zero on $\mathbb{F}_q$
2. the degree of each α_k in $f(\alpha_1, \ldots, \alpha_\eta)$ is at most d
3. $q > d$

Then there exist values $p_1, \ldots, p_\eta \in \mathbb{F}_q$ such that

$$f(\alpha_1 = p_1, \ldots, \alpha_\eta = p_\eta) \neq 0.$$

THE OPERATING FIELD SIZE

- Why not any finite field?
- Polynomial

$$x(x+1)+x+x^2$$

  is identically equal to zero over $\mathbb{F}_2$.
- A polynomial not identically equal to zero over $\mathbb{F}_q$ can evaluate to zero on all elements of $\mathbb{F}_q$.
- Consider polynomial $x(x+1)$ over $\mathbb{F}_2$

THE MAIN THEOREM

Conditions:

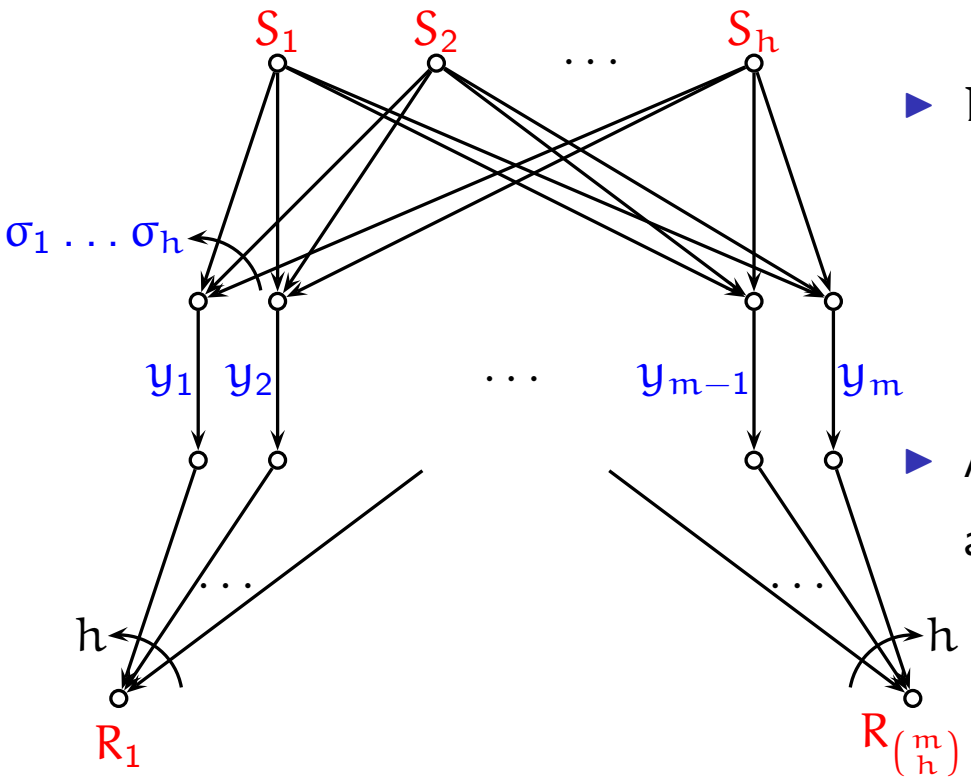
- Network is represented as a directed, acyclic graph.
- Edges have unit-capacity and parallel edges are allowed.
- There are h unit-rate information sources $S_1, \ldots, S_h$.
- There are N receivers $R_1, \ldots, R_N$ located at N distinct nodes.
- Between the sources and each receiver node,
  - the number of edges in the min-cut is h (or equivalently)
  - there are h edge-disjoint paths (S_i, R_j) for $1 \leqslant i \leqslant h$.

Claim: There exists a multicast transmission scheme of rate h.

Moreover, multicast at rate h

- cannot always be achieved by routing, but
- can be achieved by allowing the nodes to linearly combine their inputs over a sufficiently large finite field.

Combination Network $B(h, m)$



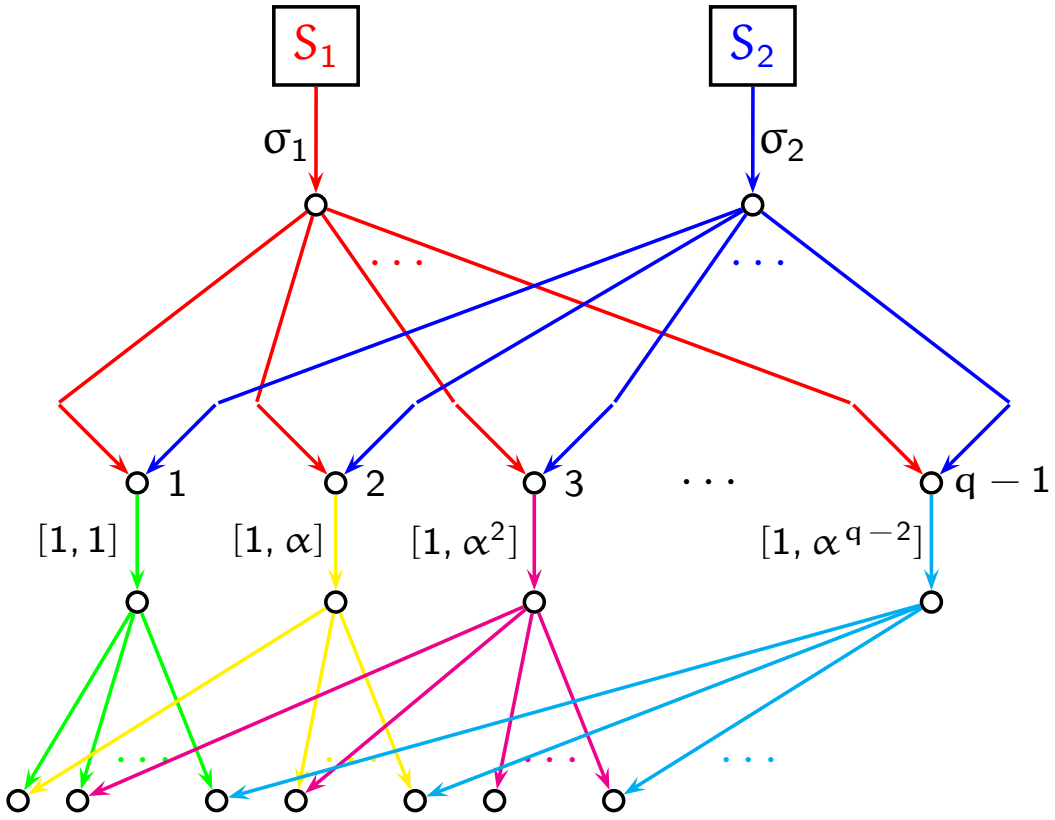
- $B(h, m)$ has
  - h information sources,
  - $\binom{m}{h}$ receivers, and
  - m bottlenecks.
- Assume $\sigma_j \in \mathbb{C}$, $j = 1, \ldots, h$, and design a rate-h multicast!

Apply an overcomplete Fourier Transform to $\{\sigma_i\}$:

$$y_k = \frac{1}{\sqrt{h}} \sum_{j=1}^{h} \sigma_j W_m^{(j-1)(k-1)}, \quad k = 1, \ldots, m, \quad W_m = e^{\imath \frac{2\pi}{m}}.$$

Each receiver can access h of the m numbers $\{y_k\}$ and recover $\{\sigma_j\}$.

$B(2, q-1)$: Coding over $\mathbb{F}_q$

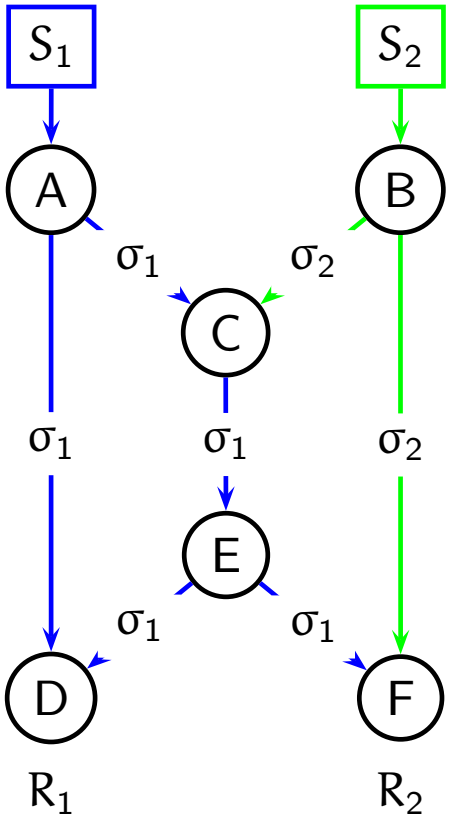


α is a primitive element of $\mathbb{F}_q$.

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

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

NETWORK MULTICAST – The Butterfly



- Sources S_1 and S_2 produce bits x_1 and x_2.
- Each receiver needs bits from both sources.
- The edges have unit capacity.

The average throughput with routing is $\frac{3}{2} = \frac{3}{4} \cdot 2$.

CODING v.s. ROUTING ALGORITHMS

- Coding is not only about increasing throughput.
- In general,
  - multicast capacity cannot be achieved by routing
  - finding an optimal routing scheme is NP-hard
- Suboptimal routing schemes are possible.
- Multicast capacity can be achieved by
  - codes that can be designed in polynomial time
  - randomized coding over a sufficiently large field

Basics
Multicast Model
Main Theorem
Summary
Examples
Benefits
Small Benefits
Coding at the Source
Large Benefits
Steiner Trees
Limited Resources
Summary

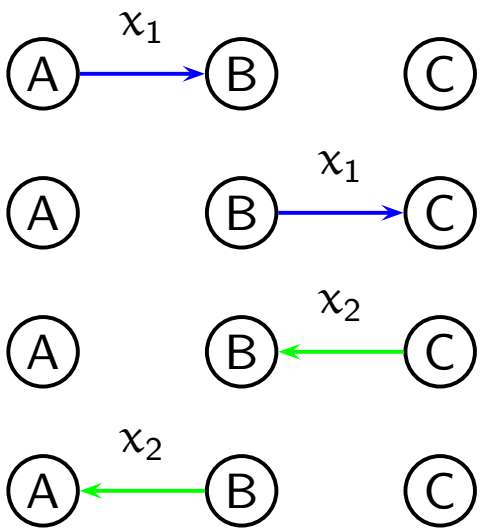
WIRELESS NETWORKS

- Example:

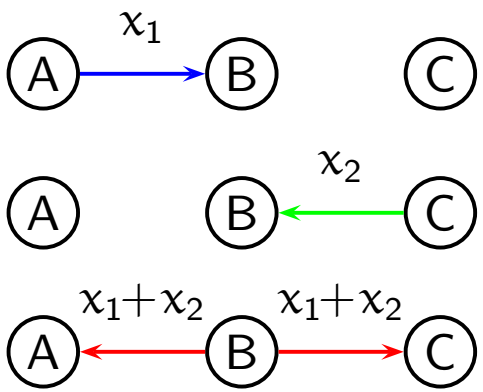


A has x_1 and C has x_2; all three nodes want both x_1 and x_2.

Without Coding



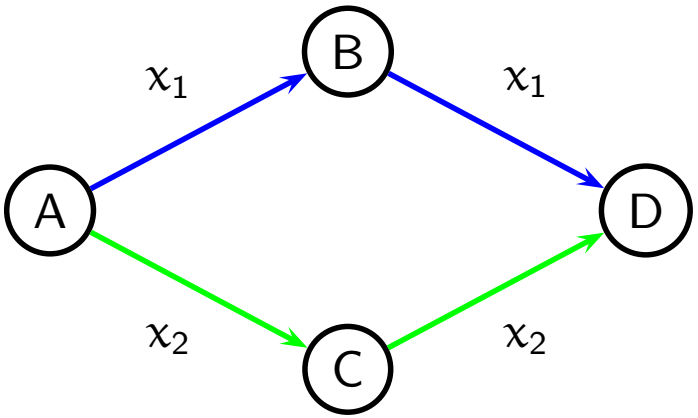
With Coding

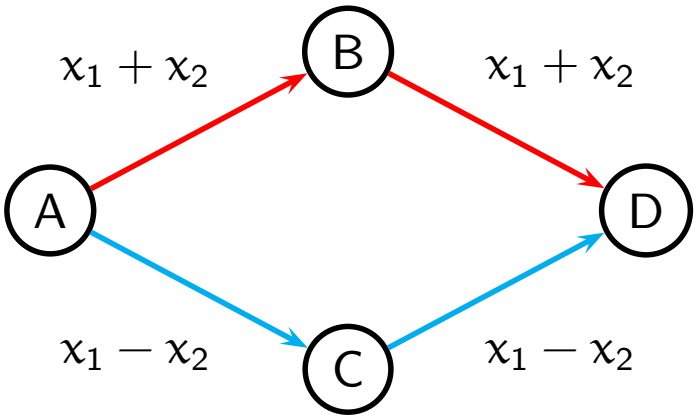


NETWORK CODING TUTORIAL
E. Soljanin
Basics
Multicast Model
Main Theorem
Summary
Examples
Benefits
Small Benefits
Coding at the Source
Large Benefits
Steiner Trees
Limited Resources
Summary

NETWORK CODING AND SECURITY

Q: What is network coding and how does it affect security?





A: Network coding affects security both favorably and adversely.

NETWORK CODING TUTORIAL
E. Soljanin
Basics
Multicast Model
Main Theorem
Summary
Examples
Benefits
Small Benefits
Coding at the Source
Large Benefits
Steiner Trees
Limited Resources
Summary

CODING v.s. ROUTING THROUGHPUT

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

$$T^r = \frac{1}{N}\sum_{j=1}^{N} T_j^r \quad \text{v.s.} \quad T^c = \frac{1}{N}\sum_{j=1}^{N} T_j^c = h.$$

(We are not interested in the maximal routing throughput.)

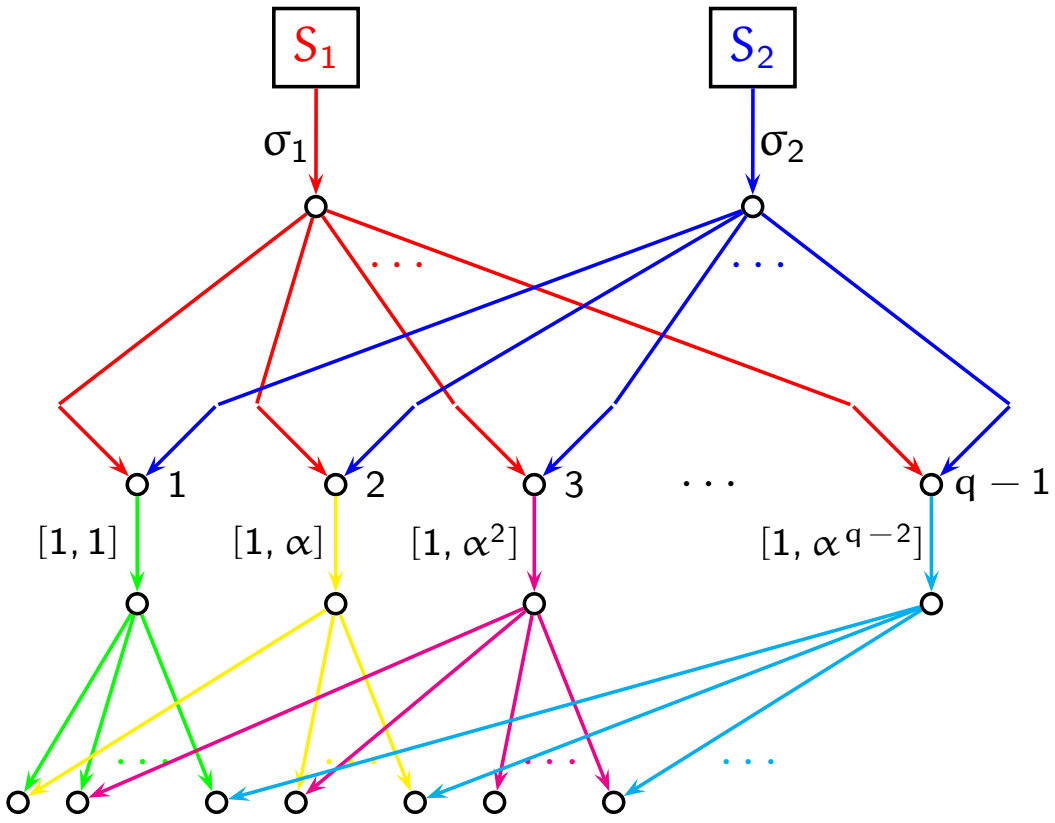
- Coding benefits T^c/T^r depend on the throughput measure.

Throughput Measures & Transmission 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.
- Results for symmetric throughput and fractional routing:
  - Undirected networks: coding at most doubles the throughput.
  - Undirected and directed networks: coding benefits are related to the integrality gap of the Steiner tree problem.

Basics
Multicast Model
Main Theorem
Summary
Examples
Benefits
Small Benefits
Coding at the Source
Large Benefits
Steiner Trees
Limited Resources
Summary

$B(2, q-1)$: Throughput with Coding over $\mathbb{F}_q$



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

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

Coding throughput is 2.

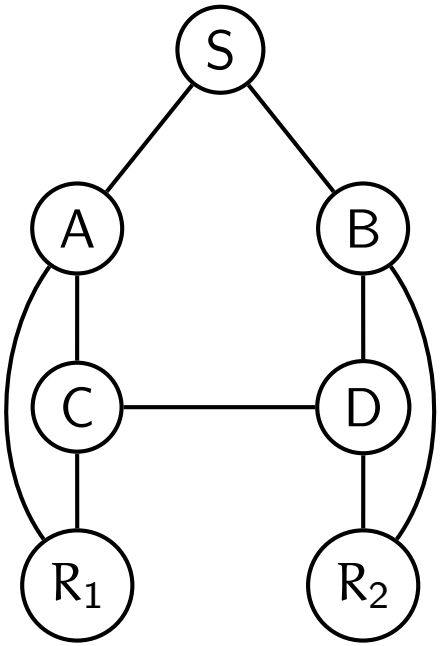
$B(2, q-1)$: Average Throughput With Routing

- The average throughput with coding is $T^c = 2$.
- Route σ_1 through one half of the $q-1$ intermediate nodes, and σ_2 through the other half.
- The average throughput with routing (for even $q-1$):

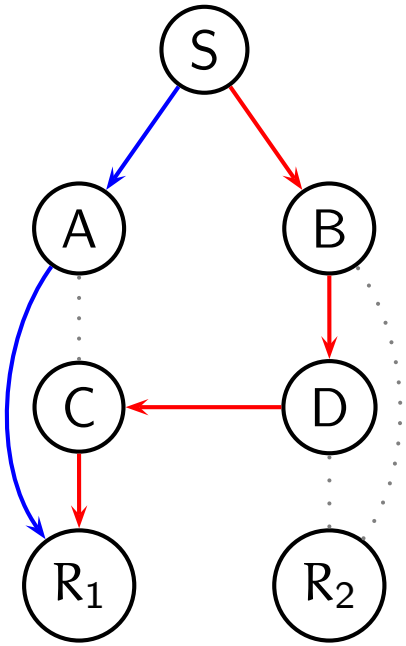
$$T^r = \frac{1}{\binom{q-1}{2}} \left[\frac{q-1}{2} \left(\frac{q-1}{2} - 1 \right) \cdot 1 + \left(\frac{q-1}{2} \right)^2 \cdot 2 \right] > \frac{3}{4} \cdot 2$$

UNDIRECTED GRAPHS

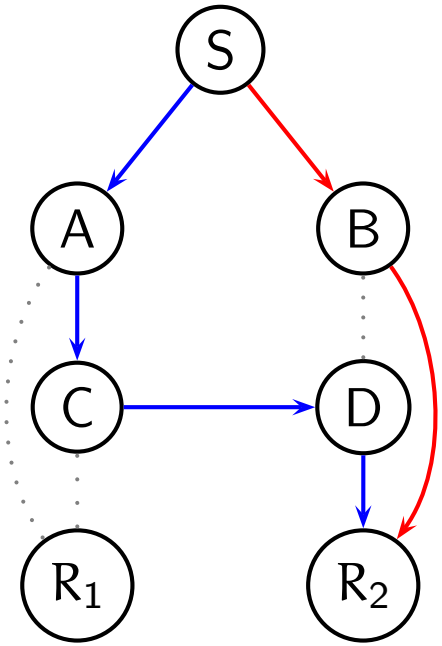
- The main theorem does not hold.
- Coding can at most double the throughput.



Original Graph

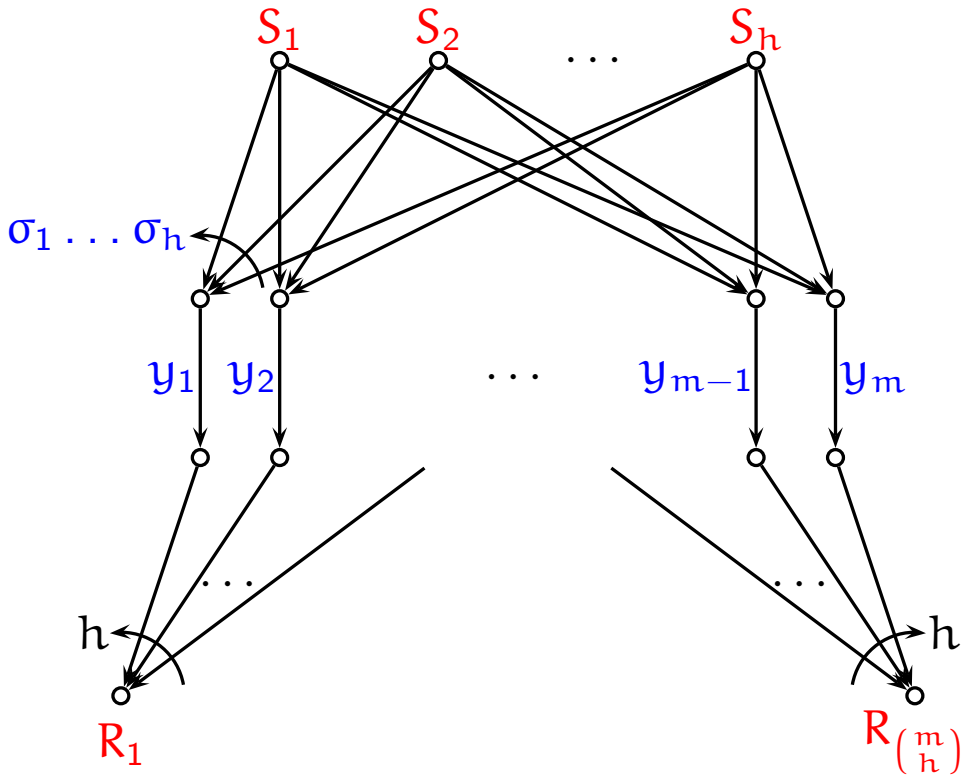


Paths to R_1



Paths to R_2

Combination Network: Coding v.s. Routing



- Coding solution:

  Encode $\sigma_1 \ldots \sigma_h$ into $y_1 \ldots y_m$ by an (m, h) RS code.

- Routing solution:

  For each bottleneck, randomly choose one of $\sigma_1 \ldots \sigma_h$.

- The routing scheme can be analyzed by an occupancy model.

AVERAGE THROUGHPUT WITH ROUTING

- There are h urns (sources) and each receiver has h balls.
- Coupon collector's problem: h possible coupons and h draws.
- $\mu_0(h)$ is the number of unobserved sources:

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

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

- A receiver observes

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

sources on the average (occupied urns, collected coupons).

NETWORK CODING TUTORIAL

E. Soljanin

Basics

Multicast Model

Main Theorem

Summary

Examples

Benefits

Small Benefits

Coding at the Source

Large Benefits

Steiner Trees

Limited Resources

Summary

A Coding Theorem for Symmetric Throughput

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

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

- The expected value of the throughput is

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

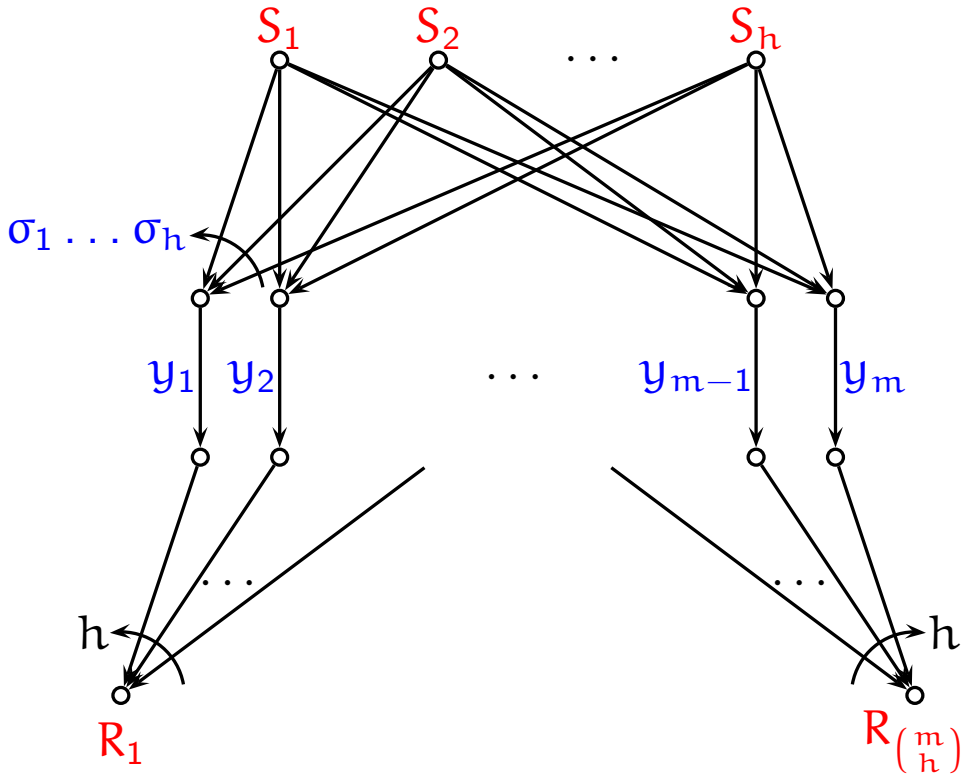
- Theorem: There exist a sequence of channel codes of rates

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

and a routing strategy such that the symmetric throughput

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

Combination Network: Symmetric Throughput



- Routing solution:

  For each bottleneck, randomly choose one of $\sigma_1 \ldots \sigma_h$.

- For $m \geqslant h^2$ and symmetric throughput measure,

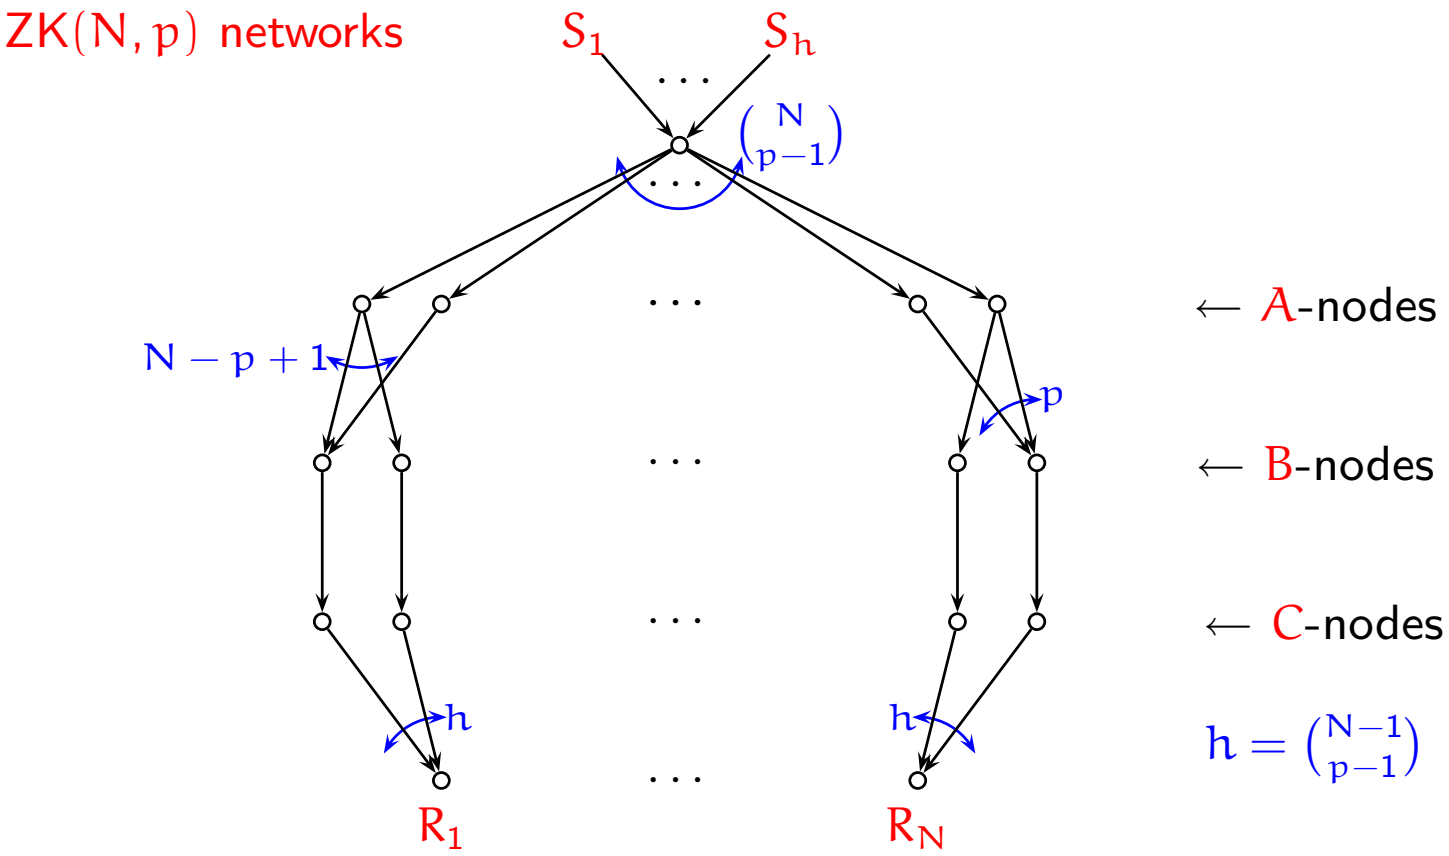
$$T^c/T^r = 1.$$

AVERAGE THROUGHPUT BENEFITS

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

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

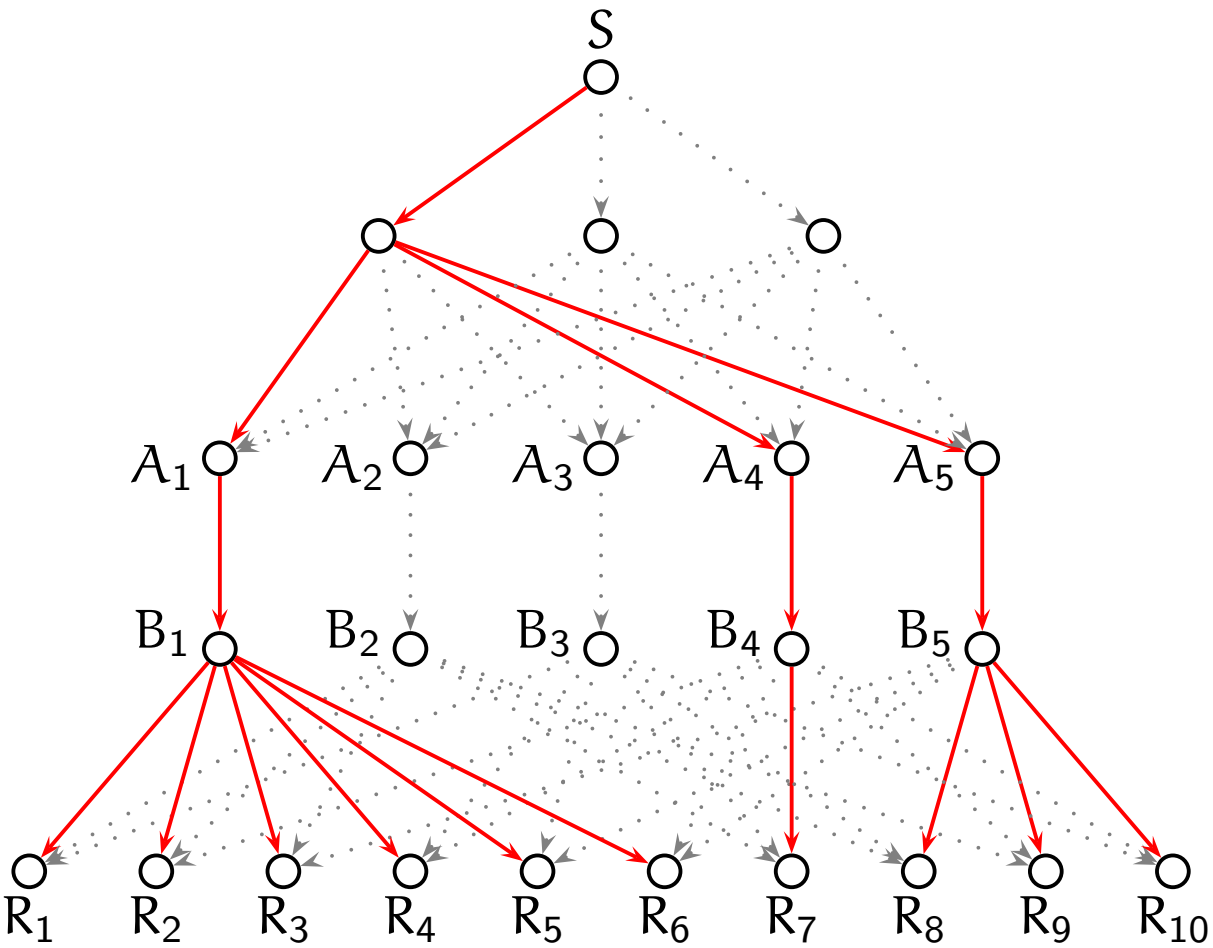
ZK Networks: Average Throughput Benefits



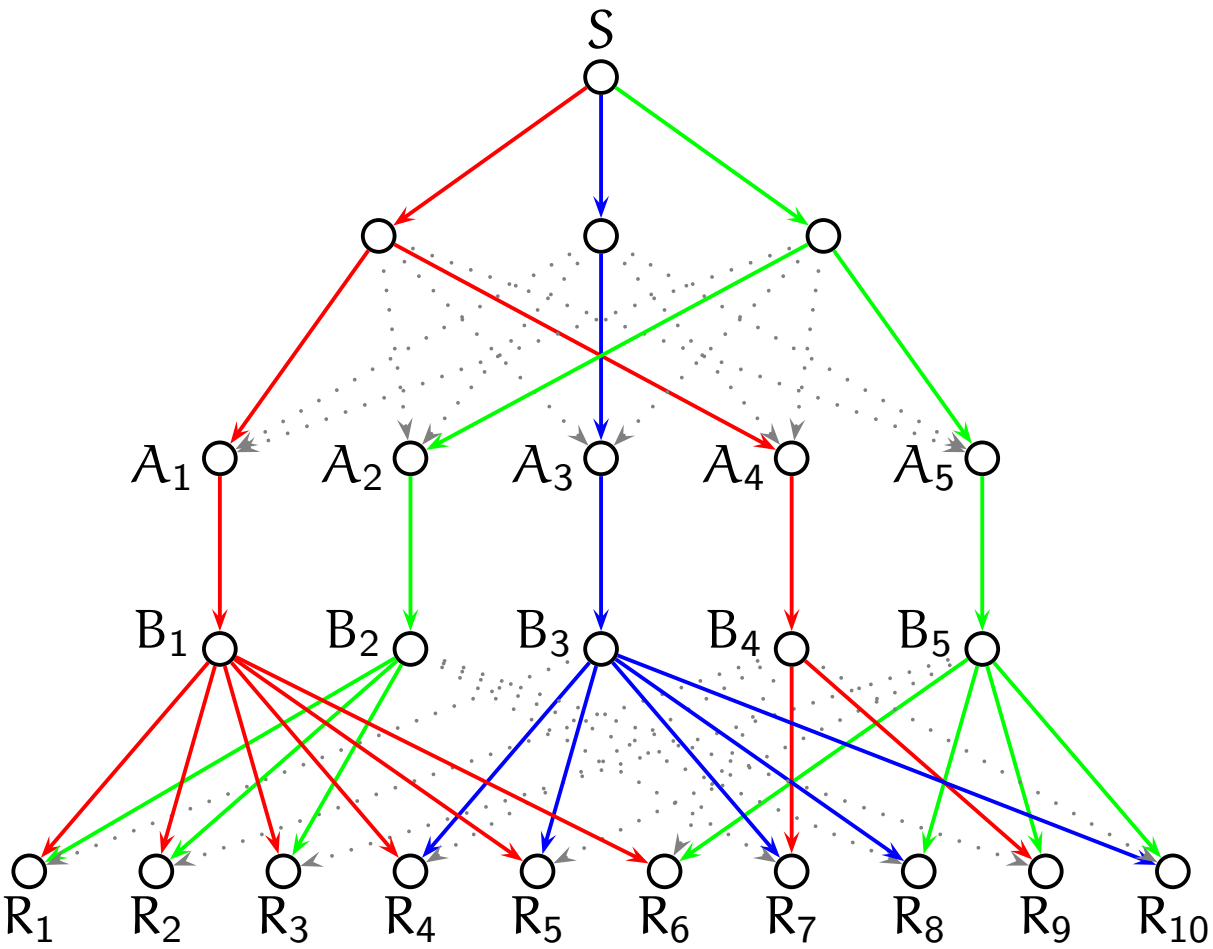
Theorem: We have $\frac{T^r}{T^c} = \frac{p-1}{N-p+1} + \frac{1}{p}$.

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

PACKING STEINER TREES



PACKING PARTIAL STEINER TREES



STEINER TREE PROBLEMS – $\{G, S, \mathcal{R}, w, c\}$

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

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

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

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

$$\beta(G, S, \mathcal{R}) = \max_{c} \frac{T^c(G, c, S, \mathcal{R})}{T^r(G, c, S, \mathcal{R})}$$

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

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

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

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

$$\alpha(G, S, \mathcal{R}) = \max_{w} \frac{\mathrm{OPT}(G, w, S, \mathcal{R})}{\mathrm{LP}(G, w, S, \mathcal{R})}$$

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

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

STEINER TREE PROBLEMS – $\{G, S, \mathcal{R}, w, c\}$

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

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

- Therefore,

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

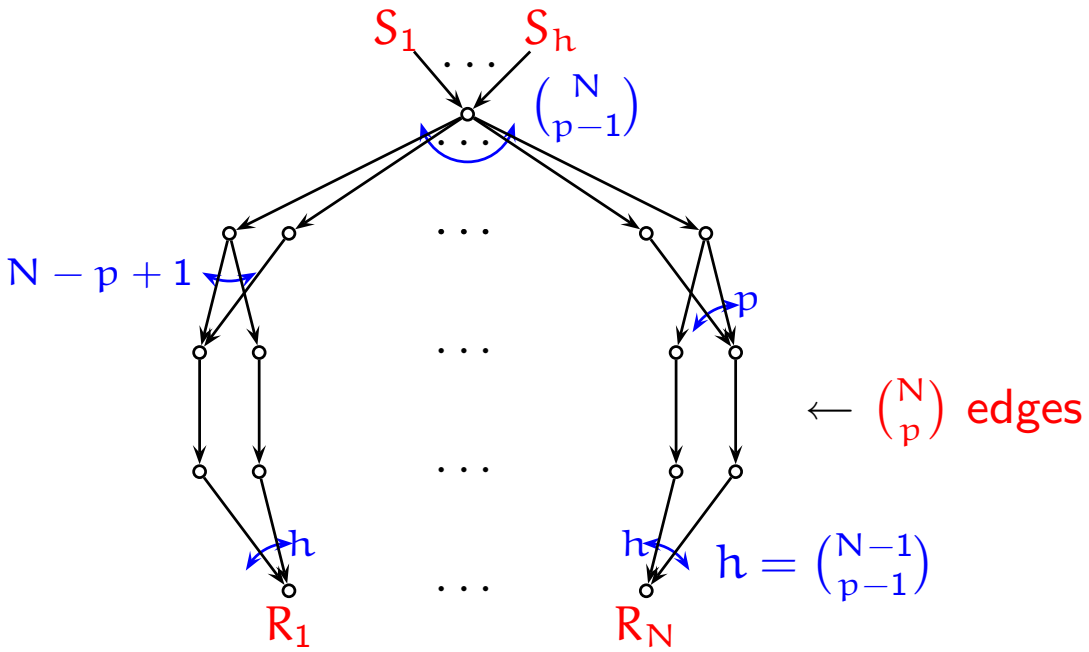
- If $n_1 = N$, we have

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

NETWORKS WITH LIMITED RESOURCES

- Network and user resources include
  - computational capabilities of nodes
  - time, power, and bandwidth
  - number of nodes that are able to code
- How do limited resources affect throughput?
- Are there any tradeoffs?

ZK Networks: Resources For Random Coding



- Nodes randomly combine their inputs.
- The probability that all N receivers can decode is at least

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

CODE-ALPHABET SIZE – Deterministic Coding

- For networks with $\mathfrak{h}$ sources and N receivers, the field of size

$$q = N$$

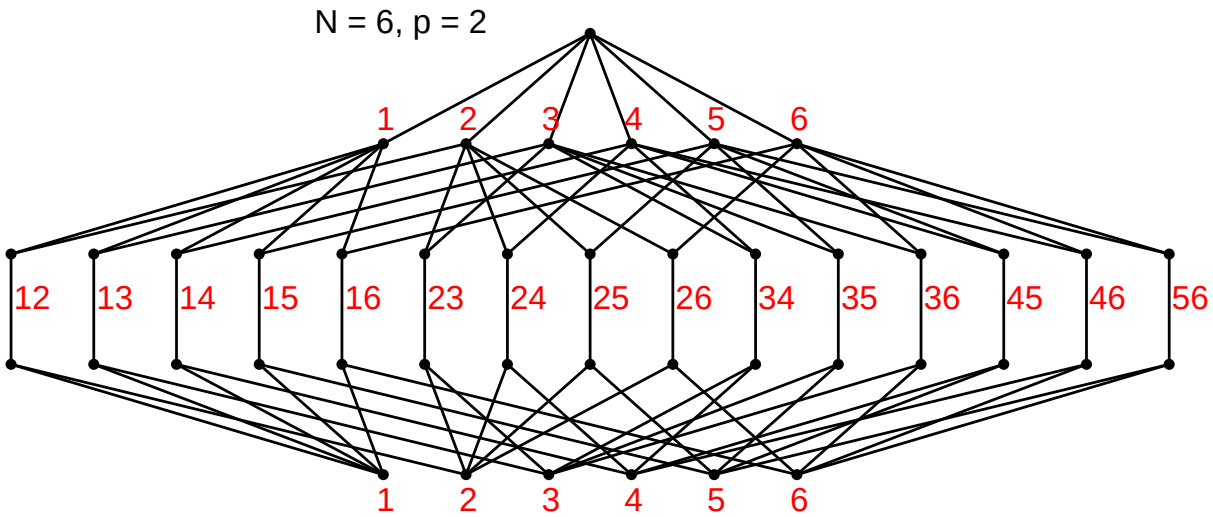
  is sufficient, and $q = O(\sqrt{2N})$ is necessary for some networks.

- For networks with 2 sources and N receivers, the field of size

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

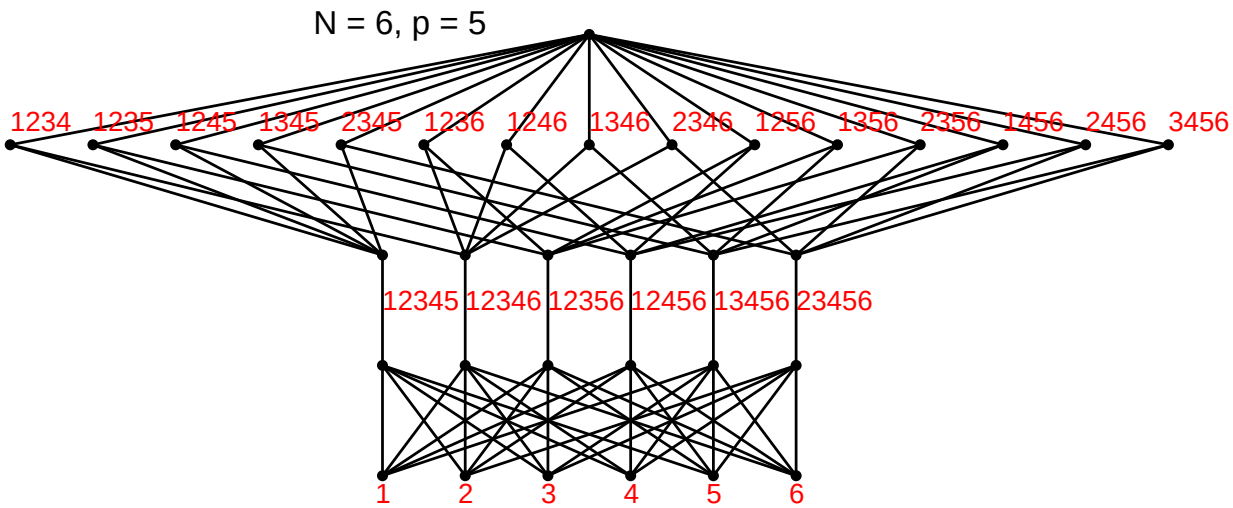
  is sufficient, and, for some networks, necessary.

Deterministic Coding for ZK(N, 2)



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

Deterministic Coding for $ZK(N, N-1)$



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

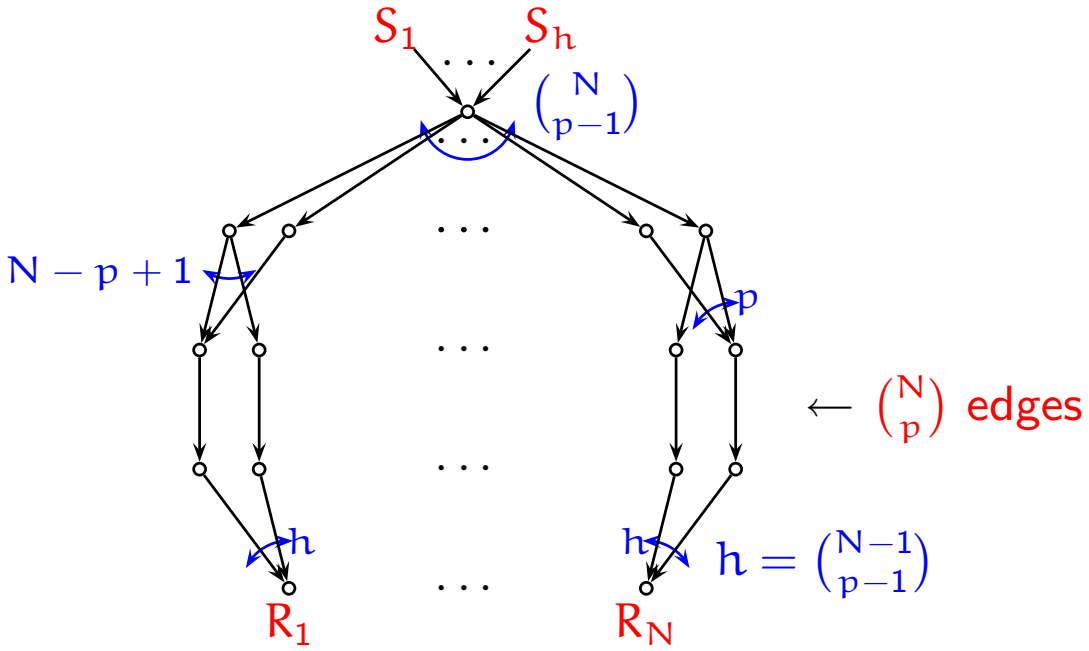
ZK Networks: Deterministic v.s. Random Coding

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

$$\left(1-\frac{N}{q}\right)^{\binom{N}{p}} \approxeq e^{-\frac{N\binom{N}{p}}{q}},$$

requires that $q > N\binom{N}{p}$.

ZK Networks with Fewer Coding Nodes



- Coding is done at $\binom{N-1}{p}$ nodes.
- Suppose that only k nodes can code.
- Theorem: This hybrid scheme achieves the average throughput of

$$T^c \cdot \frac{1}{N}\left(p + \frac{N-p}{p} + k\frac{p-1}{h}\right).$$

SUMMARY

- 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.

Basics

Multicast Model

Main Theorem

Summary

Examples

Benefits

Small Benefits

Coding at the Source

Large Benefits

Steiner Trees

Limited Resources

Summary