NETWORK CODING TUTORIAL

Emina Soljanin

Comm. & Stat. Research, Bell Labs

Murray Hill, June 2007

Theoretical Frameworks For Network Coding

- Information-Theoretic:
  the original proof of the main theorem, cutset bounds.
- Algebraic:
  randomized coding, convolutional coding, galvanized the field.
- Combinatorial:
  deterministic coding, field size, coding with limited resources
- Linear-Programming:
  throughput benefits, routing, networks with weights and costs.

Codes
Randomized Coding
Coding Vectors
LIF Algorithm
Graph Coloring
Determ. Decentralized
Costs
Field Size
Coding Points
Tradeoffs

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.

EDGE-DISJOINT PATHS

- The multicast condition:

  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 $1 \leqslant i \leqslant h$.

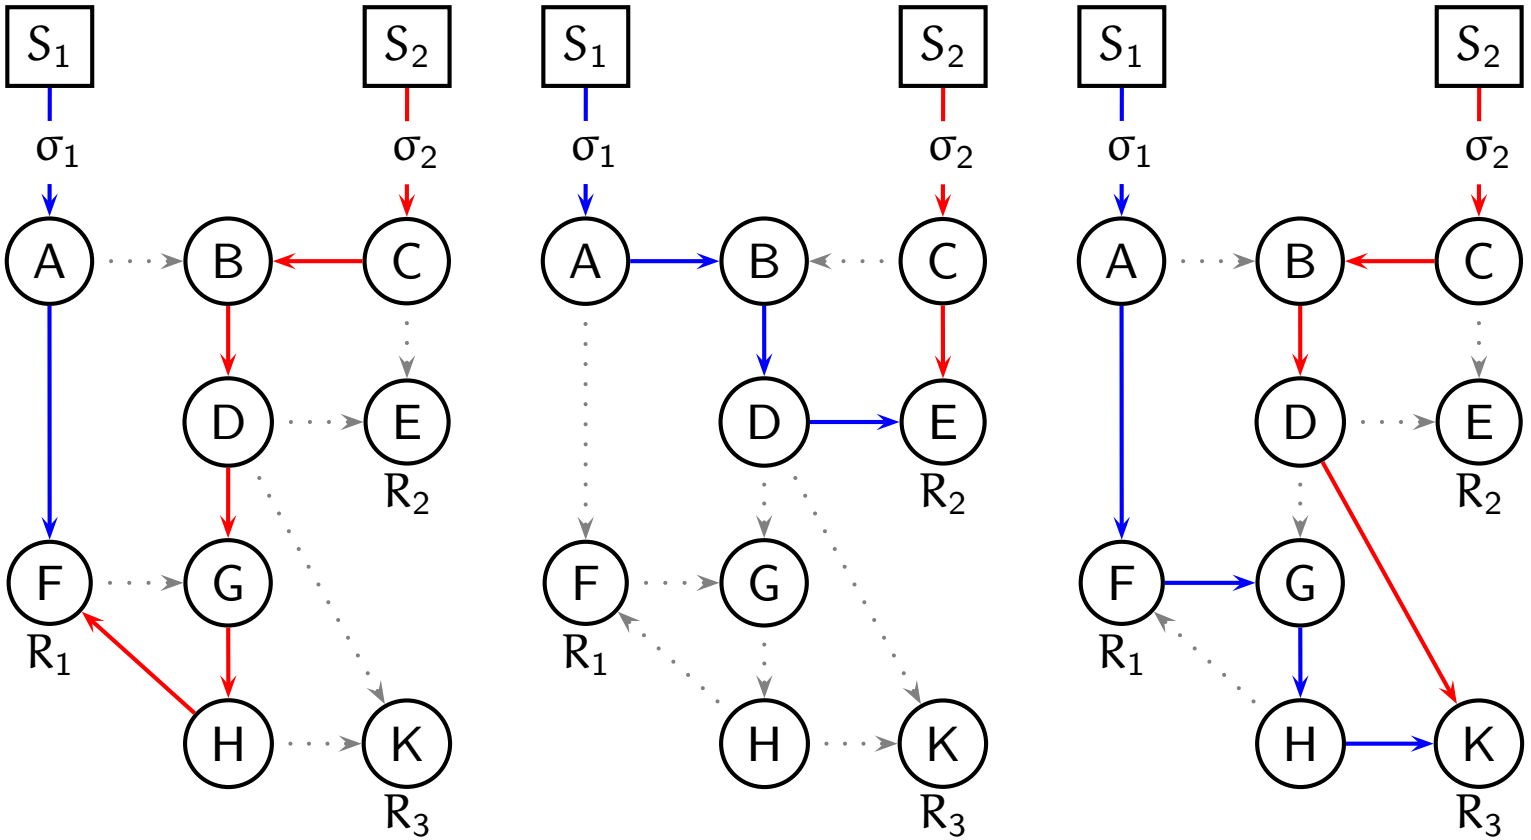
- All code design algorithms start with identifying such paths.
- The rest of the network is irrelevant.
- Coding points are edges where different paths merge.

NETWORK MULTICAST – Linear Combining

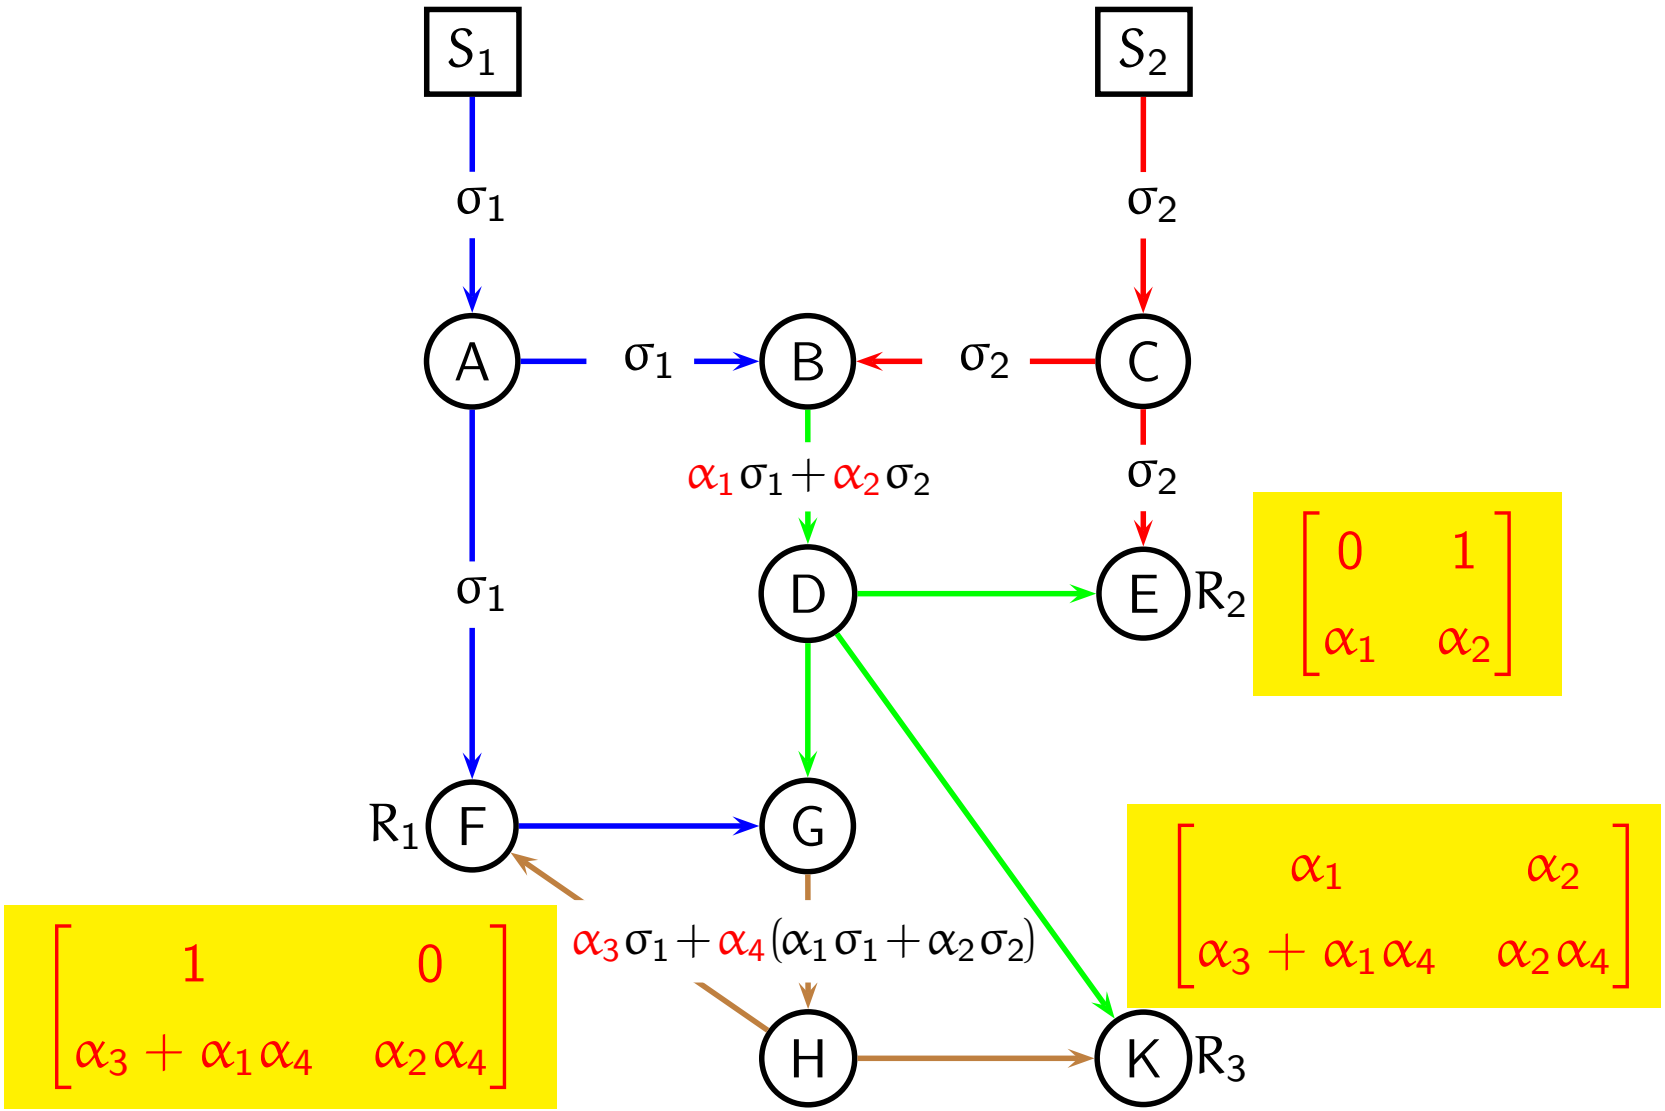
- Source S_i emits σ_i which is an element of some finite field.
- Edges carry linear combinations of their parent node inputs.
- Consequently,
  edges carry linear combinations of source symbols σ_i.
- Network Multicast Problem:

  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 $\mathbb{F}_2$

RANDOMIZED CODING

Theorem (related to the Schwartz-Zippel 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. $q > d$

If the values for $\alpha_1, \ldots, \alpha_\eta$ are chosen uniformly at random $\mathbb{F}_q$, then

$$\Pr\{f(\alpha_1, \ldots, \alpha_\eta) = 0\} \leqslant 1 - (1 - d/q)^\eta$$

RANDOMIZED CODING

We prove the theorem by induction in the number of variables η.

1. For $\eta = 1$,
   - f is a polynomial in a single variable of degree at most d.
   - An element of $\mathbb{F}_q$ is a root of f with probability of at most

$$d/q = 1 - (1 - d/q)^1.$$

2. For $\eta > 1$,
   - The claim holds for polynomials with fewer than η variables.
   - We express f as

$$f(\alpha_1, \ldots, \alpha_\eta) = \alpha_\eta^{d_1} f_1(\alpha_1, \ldots, \alpha_{\eta-1}) + f_2(\alpha_1, \ldots, \alpha_\eta),$$

   where $d_1 \leqslant d$ and f_1 is a not-identically zero polynomial.

RANDOMIZED CODING

- Consider $f(\alpha_1, \ldots, \alpha_n) = \alpha_n^{d_1} f_1(\alpha_1, \ldots, \alpha_{n-1}) + f_2(\alpha_1, \ldots, \alpha_n)$
- We have

$$\Pr[f = 0] = \Pr[f_1 = 0] \cdot \Pr[f = 0 | f_1 = 0] + \Pr[f_1 \neq 0] \cdot \Pr[f = 0 | f_1 \neq 0].$$

- Furthermore,
  1. $\Pr[f_1 = 0] \leqslant 1 - (1 - d/q)^{(n-1)}$ by the inductive hypothesis.
  2. $\Pr[f = 0 | f_1 = 0] \leqslant 1$.
  3. $\Pr[f = 0 | f_1 \neq 0] \leqslant d/q$, as a polynomial in α_n
- Therefore,

$$\begin{aligned}
\Pr[f = 0] &\leqslant \Pr[f_1 = 0] + (1 - \Pr[f_1 = 0])\frac{d}{q} && \text{by 2. and 3.} \\
&= \Pr[f_1 = 0]\left(1 - \frac{d}{q}\right) + \frac{d}{q} \\
&\leqslant \left[1 - \left(1 - \frac{d}{q}\right)^{(n-1)}\right]\left(1 - \frac{d}{q}\right) + \frac{d}{q} && \text{by 1.} \\
&= 1 - \left(1 - \frac{d}{q}\right)^{n}.
\end{aligned}$$

Codes
Randomized Coding
Coding Vectors
LIF Algorithm
Graph Coloring
Determ. Decentralized
Costs
Field Size
Coding Points
Tradeoffs

RANDOMIZED CODING

- Recall that

$$f(\alpha_1, \ldots, \alpha_\eta) \triangleq \det(\mathbf{C}_1) \cdots \det(\mathbf{C}_N).$$

- What is d?
- Because of the multicast condition, we have $d \leqslant N$.
- What is η?
- $\eta \leqslant h \cdot n_c$, where n_c is the number of coding points.

GLOBAL CODING VECTORS

- Edges carry linear combinations of their parent node inputs.
- $\{\alpha_k\}$ are the local coding coefficients.
- Each edge e carries a linear combination of source symbols:

$$c_1(e)\sigma_1 + \cdots + c_h(e)\sigma_h = \begin{bmatrix} c_1(e) \dots c_h(e) \end{bmatrix} \begin{bmatrix} \sigma_1 \\ \vdots \\ \sigma_h \end{bmatrix}$$

- $[c_1(e) \dots c_h(e)] \in \mathbb{F}_q^h$ is the global coding vector of edge e.

DECODING FOR RECEIVER j

- ρ_i^j is the symbol on the last edge on the path (S_i, R_j).
- c_i^j is the coding vector of the last edge on the path (S_i, R_j).
- $\mathbf{C}_j$ is the matrix whose i-th row is c_i^j.
- Receiver j has to solve the following system of equations:

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

NETWORK MULTICAST – Code Design

Select a coding vector for each edge e of the network so that

1. the matrices $\mathbf{C}_1 \ldots \mathbf{C}_N$ are full rank.
2. the coding vector of e is in the linear span of the coding vectors of the input edges to the parent node of e.

LINEAR INFORMATION FLOW ALGORITHM

- Initially,
  - for $1 \leqslant j \leqslant N$, find h edge-disjoint paths (S_i, R_j), $1 \leqslant i \leqslant h$.
  - find coding points (edges where the paths merge).
  - assign coding vectors to all edges with no upstream coding points. How is this done?
  - Set $\mathbf{C}_j^{(0)} = I$ for $1 \leqslant j \leqslant N$.
- At each stage $k > 1$ of the algorithm,
  - coding vectors of certain coding points are determined.How?
  - Q: Which ones? A: Those with coding vectors of all input edges already determined.
  - Corresponding rows in $\mathbf{C}_j^{(k)}$ are updated. How?
  - For which j? Which rows?

LINEAR INFORMATION FLOW ALGORITHM

- The i-th row of $\mathbf{C}_j^{(k)}$ represents the global coding vector of the edge of (S_i, R_j) processed at stage k.
- Note that the i-th raw of $\mathbf{C}_j^{(k)}$ will be identical to, say, l-th raw of some $\mathbf{C}_m^{(k)}$.
- The coding vectors determined at stage k ensure that matrices $\mathbf{C}_j^{(k)}$, $1 \leqslant j \leqslant N$, are regular. How?
- Algorithms for finding such vectors? Field size? Scalability?

LINEAR INFORMATION FLOW ALGORITHM

- Let e be a coding point with m inputs.
- There are $q^m - 1$ feasible coding vectors for e in $\mathbb{F}_q^h$.
- A receiver using a path containing e eliminates up to $q^{m-1} - 1$ vectors. Q: How many paths can intersect on e?
- A: At most N. Therefore, there are at least

$$q^m - Nq^{m-1}$$

  valid vectors for e.
- The field of size $q > N$ is sufficient.

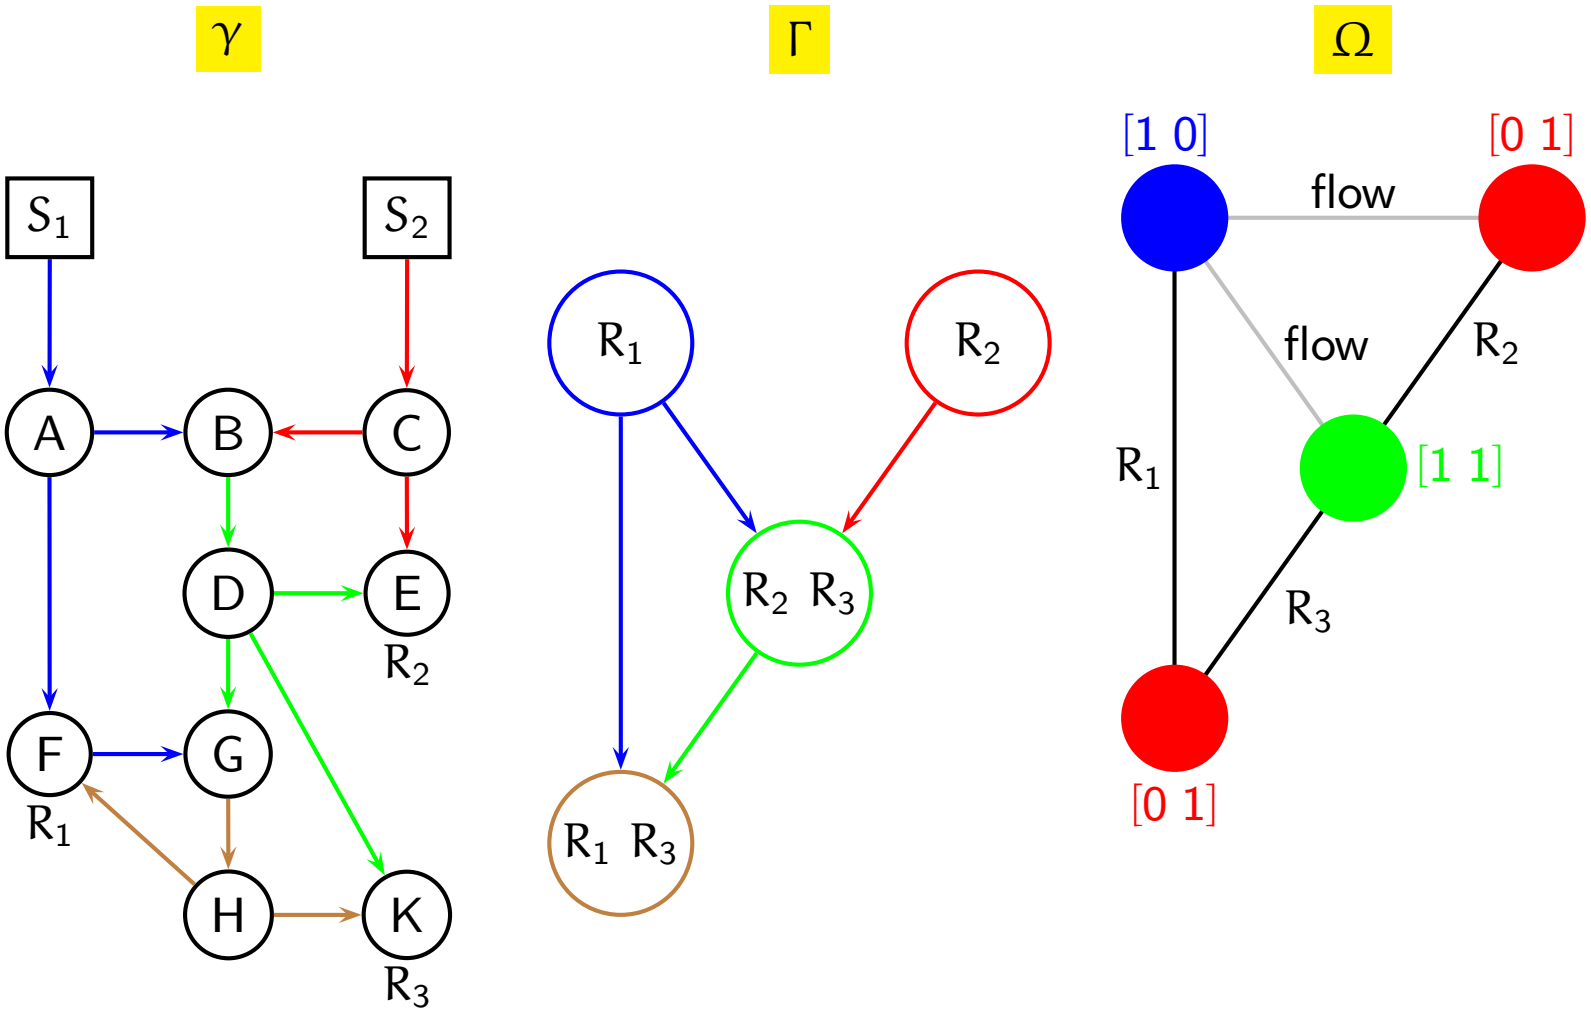
Coding Vectors for Networks with Two Sources

- For edge e, coding vector $[c_1(e)\, c_2(e)]$ is in $\mathbb{F}_q^2$.
- Let $\mathcal{A}$ be the following set of $(q+1)$ vectors:

$$[0\,1],\ [1\,0],\ \text{and } [1\,\alpha^i] \text{ for } 0 \leqslant i \leqslant q-2,$$

  where α is a primitive element of $\mathbb{F}_q$.
- Consider any two different vectors in $\mathcal{A}$:
  - they are linearly independent, and
  - any vector in $\mathcal{A}$ is in their linear span.

VERTEX COLORING AND CODE DESIGN



Deterministic Decentralized Algorithms

- Networks with two sources and N receivers:

  To each coding point, assign a different vector from the set

  $$[0\,1],\ [1\,0],\ \text{and}\ [1\,\alpha^i]\ \text{for}\ 0 \leqslant i \leqslant q-2,$$

  as its global coding vector.
- For general networks, use RS codes.
- Scalability: such codes are valid as long as Γ stays the same.
- The price to pay is a larger field size.

ENCODING/DECODING COMPLEXITY

If the operating field is $\mathbb{F}_{2^m}$, we need

- $\mathcal{O}(h^2)$ operations in $\mathbb{F}_{2^m}$ to form a linear combination.
- $\mathcal{O}(h^3)$ operations in $\mathbb{F}_{2^m}$ for a matrix inversion by Gaussian elimination.
- $\mathcal{O}(m^2)$ binary operations for a multiplication or inversion.

CODE-ALPHABET SIZE – Deterministic Coding

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

$$\mathrm{q} = \mathrm{N}$$

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

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

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

  is sufficient, and, for some networks, necessary.

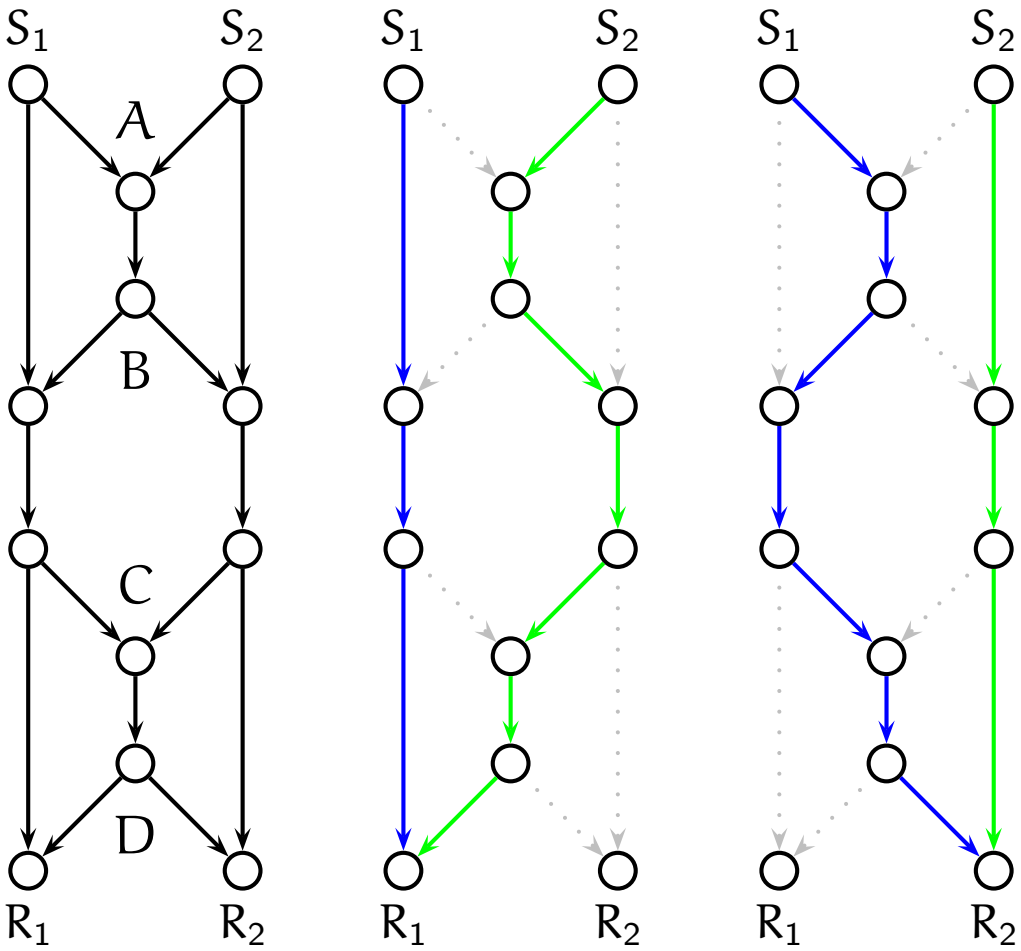
Field Size For Networks With Two Sources

Elements of the Proof:

- Lemma: Every vertex in an Ω has degree at least two.
- Lemma: Every k-chromatic graph has at least k vertices of degree at least $k-1$.
- For an Ω with n nodes and chromatic number k:
  1. $E(\Omega) \geqslant [k(k-1) + (n-k)2]/2$,
  2. $E(\Omega) \leqslant N + n - 2$.
- Therefore $N \geqslant \frac{k(k-1)}{2} - k + 2$, and thus

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

MINIMAL CONFIGURATIONS – Example



MINIMAL CONFIGURATIONS

- A configuration is minimal with the multicast property if removing any edge violates the multicast condition.
- A minimal configuration can be found in polynomial time.
- Minimal configurations have certain structural properties.
- A minimal configuration may have more than the minimum number of coding points possible for a given network.
- Finding a configuration with the minimum number of coding points is in general NP hard.
- Even deciding whether coding is necessary is in general NP.

MINIMAL CONFIGURATIONS – Properties

- A network with two sources and two receivers
  - either does not require coding or
  - is equivalent to the butterfly network.
- In networks with N receivers, the number of coding points
  - $\mathrm{n_c} \leqslant \mathrm{N}$, for networks with two sources, and
  - $\mathrm{n_c} \geqslant \Omega(\mathrm{h}^2\mathrm{N})$ for some networks with $\mathrm{h} > 2$ sources.

Tradeoffs in Networks with Limited Resources

- design complexity v.s. the alphabet size
- throughput v.s. the number of coding points
- min-cut v.s. the alphabet size
- throughput v.s. the alphabet size