Introduction to Network Coding

Christina Fragouli

School of Computer and Communication Sciences

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Introduction to Network Coding

Monograph:

Network Coding: Theory and Applications

C. Fragouli and E. Soljanin

Foundations and Trends in Networking

2007-2008

http://arni.epfl.ch

Introduction to Network Coding

General Context:

How we treat information when we want to communicate over a network.

Traditionally information treated as fluid through pipes.

Network Coding

Ahlswede,Cai,Li,Yeung 2000

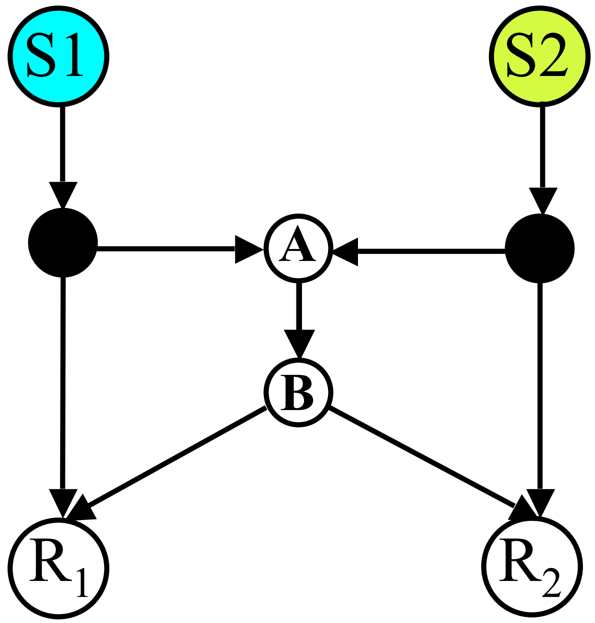

Receiver 1 *Receiver 2*

Network Coding

Ahlswede,Cai,Li,Yeung 2000

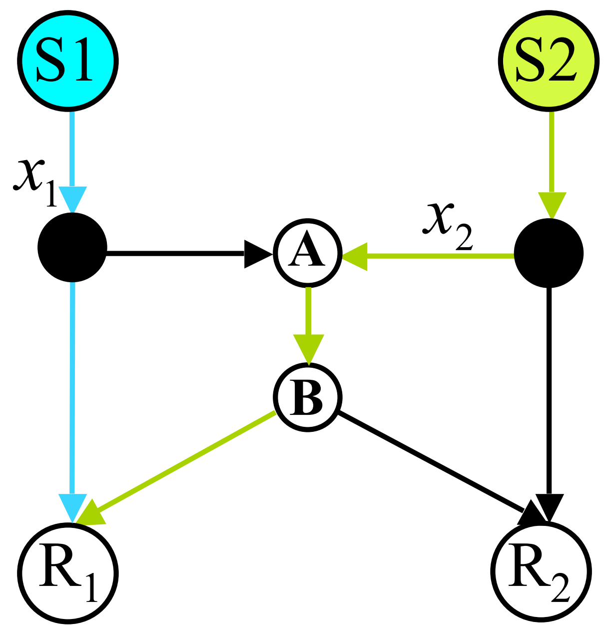

Receiver 1 *Receiver 2*

Network Coding

Ahlswede,Cai,Li,Yeung 2000

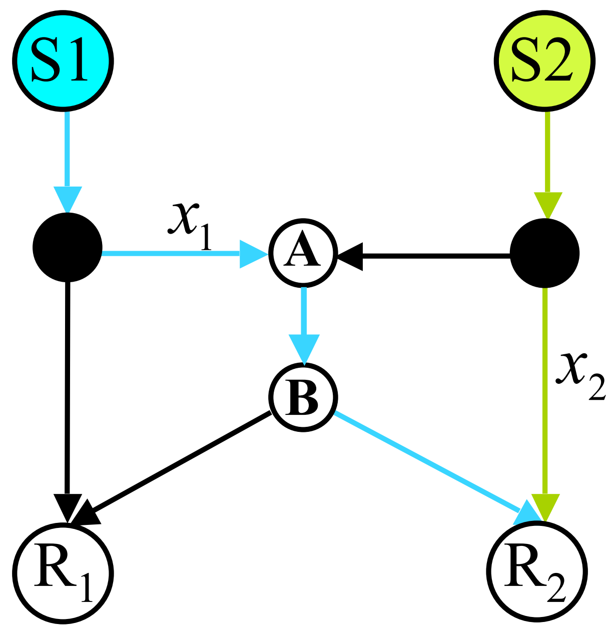

Receiver 1

Receiver 2

Network Coding

Ahlswede,Cai,Li,Yeung 2000

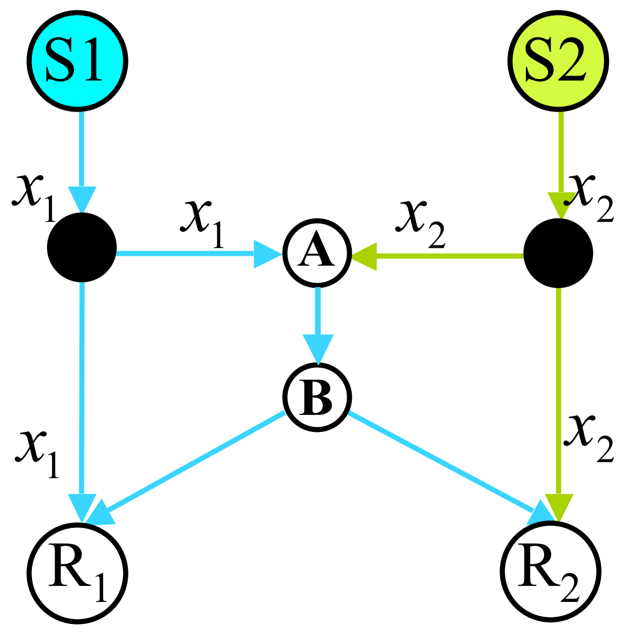

Receiver 1 *Receiver 2*

Network Coding

Ahlswede,Cai,Li,Yeung 2000

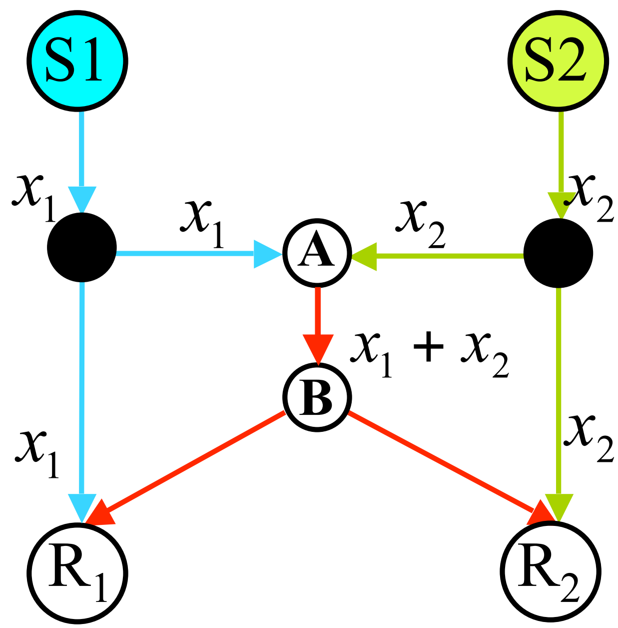

Receiver 1 *Receiver 2*

Network Coding

Ahlswede,Cai,Li,Yeung 2000

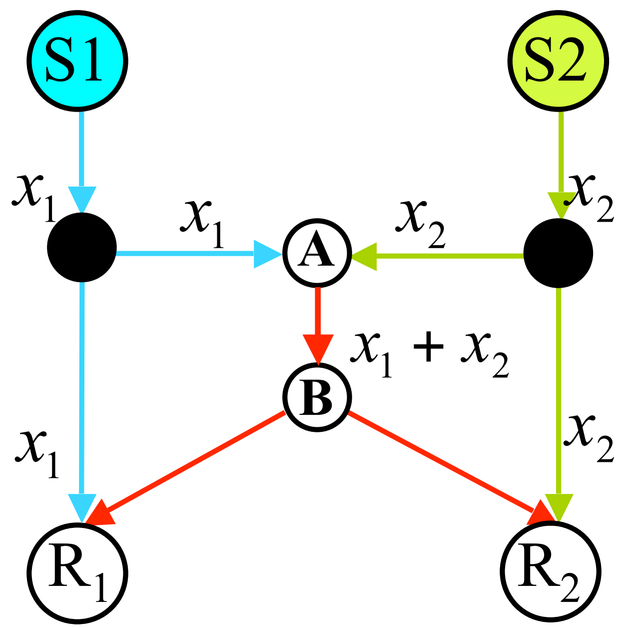

$\{x_1, x_1 + x_2\}$

Receiver 1

$\{x_2, x_1 + x_2\}$

Receiver 2

New idea

Ahlswede,Cai,Li,Yeung 2000

Routing

Nodes in the network are only allowed to forward the incoming information flows

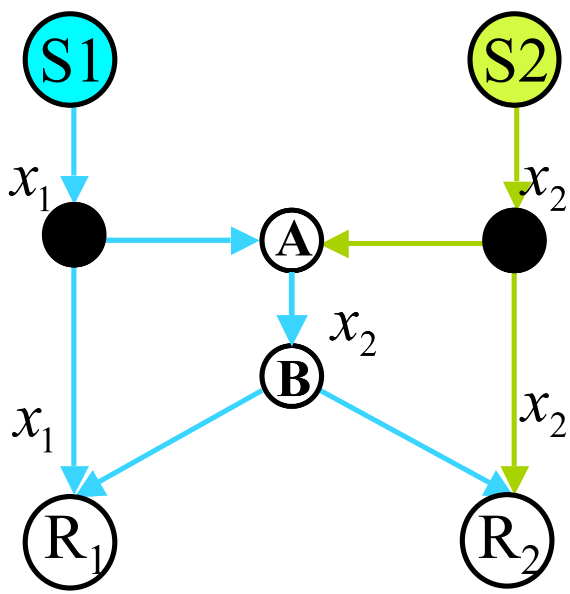

Network Coding

Nodes in the network are allowed to process the incoming information flows

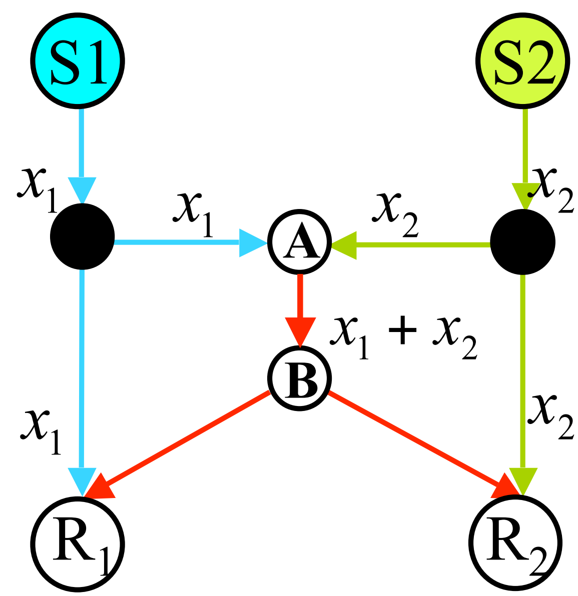

Network Coding Research

- To explore and develop new fundamental approaches in information flow through networks.
- To have impact on applications: make new generation networks more efficient, reliable, and secure.

Wireless Resources

Introductory Example

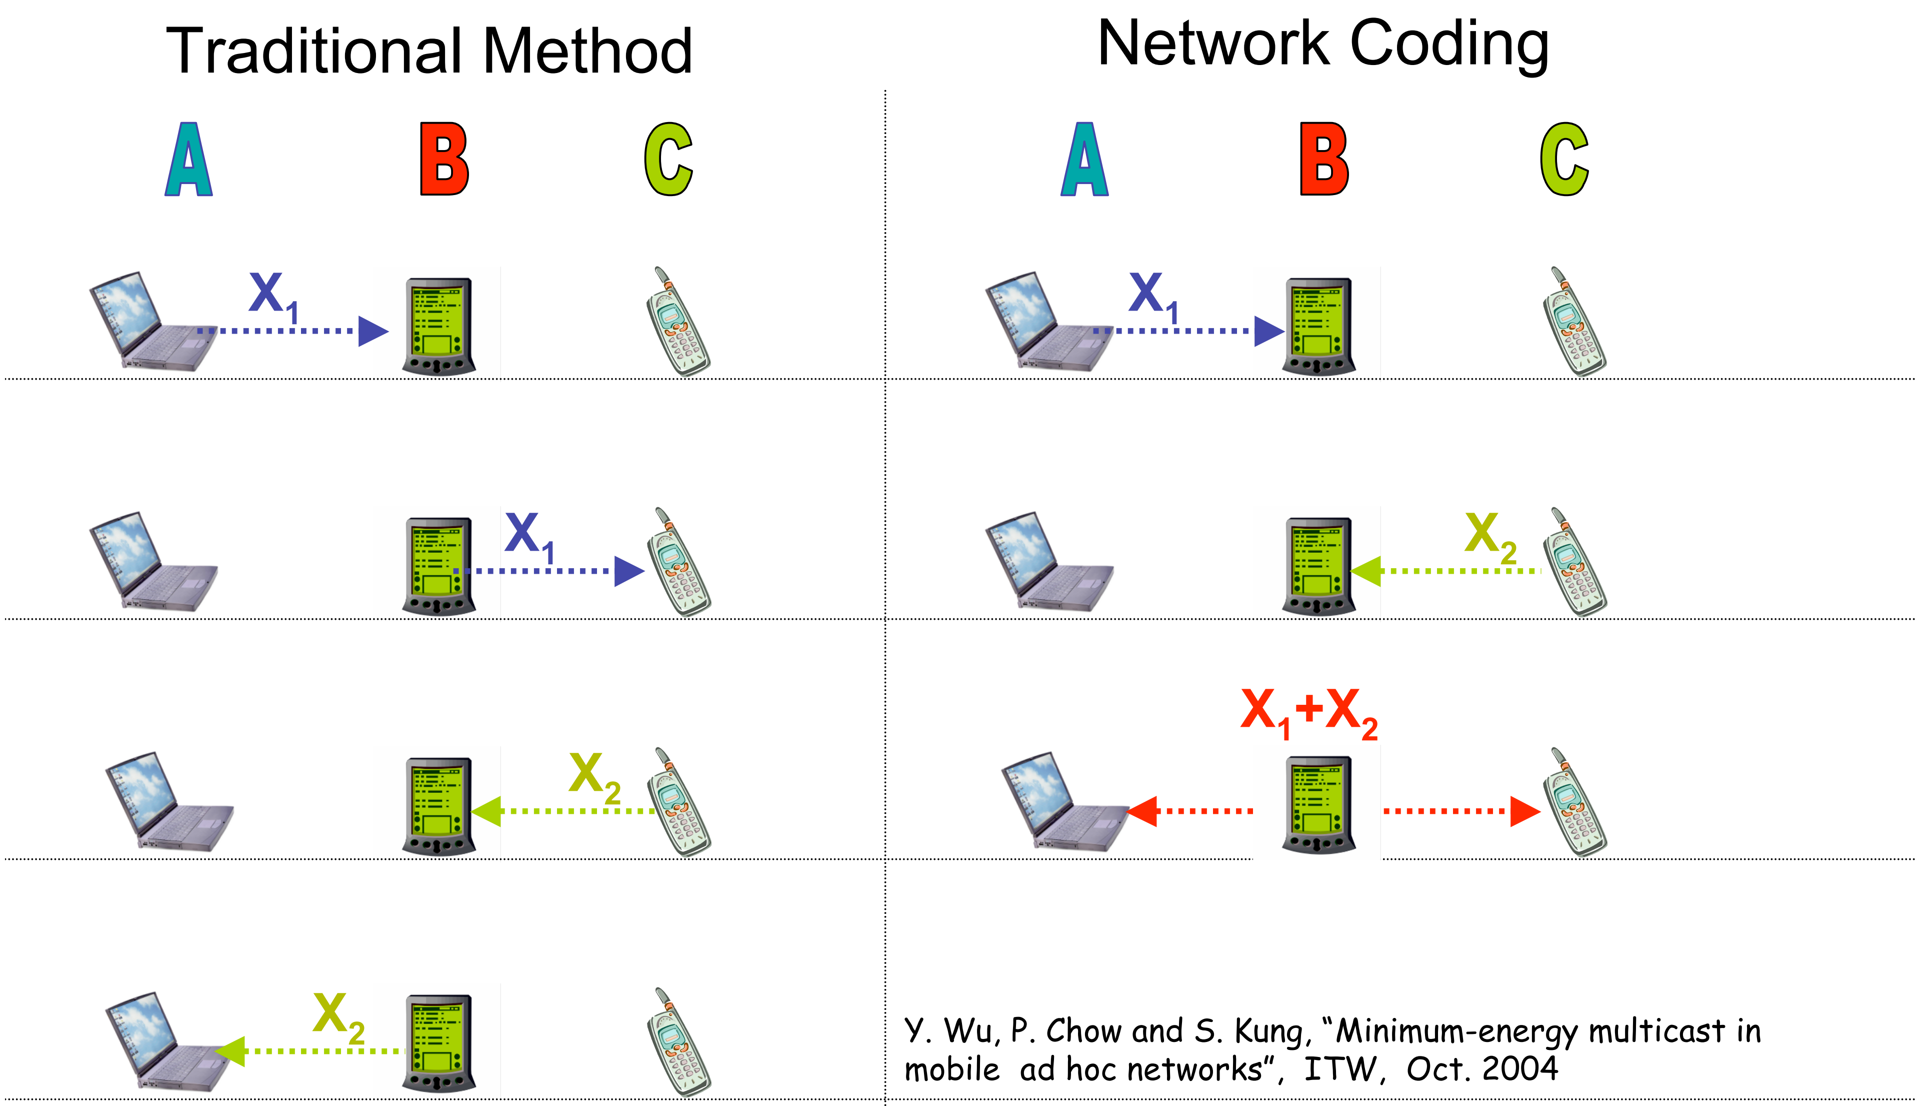

Y. Wu, P. Chow and S. Kung, "Minimum-energy multicast in mobile ad hoc networks", ITW, Oct. 2004

Security

Introductory Example

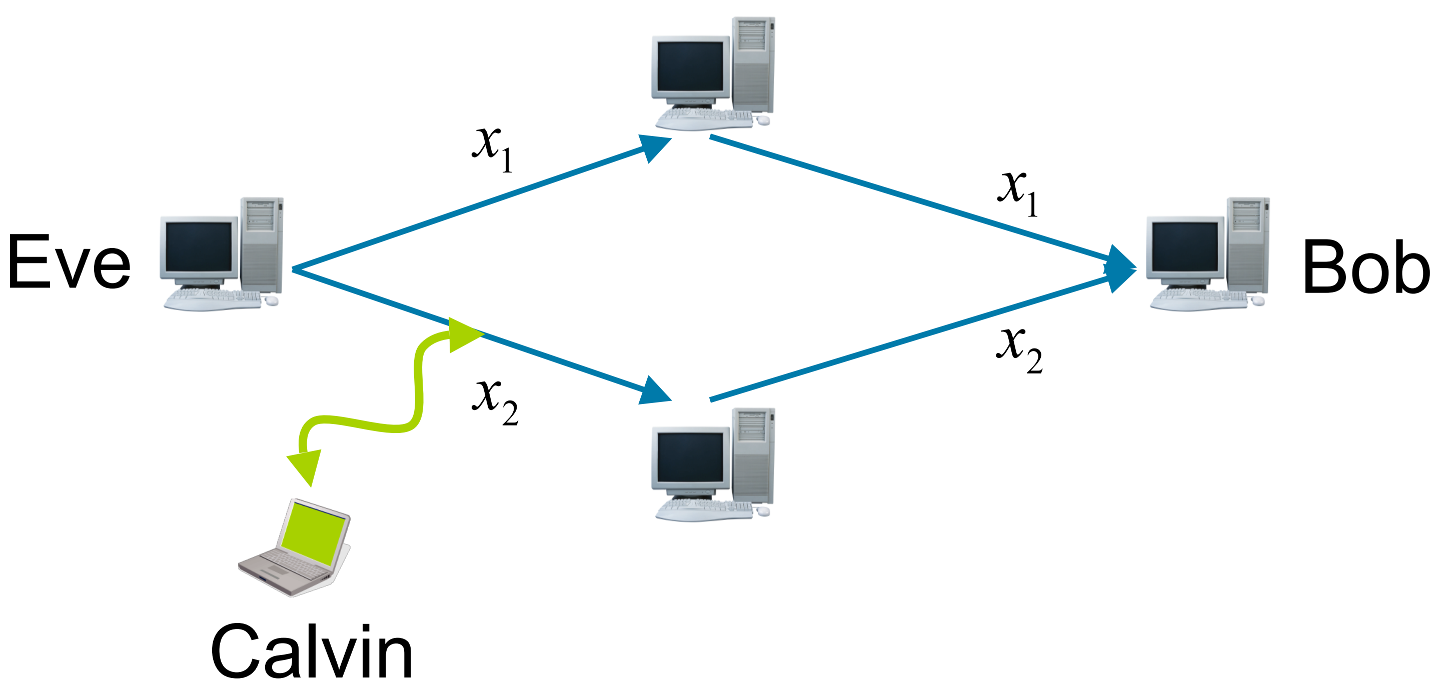

Security

Introductory Example

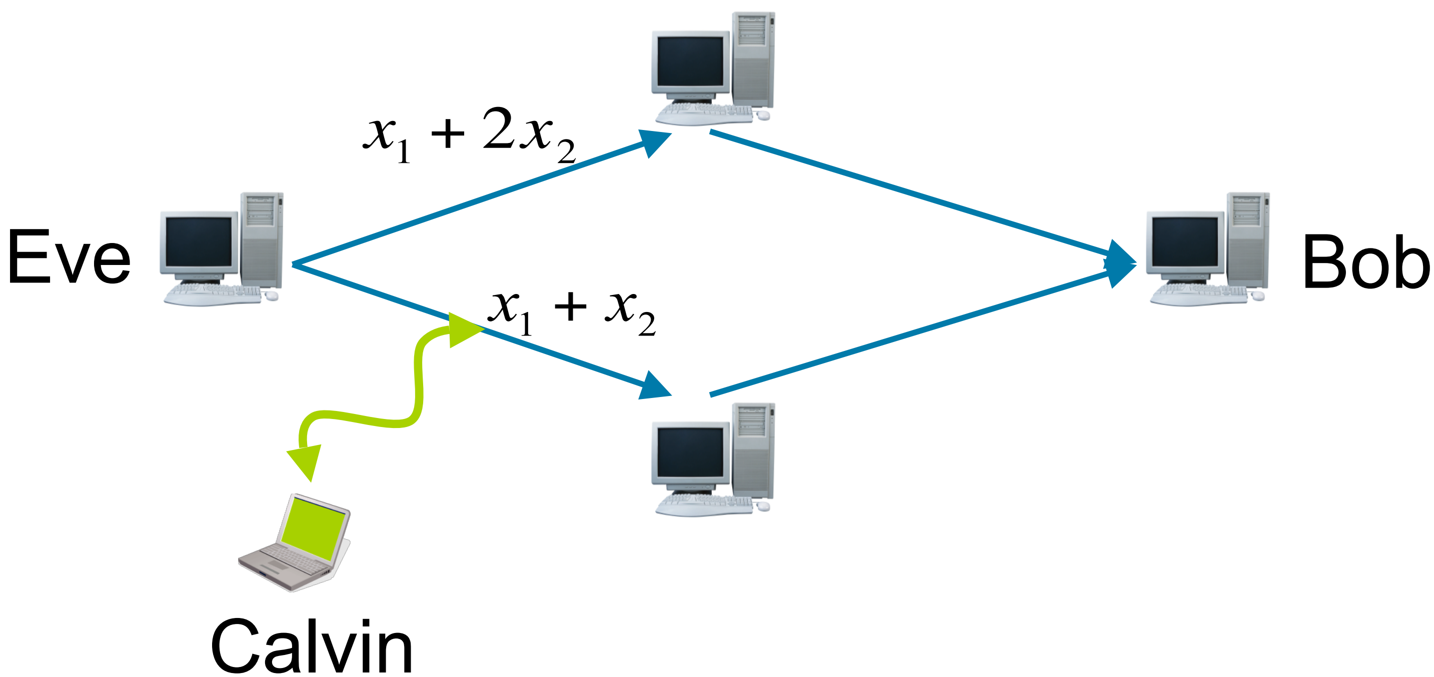

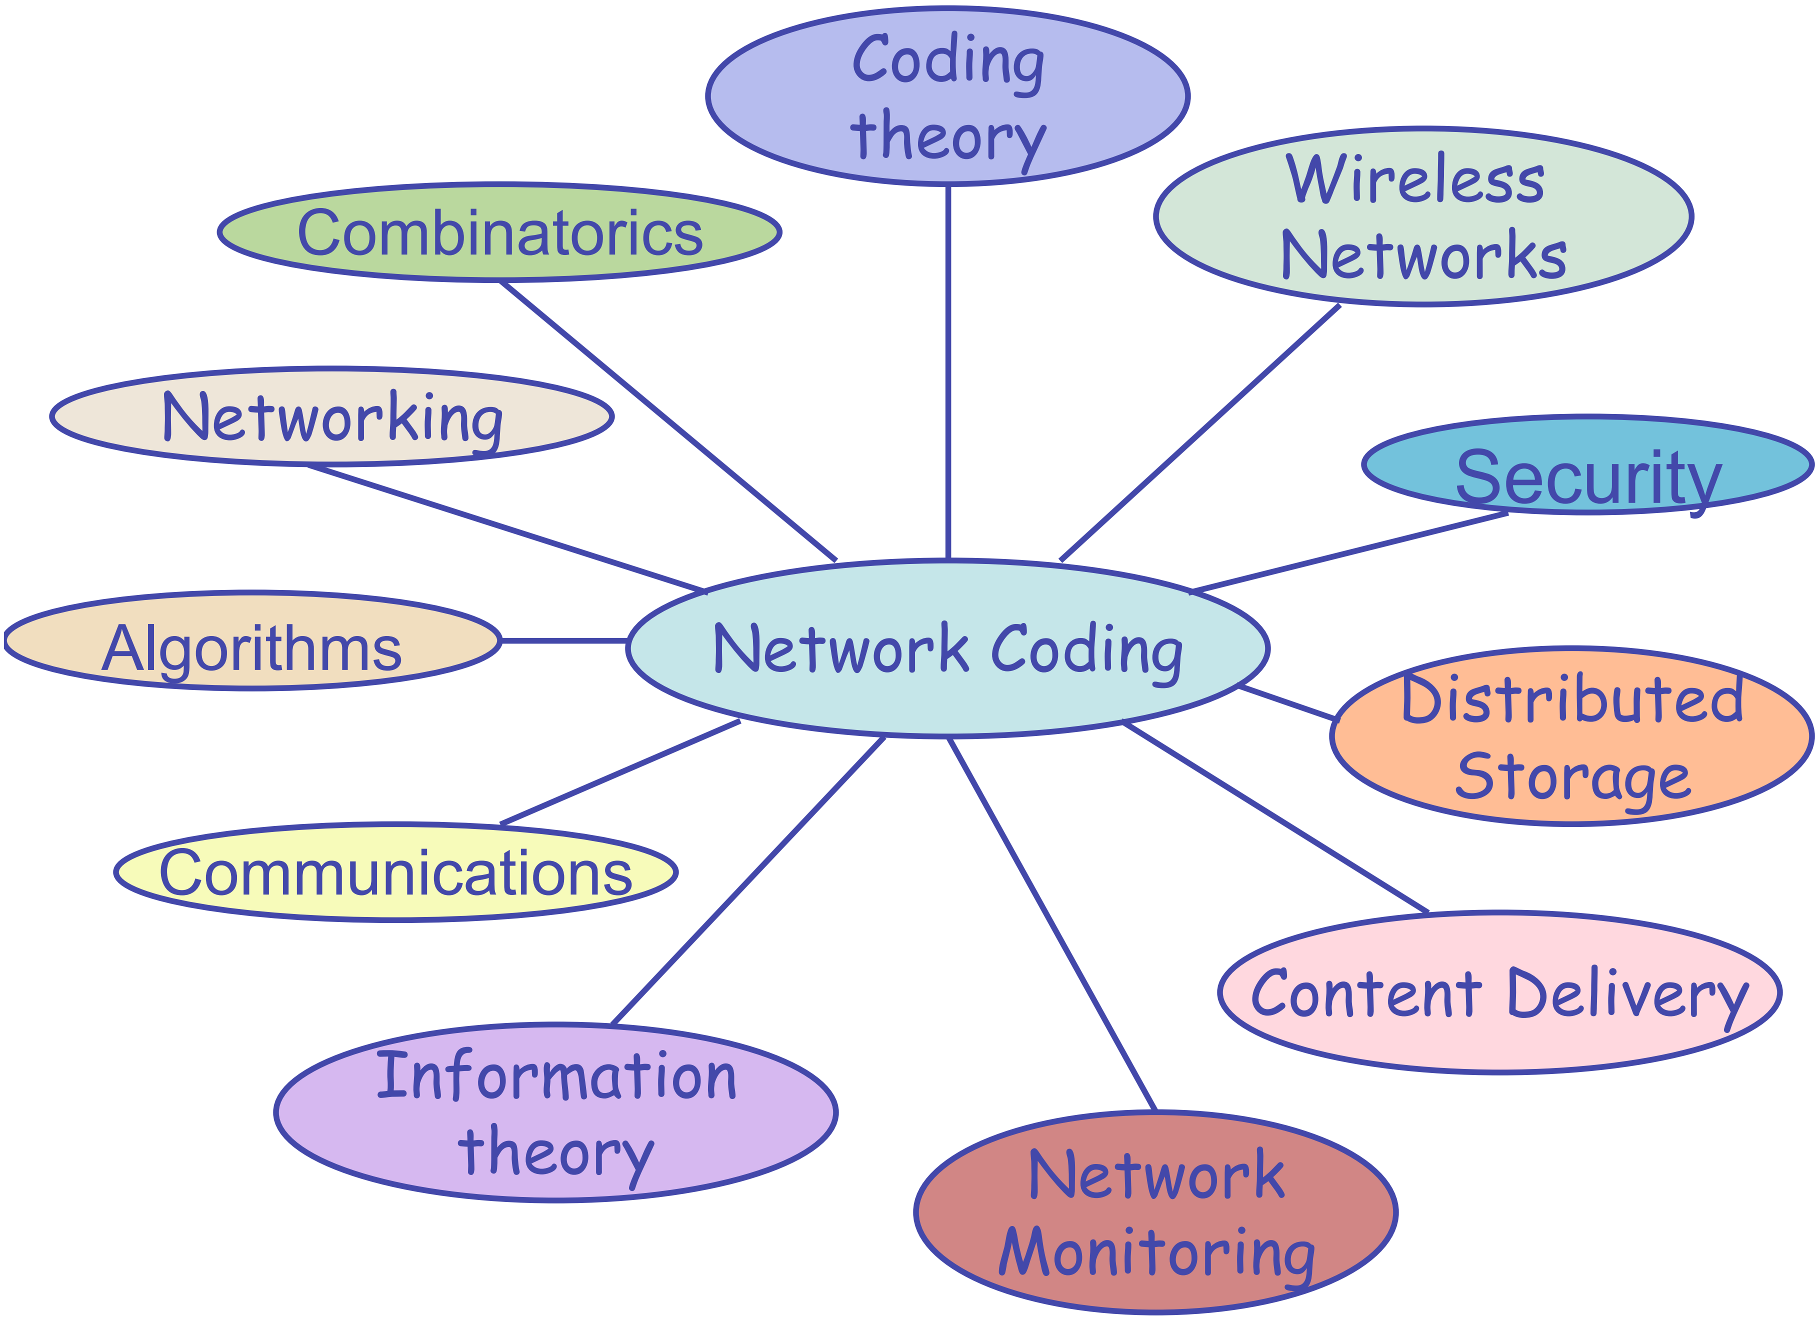
Coding theory
Combinatorics
Wireless Networks
Networking
Security
Algorithms
Network Coding
Distributed Storage
Communications
Content Delivery
Information theory
Network Monitoring

Outline

1. Main Theorem in Multicasting

2. Benefits and Requirements

3. Network Code Design

4. Applications

Outline

1. Main Theorem in Multicasting

- Min-cut Max-flow Theorem
- Statement of the Main Theorem
- Proof Using the Algebraic Framework
- Discussion on Theorem Assumptions

Min-Cut Max-Flow Theorem

[Ford, Fulkerson] ~1950

- Consider a network represented as a directed acyclic graph G=(V,E) with unit-capacity edges.
- Assume a source node S wants to transmit information to a receiver node R.

- If the min-cut between S and R equals h, then information can be send from S to R at a maximum rate of h.

Equivalently, there exist h edge-disjoint paths from the source R to the receiver j.

Min-Cut Max-Flow Theorem

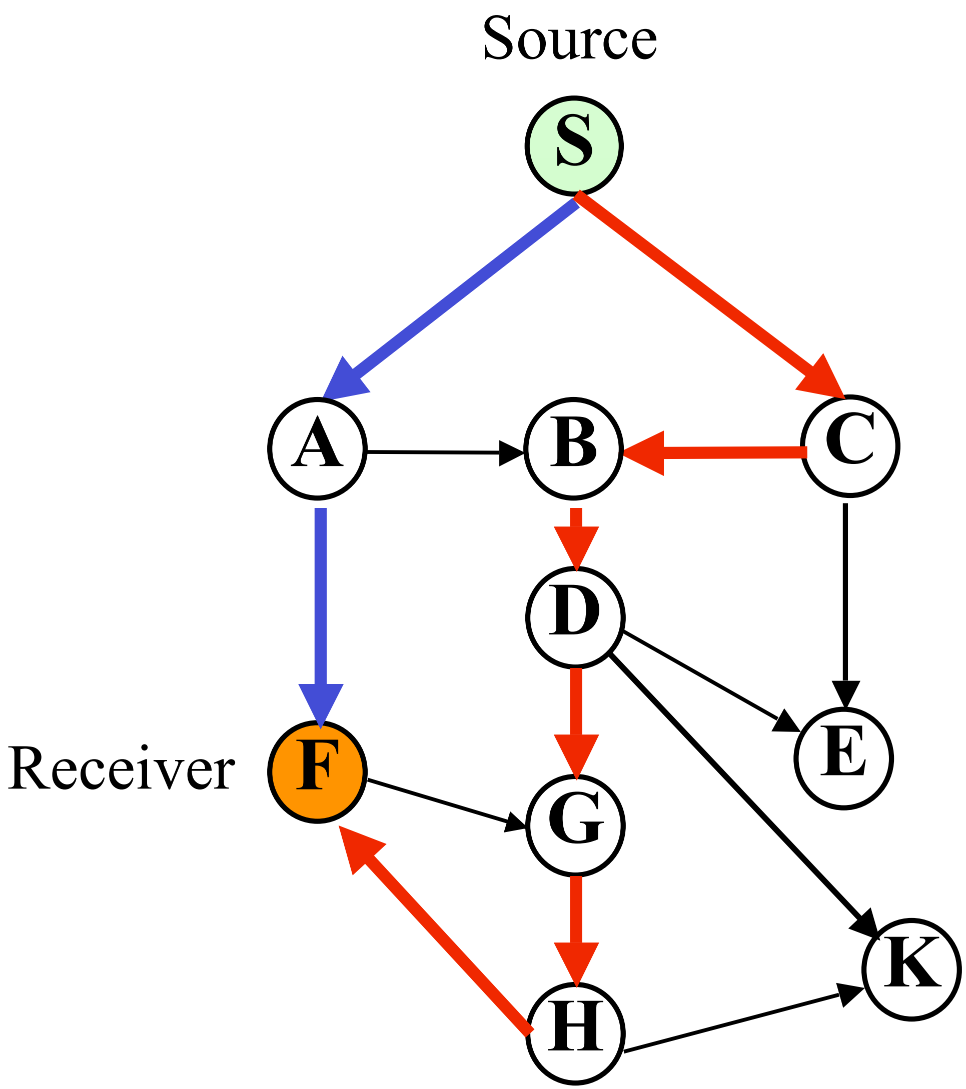

Network Multicast

- A network is represented as a directed acyclic graph with unit-capacity edges.
- There are h unit-rate information sources $S_1, \ldots, S_h$ and N receivers $R_1,\ldots,R_N$

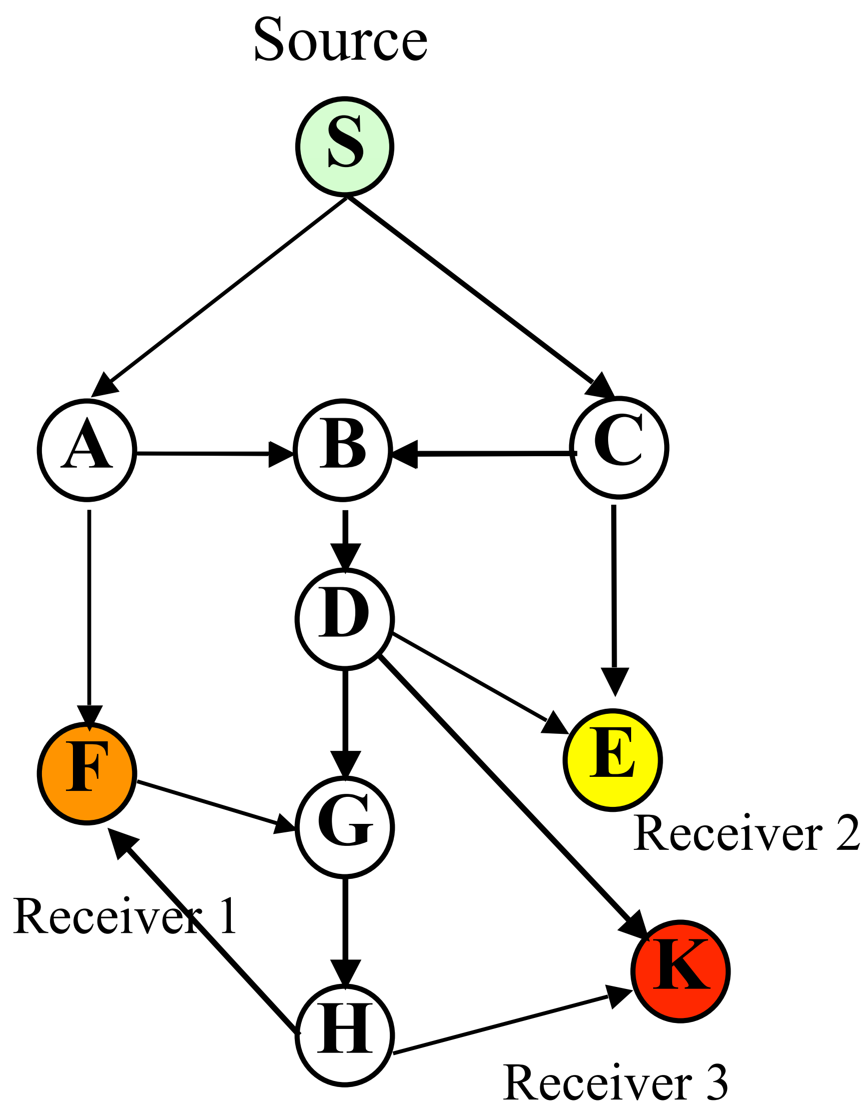
Source
S
A
B
C
D
E
F
G
H
K
Receiver 1
Receiver 2
Receiver 3

Network Multicast

- A network is represented as a directed acyclic graph with unit-capacity edges.
- There are h unit-rate information sources $S_1, \ldots, S_h$ and N receivers $R_1,\ldots,R_N$

- Max-flow min-cut theorem [Ford, Fulkerson] ~1950

We can transmit rate h to receiver j, if the min-cut to receiver j is h, i.e., there are h edge-disjoint paths from the sources to the receiver j.

Network Multicast

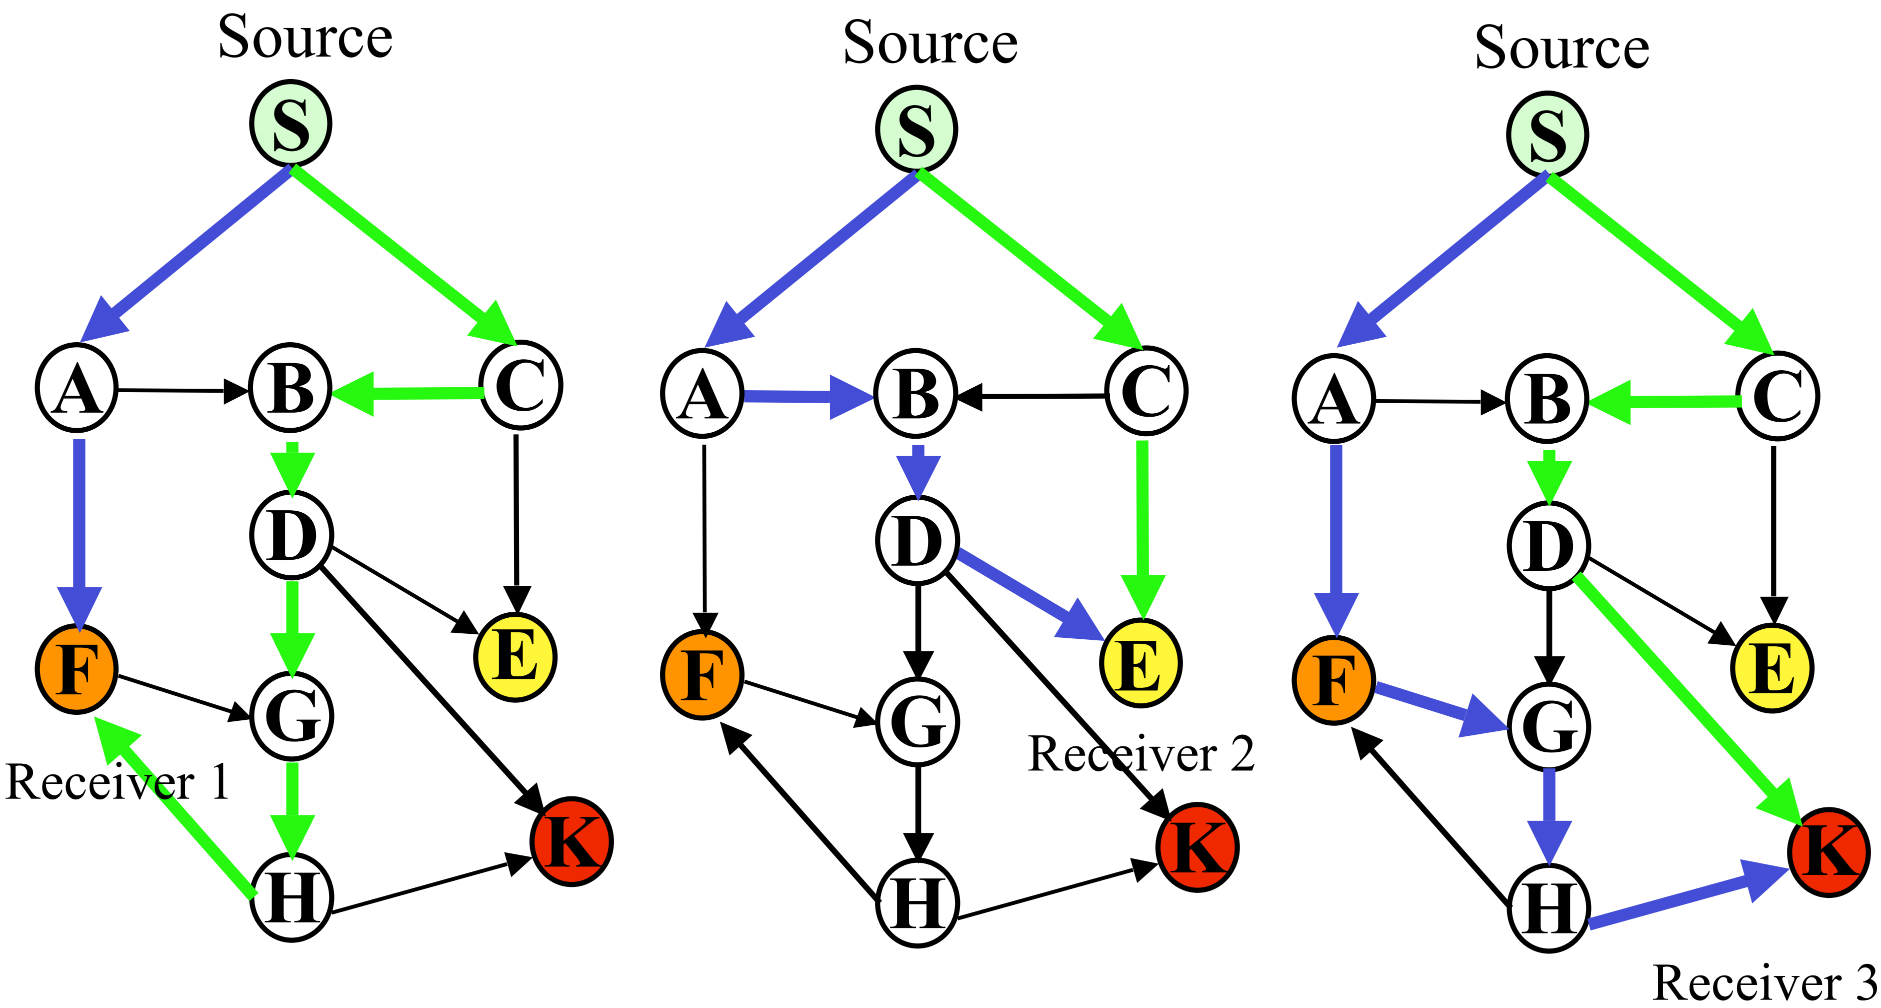

Network Multicast

- A network is represented as a directed acyclic graph with unit-capacity edges.
- There are h unit-rate information sources $S_1, \ldots, S_h$ and N receivers $R_1,\ldots,R_N$

- Max-flow min-cut theorem [Ford, Fulkerson] ~1950

We can transmit rate h to receiver j, if the min-cut to receiver j is h, i.e., there are h edge-disjoint paths from the sources to the receiver j.

- Using Network Coding: (Theorem [Alshwede,Cai,Li,Yeung] ~2000)

If the min-cut to each receiver is h, we can simultaneously transmit rate h to all receivers if each node is G can linearly re-encode information.

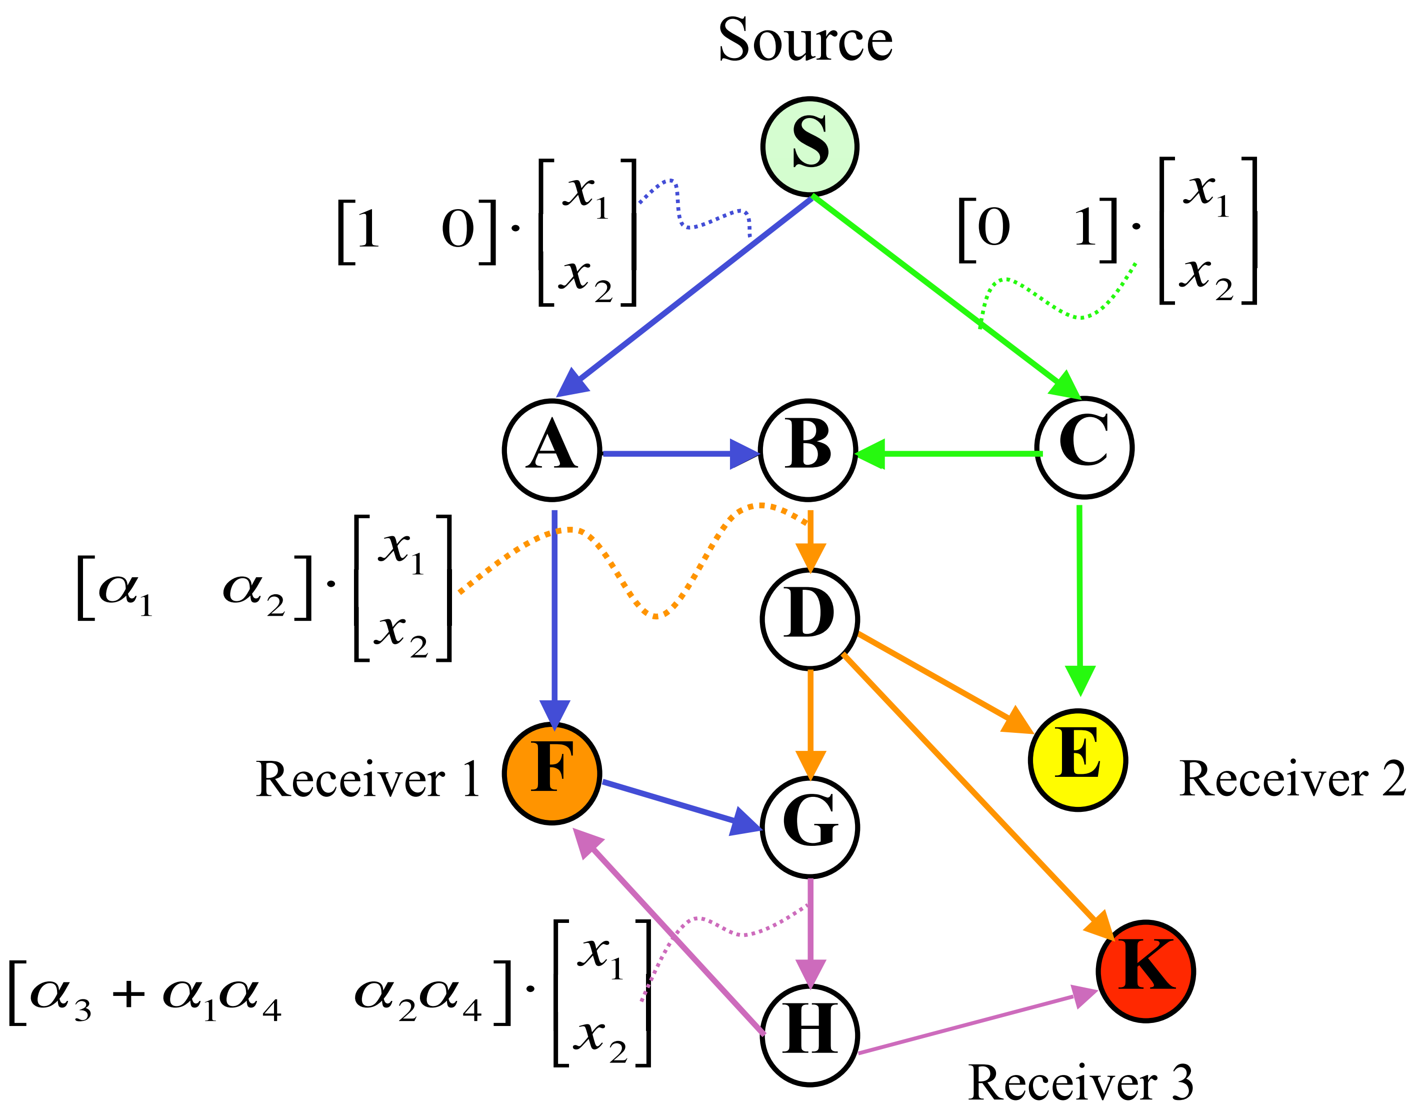

Source
S
$[1 \quad 0] \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$
$[0 \quad 1] \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$
A
B
C
$[\alpha_1 \quad \alpha_2] \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$
D
Receiver 1
F
E
Receiver 2
G
K
$[\alpha_3 + \alpha_1\alpha_4 \quad \alpha_2\alpha_4] \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$
H
Receiver 3

Coding vectors:

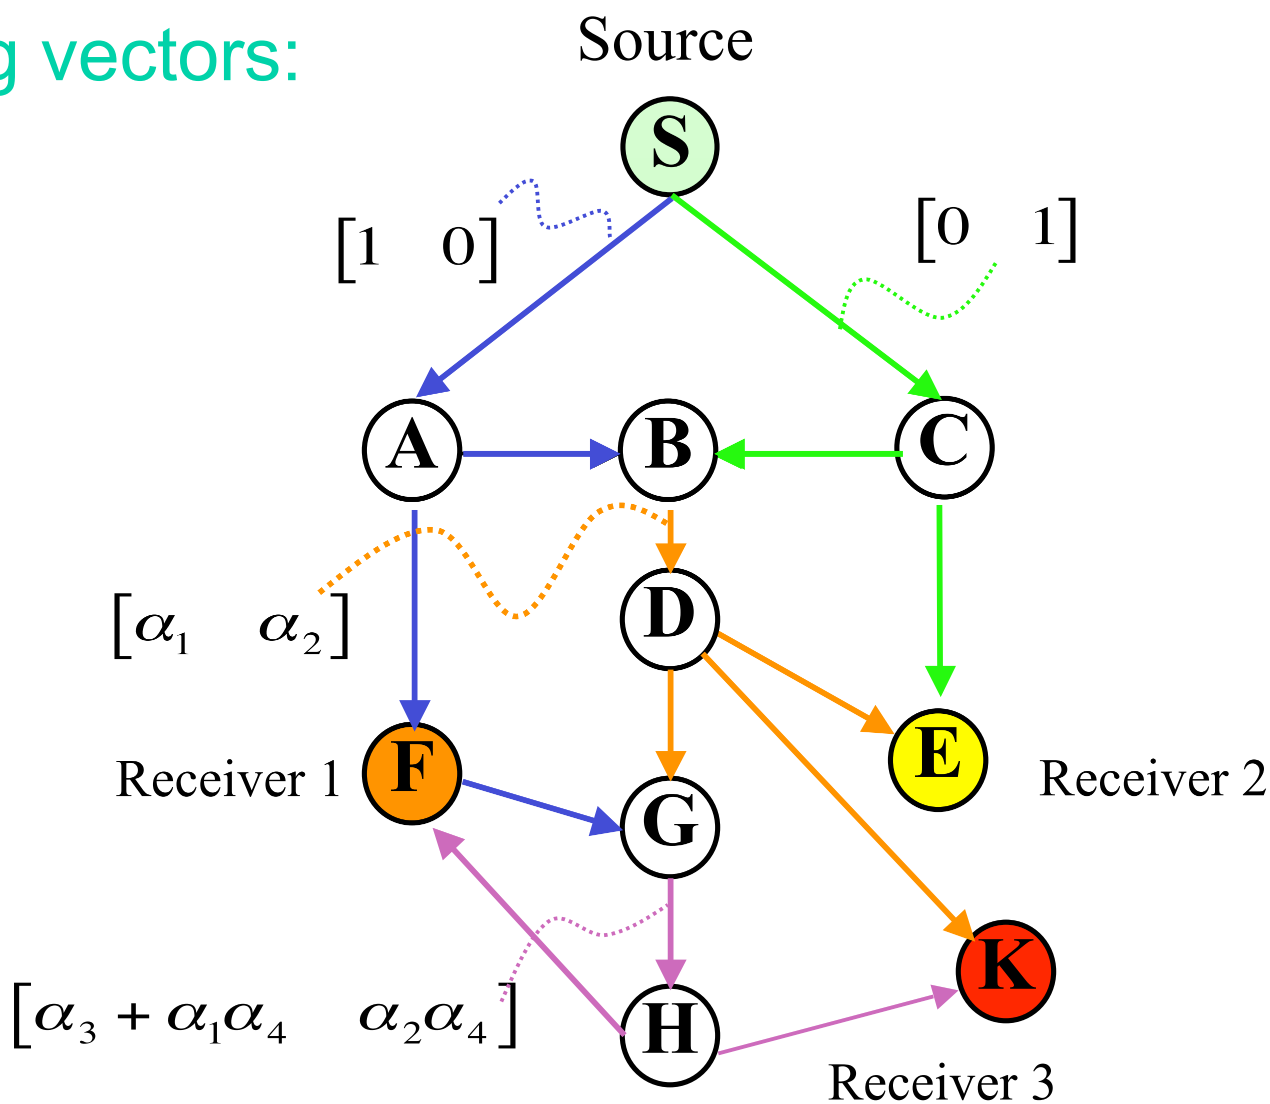

Network Coding
Linear Combining

- Source S_i emits a symbol x_i which is an element of some finite field F_q .
- Each edge carries a linear combination of its parent nodes inputs.
- Consequently, each edge carries a linear combination of the source symbols.
- The h edges a receiver observes should carry independent linear combinations of source symbols.

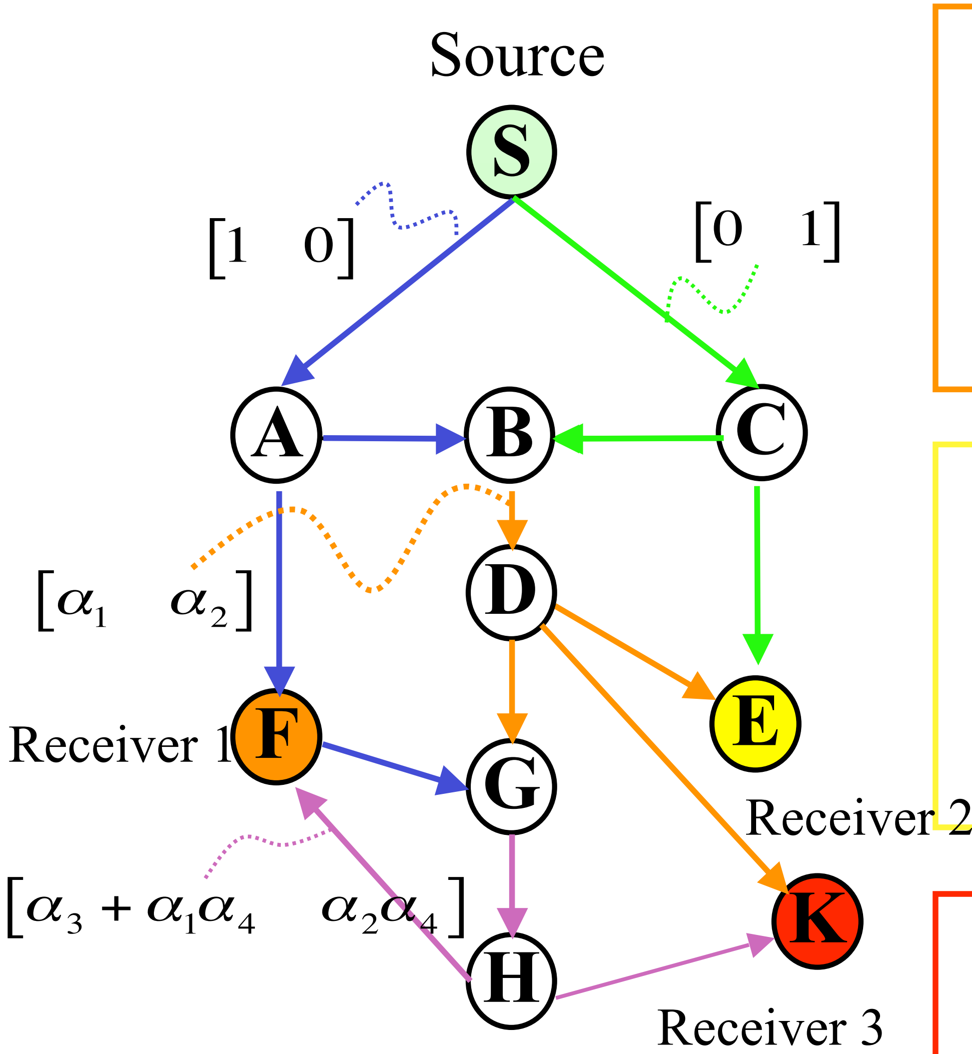

Receiver 1

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \alpha_3 + \alpha_1\alpha_4 & \alpha_2\alpha_4 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 2

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \alpha_1 & \alpha_2 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 3

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 \\ \alpha_3 + \alpha_1\alpha_4 & \alpha_2\alpha_4 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Algebraic Statement of the Main Theorem

- There exist linear coefficients $\{\alpha_i\}$ so that each receiver has a full rank of equations to solve

Receiver 1

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \alpha_3 + \alpha_1\alpha_4 & \alpha_2\alpha_4 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 2

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \alpha_1 & \alpha_2 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 3

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} \alpha_1 & \alpha_2 \\ \alpha_3 + \alpha_1\alpha_4 & \alpha_2\alpha_4 \end{bmatrix} \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Algebraic Statement of the Main Theorem

- There exist linear coefficients $\{\alpha_i\}$ so that

$$\det \mathrm{A}_1 \cdot \det \mathrm{A}_2 \cdot \det \mathrm{A}_3 \neq 0$$

Receiver 1

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_1 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 2

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_2 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 3

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_3 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Sparse Zeros Lemma

Let $f(\alpha_1,\ldots,\alpha_n)$ be a multivariate polynomial with maximum degree in each variable of at most d.

Then, in any finite field of size q where the polynomial is not identically zero, there exist values such that

$$f(\alpha_1 = p_1,\ldots,\alpha_n = p_n) \neq 0$$

Sparse Zeros Lemma

Proof

- For k=1, a polynomial of degree d can have at most d roots.
- Assume it holds for k=n-1.
- For k=n, expand the polynomial as

$$f(\alpha_1,\ldots,\alpha_n) = \sum_{i=0}^{d} f_i(\alpha_1,\ldots,\alpha_{n-1}) \cdot \alpha_n^i$$

- From induction, there exist values so that at least one coefficient polynomial is nonzero. Substituting these values, we get a polynomial in one variable.

Given a graph, how do we calculate the transfer matrices?

Receiver 1

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_1 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 2

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_2 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Receiver 3

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \mathbf{A}_3 \cdot \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Transfer Matrix Calculation

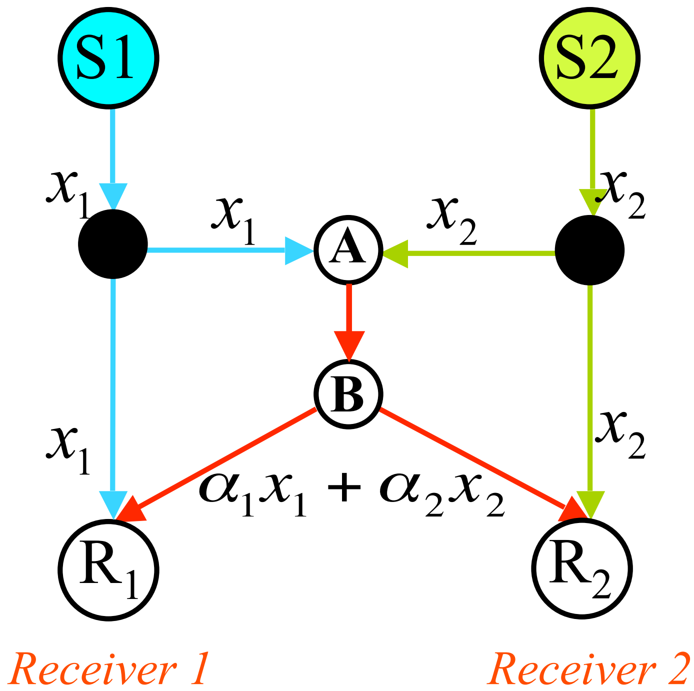

- Select the linear combinations so that each receiver has a full rank system of equations to solve. In the example:

Receiver 1 $\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} 1 & 0 \\ \alpha_1 & \alpha_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$

Receiver 2 $\begin{bmatrix} y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} 0 & 1 \\ \alpha_1 & \alpha_2 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$

Transfer matrix calculation

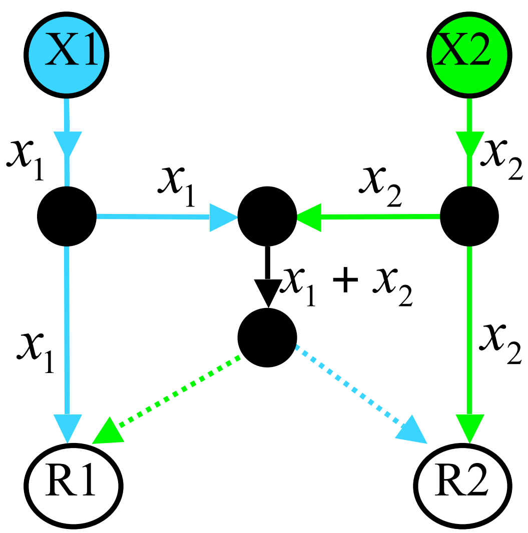

$\{x_1, x_1 + x_2\}$ $\{x_2, x_1 + x_2\}$

Underlying Assumption:

All nodes simultaneously receive their inputs and produce their outputs.

Connection with convolutional codes

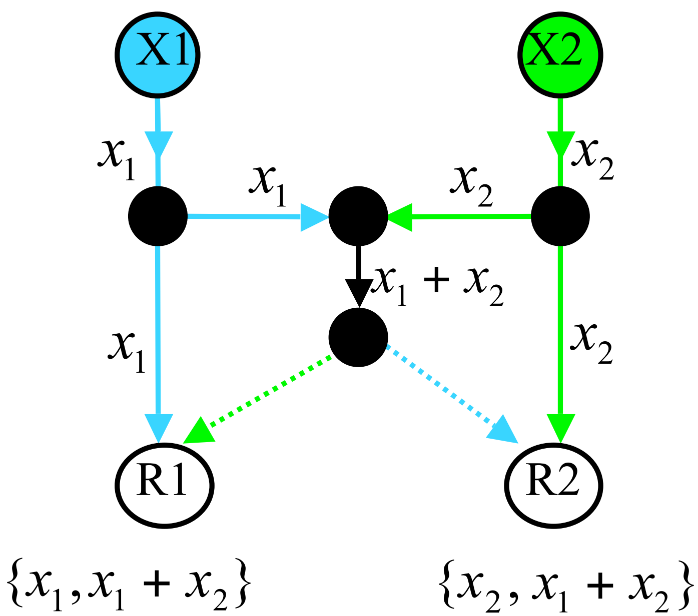

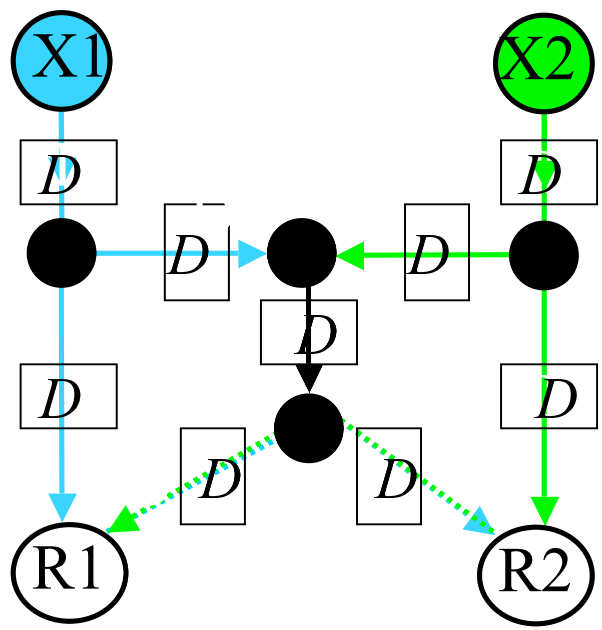

$$\begin{bmatrix} y_1 \\ y_2 \end{bmatrix} = \begin{bmatrix} D^2 & 0 \\ \alpha_1 D^4 & \alpha_2 D^4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \quad \begin{bmatrix} y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} 0 & D^2 \\ \alpha_1 D^4 & \alpha_2 D^4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Transfer Matrix

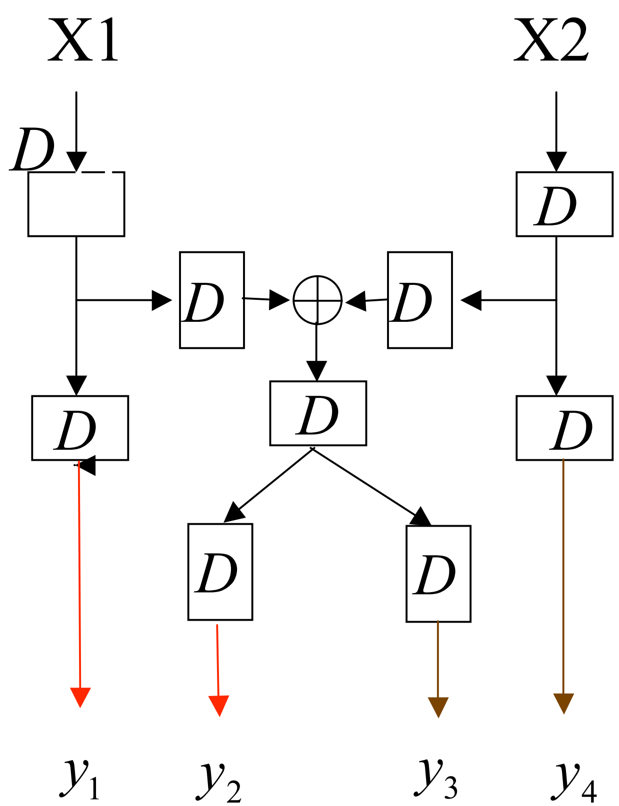

$$s_{j+1} = As_j + Bx_j$$

$$y_{j+1} = Cs_j + \Delta x_j$$

$$y = \boxed{(\Delta + C(\mathrm{D}^{-1}I - A)^{-1}B)}x$$

$$\begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} D^2 & 0 \\ \alpha_1 D^4 & \alpha_2 D^4 \\ 0 & D^2 \\ \alpha_1 D^4 & \alpha_2 D^4 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix}$$

Transfer Matrix

For receiver j:

$$\mathrm{A}_j = C_j (I - A)^{-1} B$$

$$s_{j+1} = A s_j + B x_j$$

$$y_{j+1} = C_j s_j$$

Theorem

In a network with N receivers an alphabet of size q>N is always sufficient.

$$\det \mathrm{A}_1 \cdot \det \mathrm{A}_2 \cdots \det \mathrm{A}_N \neq 0$$

Lemma

Let $\mathrm{A}_1 = C_1 (I - A)^{-1} B$

Then $|\det(\mathrm{A}_1)| = |\det(N_1)|$ where $N_1 = \begin{bmatrix} C_1 & 0 \\ I - A & B \end{bmatrix}$

$$\det N_1 \cdot \det N_2 \cdots \det N_N \neq 0$$

Design: Matrix A

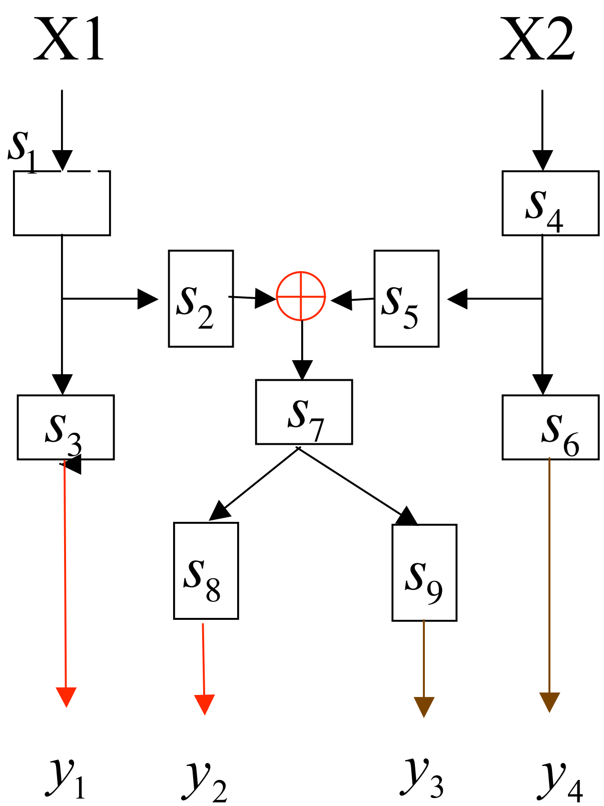

$$s_{j+1} = \boxed{A} s_j + Bx_j$$

$$y_{j+1} = Cs_j + Dx_j$$

$$\begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \end{bmatrix} = \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & * & 0 & 0 & * & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \cdot \begin{bmatrix} s_1 \\ s_2 \\ s_3 \\ s_4 \\ s_5 \\ s_6 \\ s_7 \\ s_8 \\ s_9 \end{bmatrix} + \ldots$$

Discussion on Main Theorem

How restrictive are the assumptions?

1) Directed graph
2) Acyclic graph
3) Same min-cut to all receivers

Undirected Graphs

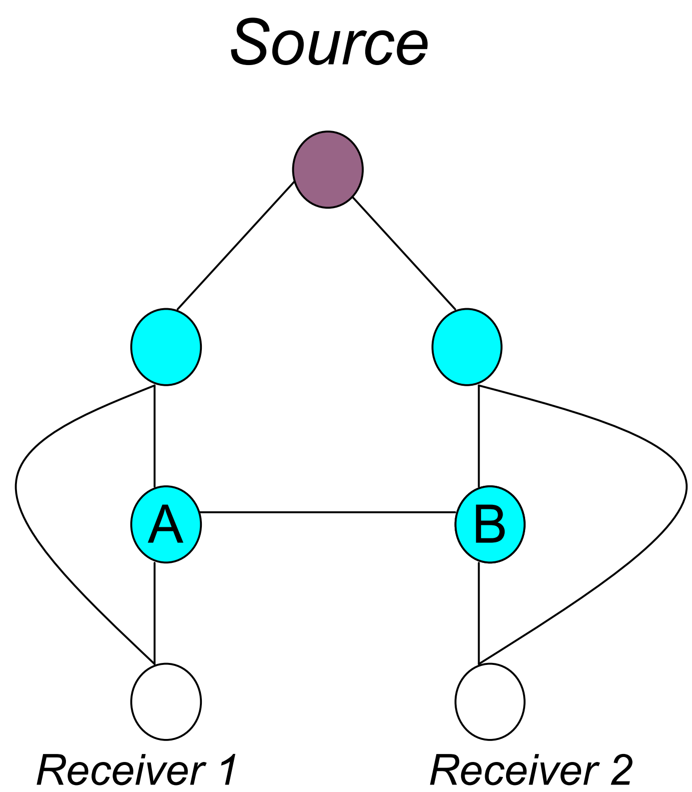

Undirected Graphs

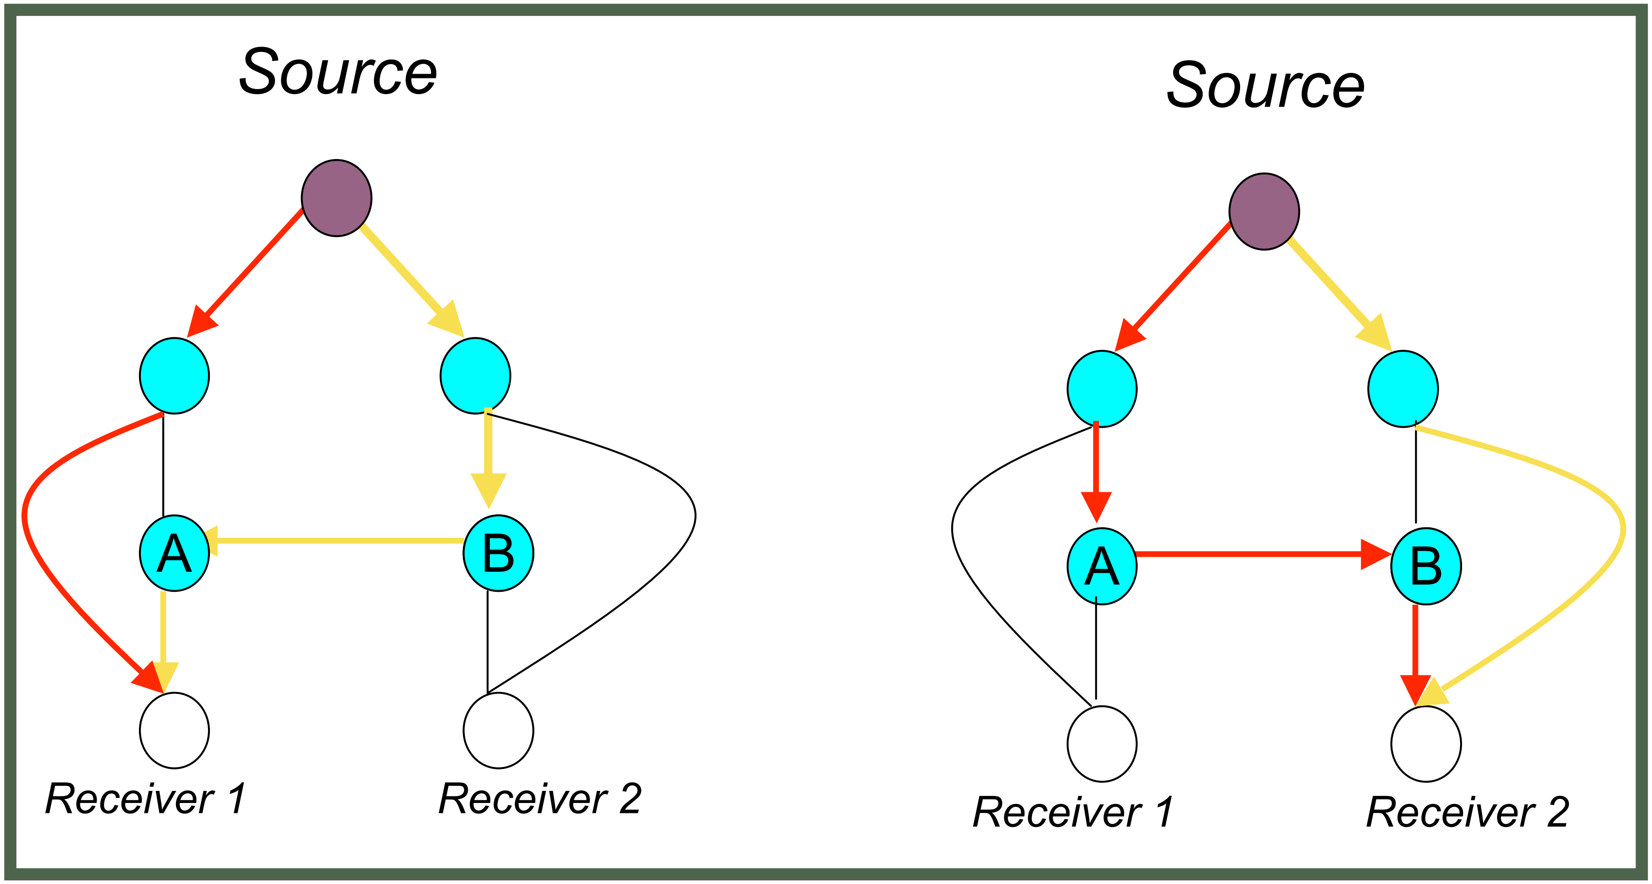

Undirected Graphs

Consider a network represented as an undirected graph with unit-capacity edges, h unit-rate information sources $S_1, \ldots, S_h$ located on the same vertex of the graph and N receivers $R_1, \ldots, R_N$. Assume the min-cut to each receiver is h.

We do not know the solution in general

Undirected Graphs

Consider a network represented as an undirected graph with unit-capacity edges, h unit-rate information sources $S_1, \ldots, S_h$ located on the same vertex of the graph and N receivers $R_1,\ldots,R_N$. Assume the min-cut to each receiver is h.

[Li and Li] ~2003

We can simultaneously transmit rate h/2 to all receivers, even when only using routing.

Undirected Graphs

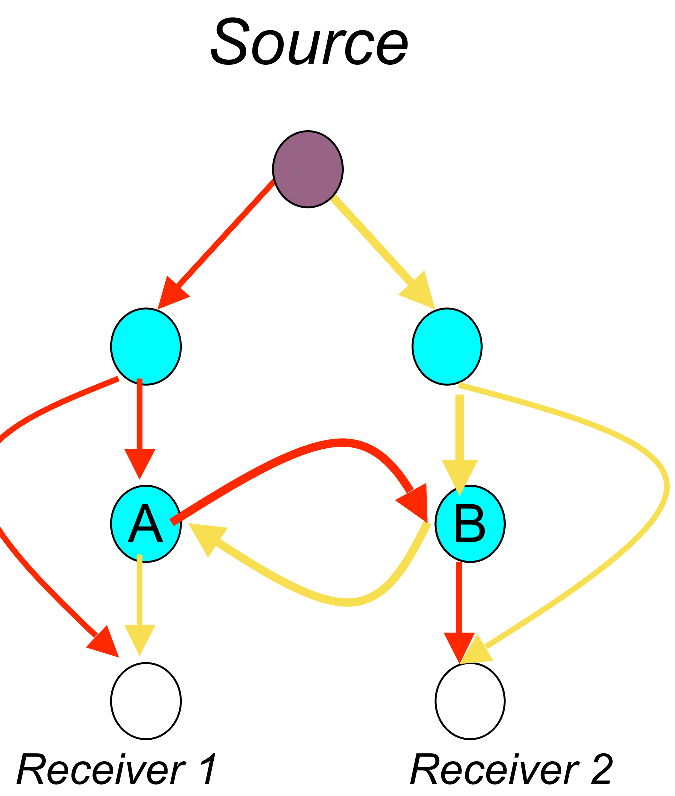

Discussion on Main Theorem

How restrictive are the assumptions?

1) Directed graph Restrictive
2) Acyclic graph Not Restrictive
3) Same min-cut to all receivers Restrictive

Receivers with different min-cut

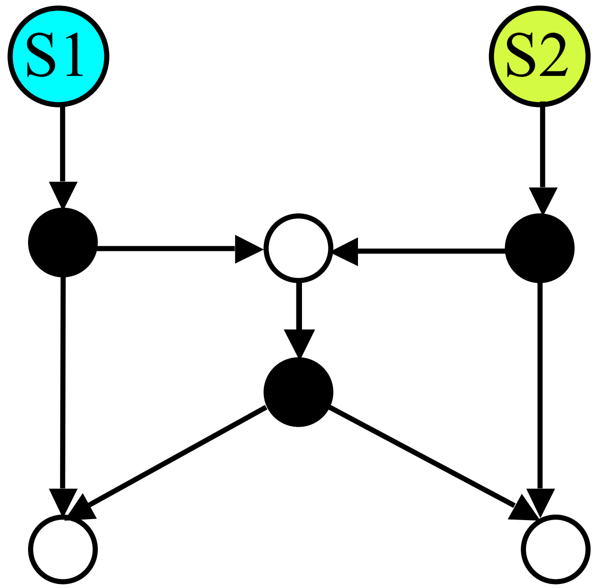

Open Questions

1. How do we design network codes.

2. Relationship between graph structure and network code design.

3. How large an alphabet do we need to use?

4. What are the throughput benefits that we may hope to get?

5. ….

Outline

1. Main Theorem in Multicasting

2. Benefits and Requirements

2. Benefits and Requirements

1. Throughput
2. Routing complexity
3. Energy
4. Delay
5. …..

Complexity
(operational and set-up complexity)

Combinatorial framework:

Information flow decomposition

Complexity Requirements

- How large a finite field do we need?
- How many nodes in the network need to perform linear combining operations?
- How difficult is it to design network codes?
- What is the encoding and decoding complexity?

Motivation

Are there common structural properties in multicast configurations with the same number of sources and receivers?

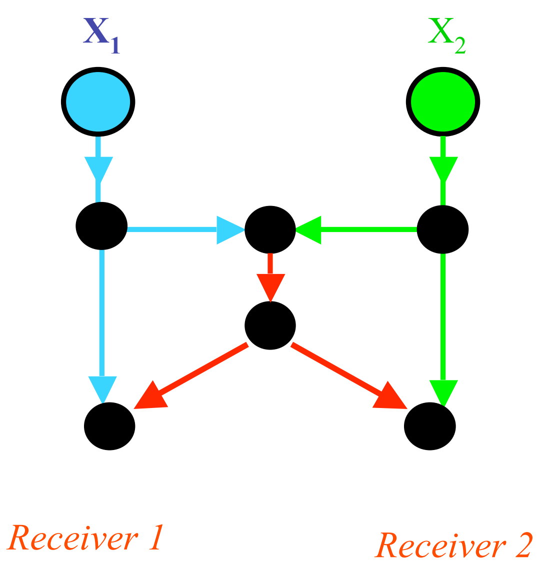

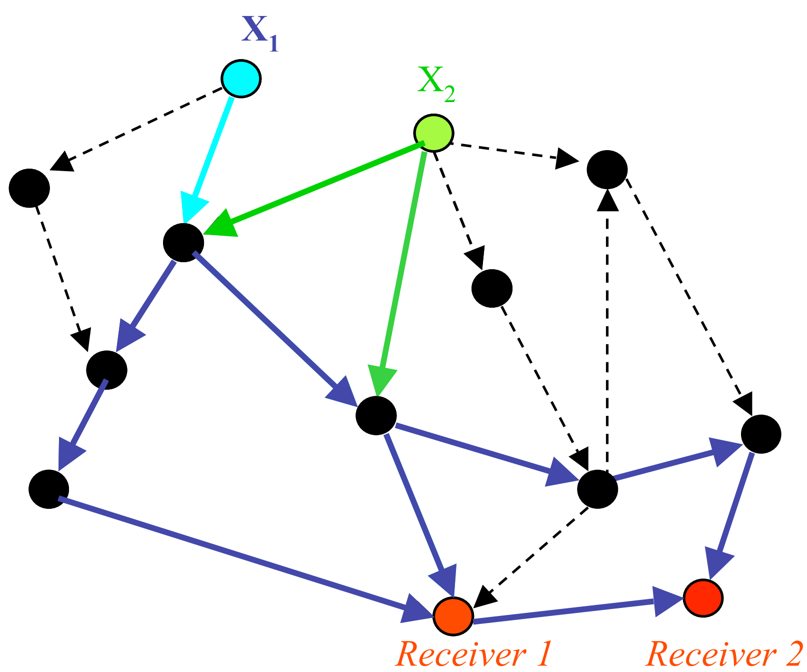

Information Flow Decomposition

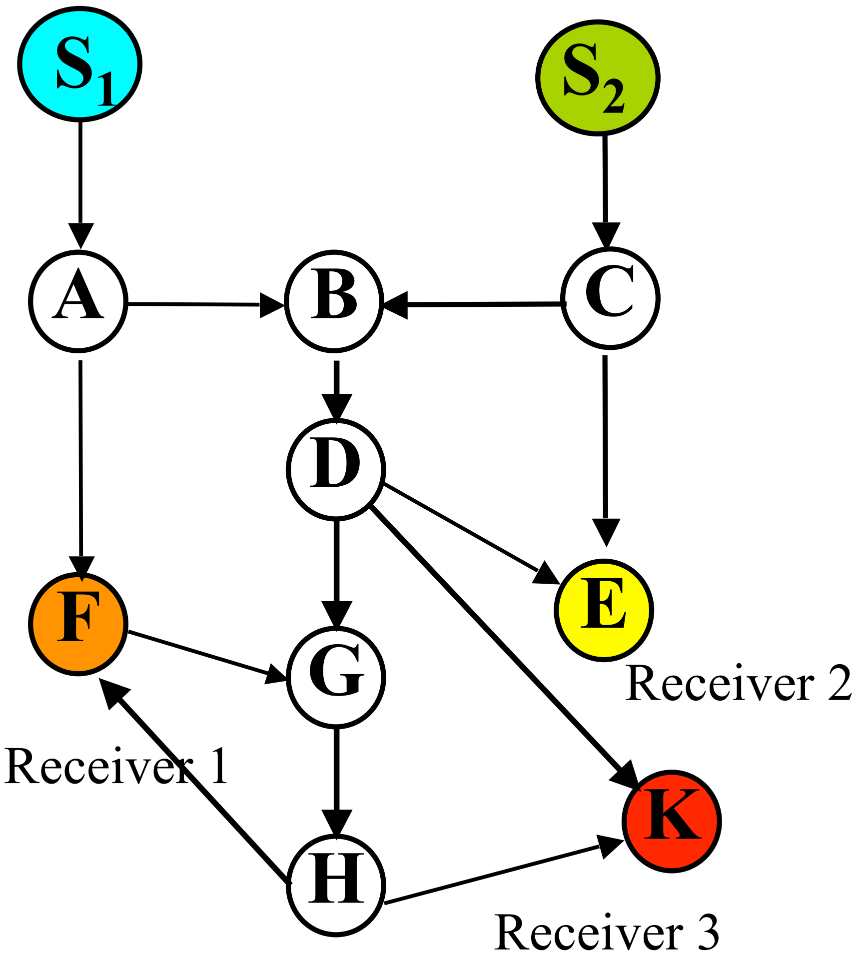

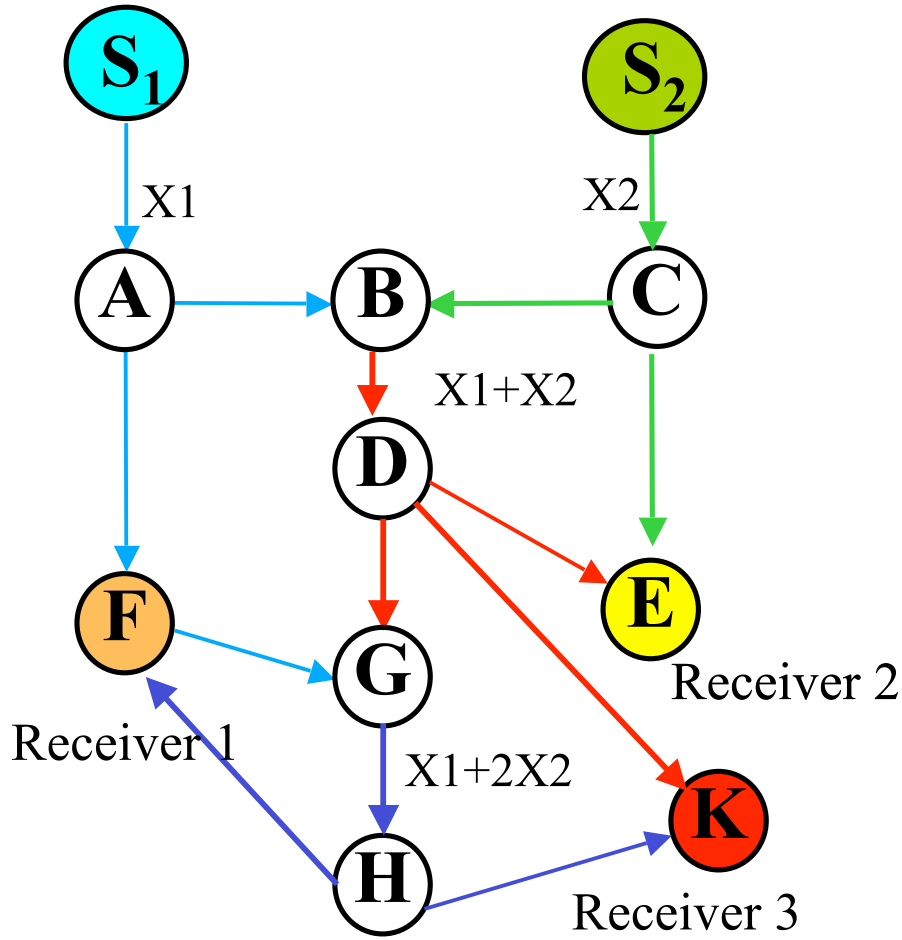

Subtree graph

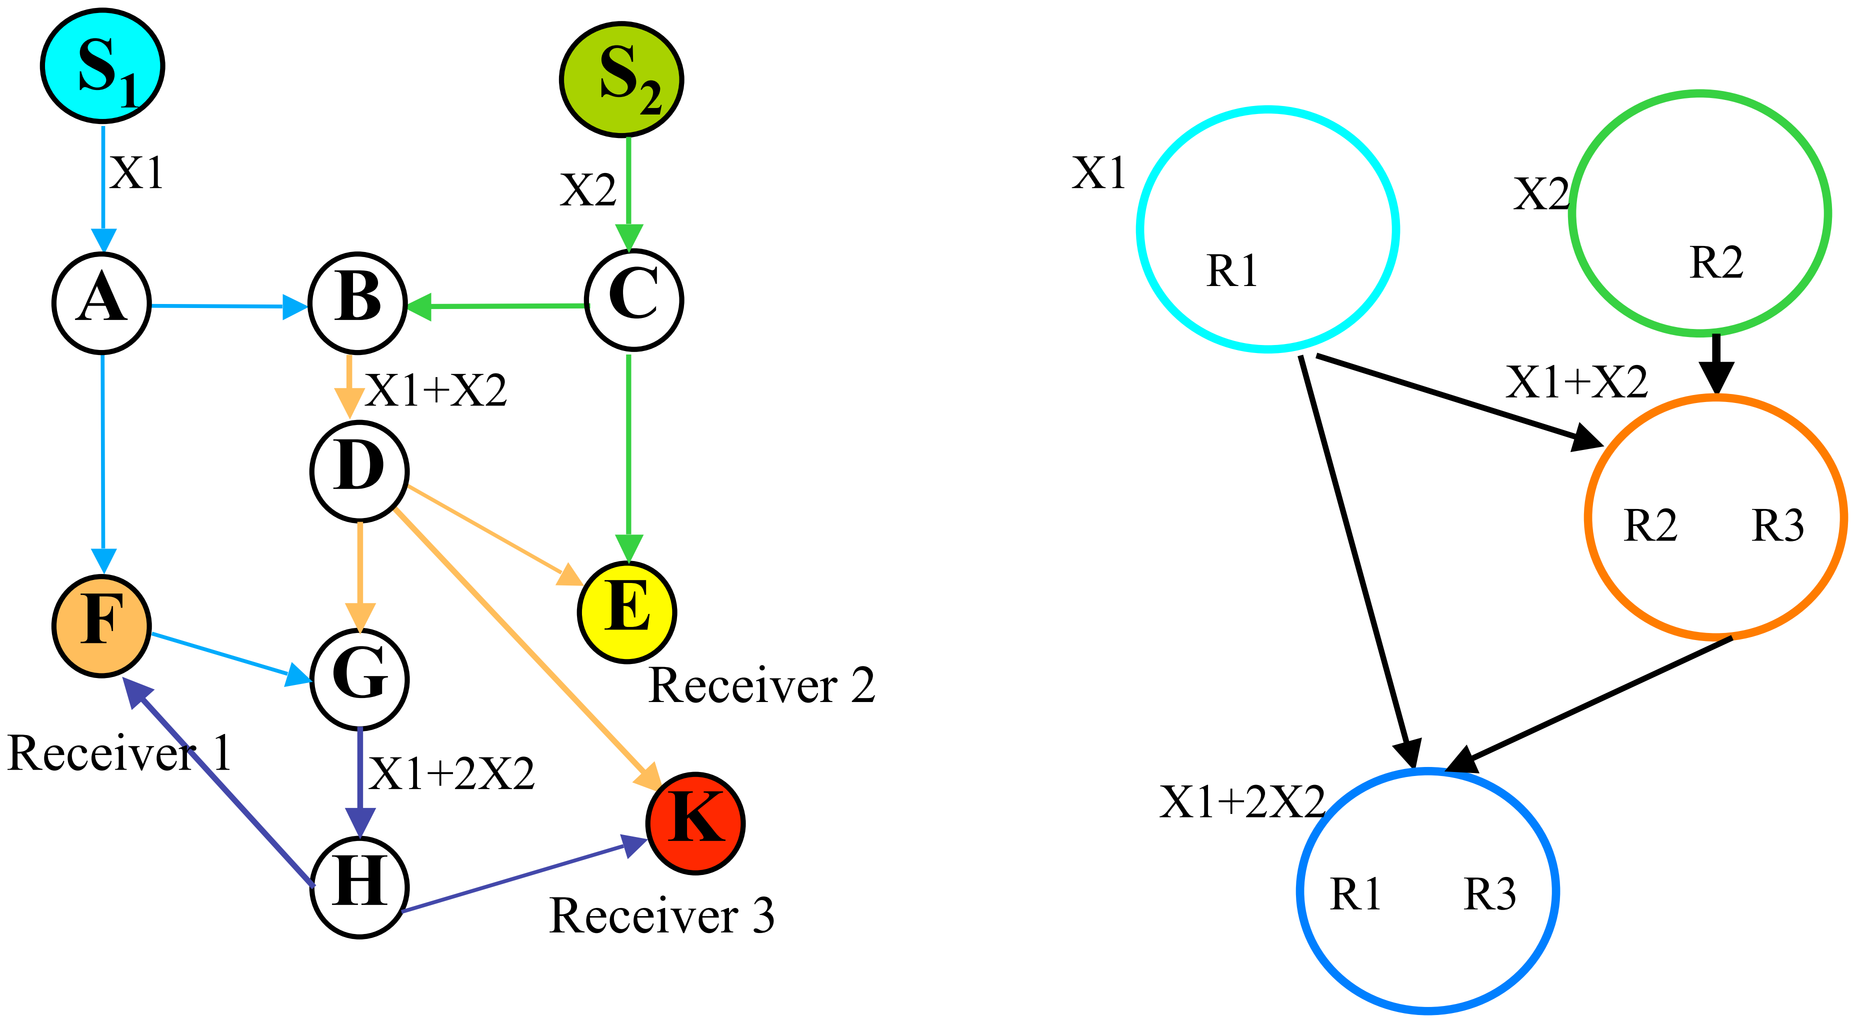

Contract each area of the network through which the same information flows to a vertex

Example: Two Sources-Two Receivers

There exist only two cases:

No network coding

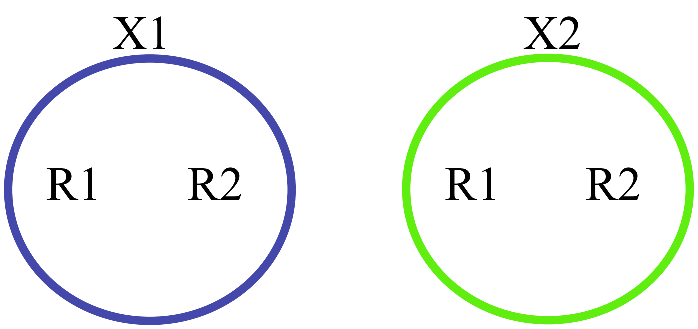

Network coding at one edge

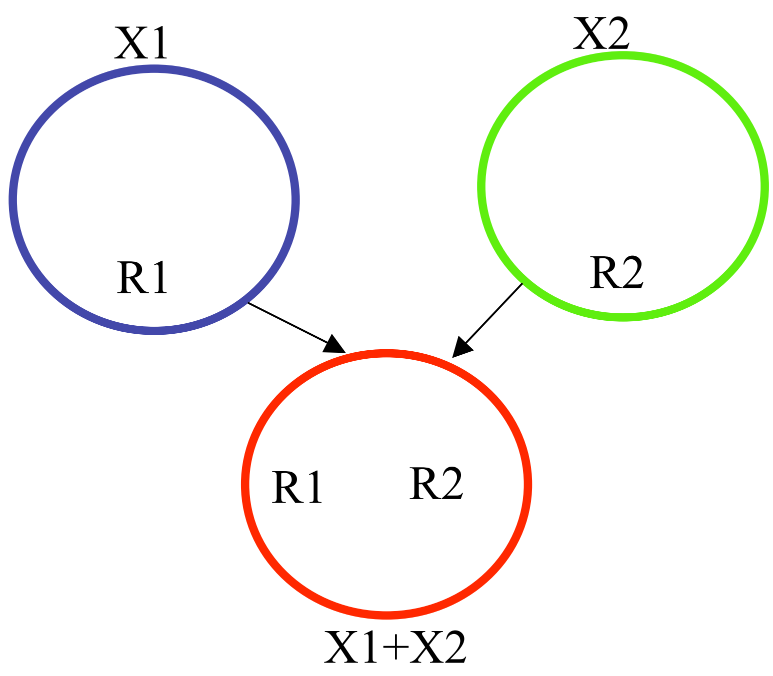

Butterfly Network

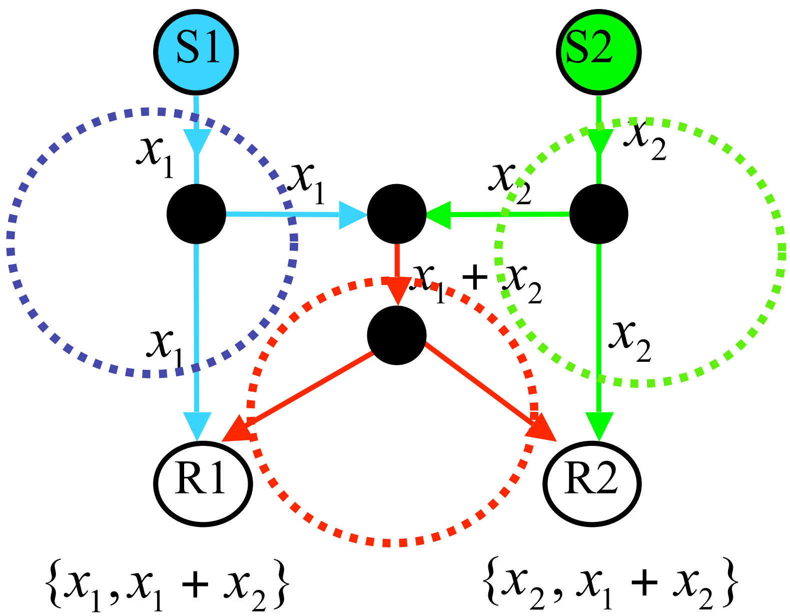

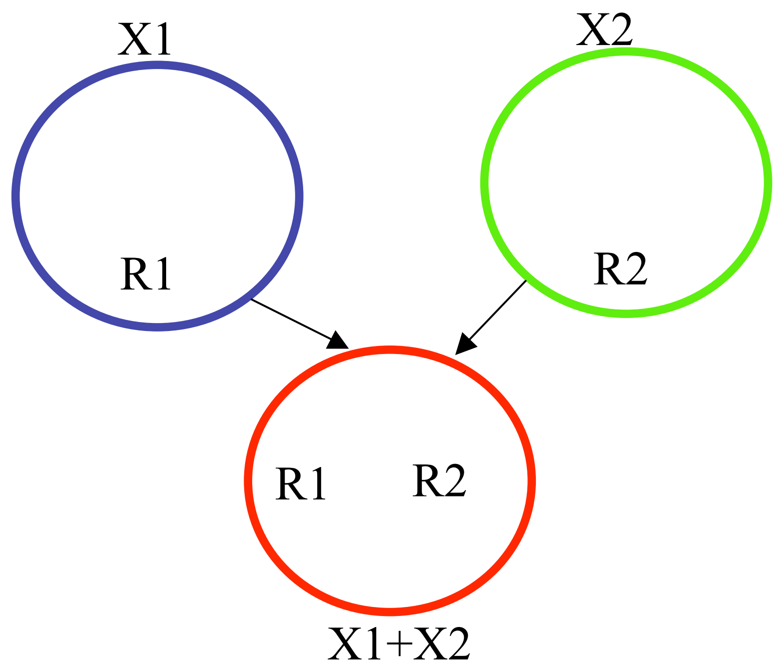

Two Sources-Three Receivers

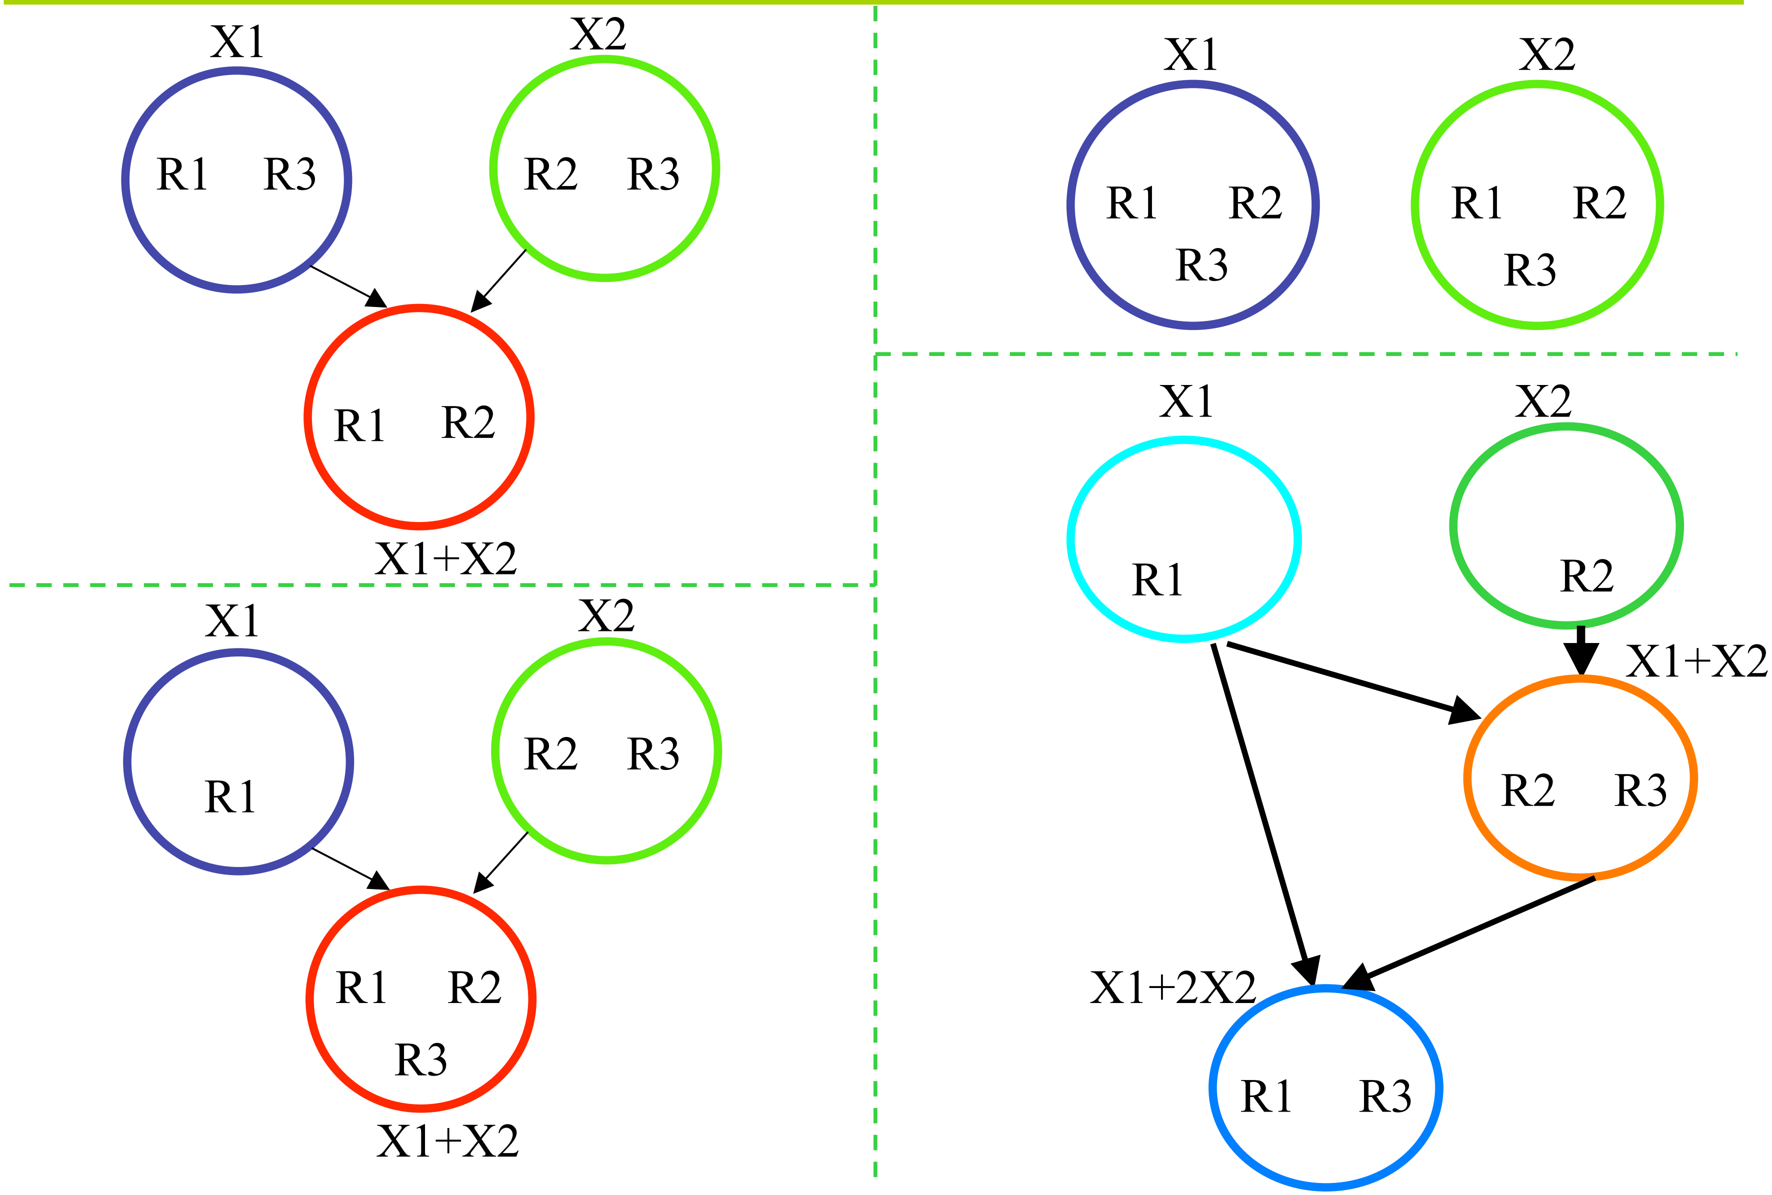

Theorem

(Fragouli, Soljanin 2004)

Starting from multicast configuration with h sources and N receivers over an arbitrary graph we can find the associated subtree graph in polynomial time.

- Very distinct graphs correspond to the same subtree graph.
- The subtree graph has a much smaller number of vertices.
- Precompute network codes.

Number of “coding points”

In configurations with h=2 sources and N receivers we have at most N-1 coding points.

Generally, in configurations with h sources and N receivers we have at most h^2N^3 coding points.

Applications of the information flow decomposition

- Derive theoretical results, for example
  1. Alphabet size bounds
  2. Throughput benefits

- Design practical network coding schemes
  -distributed algorithms
  -convolutional network codes

Alphabet size

- Directed graph with unit capacity edges, coding over F_q.
- What alphabet size q is sufficient for all possible configurations with h sources and N receivers?

Bounds on Alphabet Size

- Infinite (Ahlswede, Yeung, Cai et al. 2000)
- N (Sanders, Egner, Koetter, Medard, et al. 2003)
- For any configuration with h=2 sources and N receivers

an alphabet of size $\left\lfloor \sqrt{2N - \frac{7}{4}} + \frac{1}{2} \right\rfloor$ is always sufficient
(Fragouli, Soljanin, Shokrollahi 2004)

Alphabet size

[CISS 2004, Trans. IT 2005]

- Directed graph with unit capacity edges, coding over F_q.
- What alphabet size q is sufficient for all possible configurations with h=2 sources and N receivers?

Theorem: For any configuration with h=2 sources and N receivers an alphabet of size $\left\lfloor \sqrt{2N - \frac{7}{4}} + \frac{1}{2} \right\rfloor$ is always sufficient.

We will show that the problem of designing a network code for h=2 sources can be reduced to the problem of coloring an appropriately defined graph.

Network code design:
satisfy some linear independence conditions

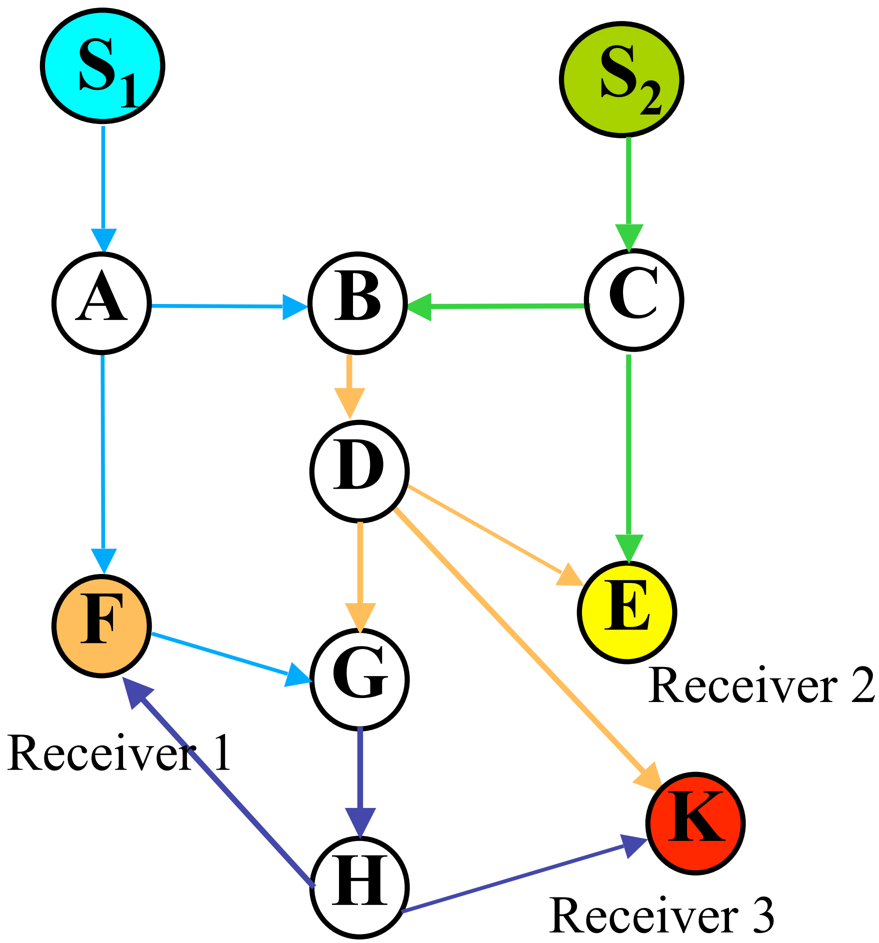

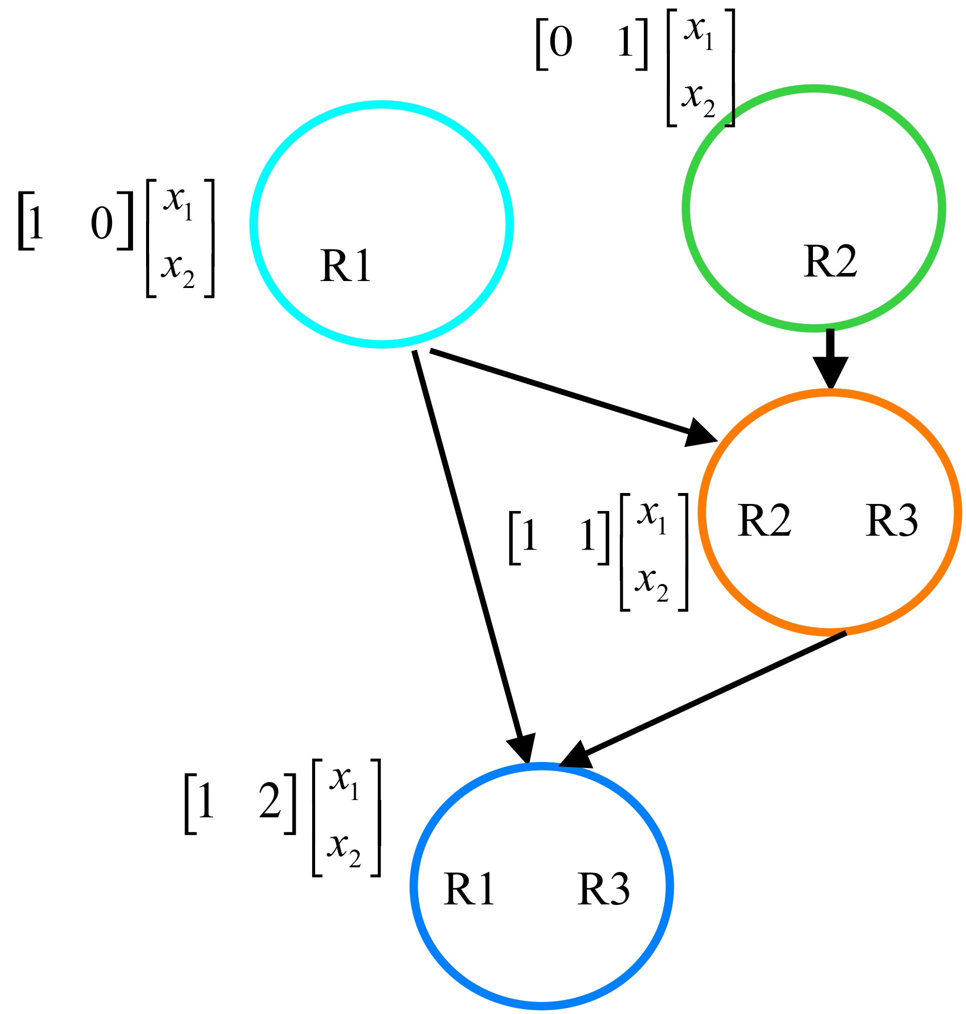

Colors

For h=2, it is sufficient to consider q+1 coding vectors over F_q:

$$[0 \quad 1] \quad [1 \quad 0] \quad [1 \quad a] \quad [1 \quad a^2] \quad \cdots \quad [1 \quad a^{q-1}]$$

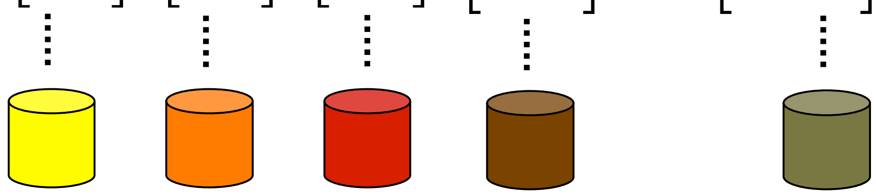

Any two such vectors form a basis of the 2-dimensional space

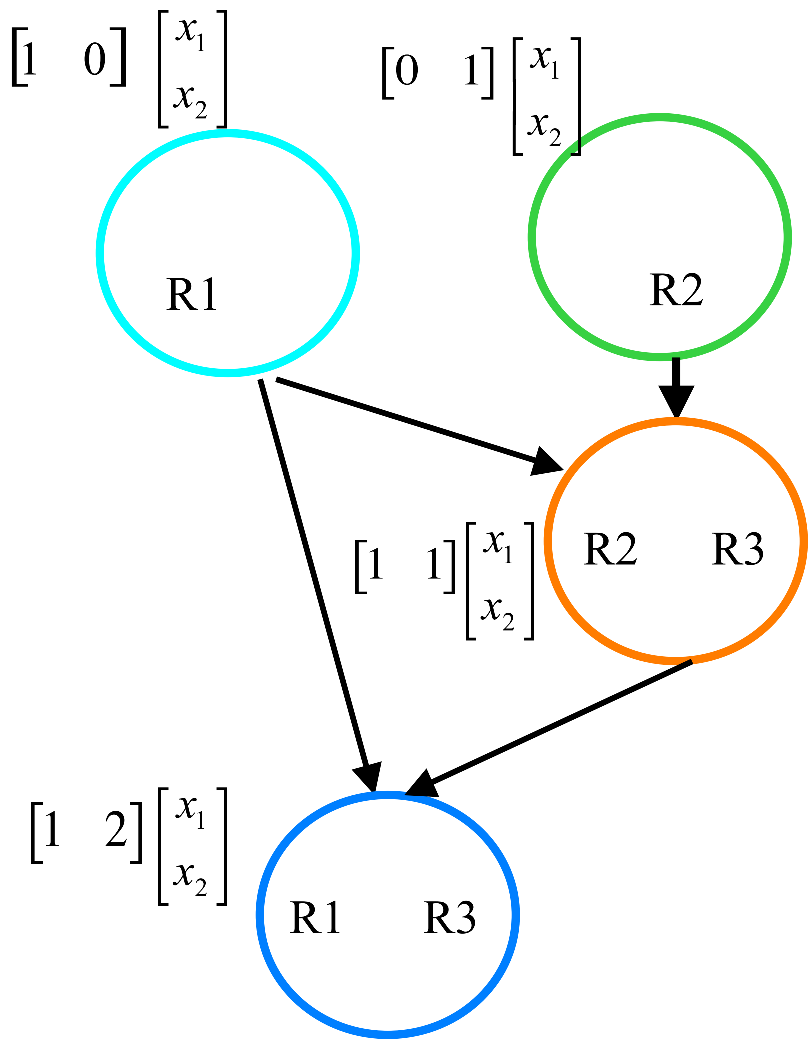

Graph to color

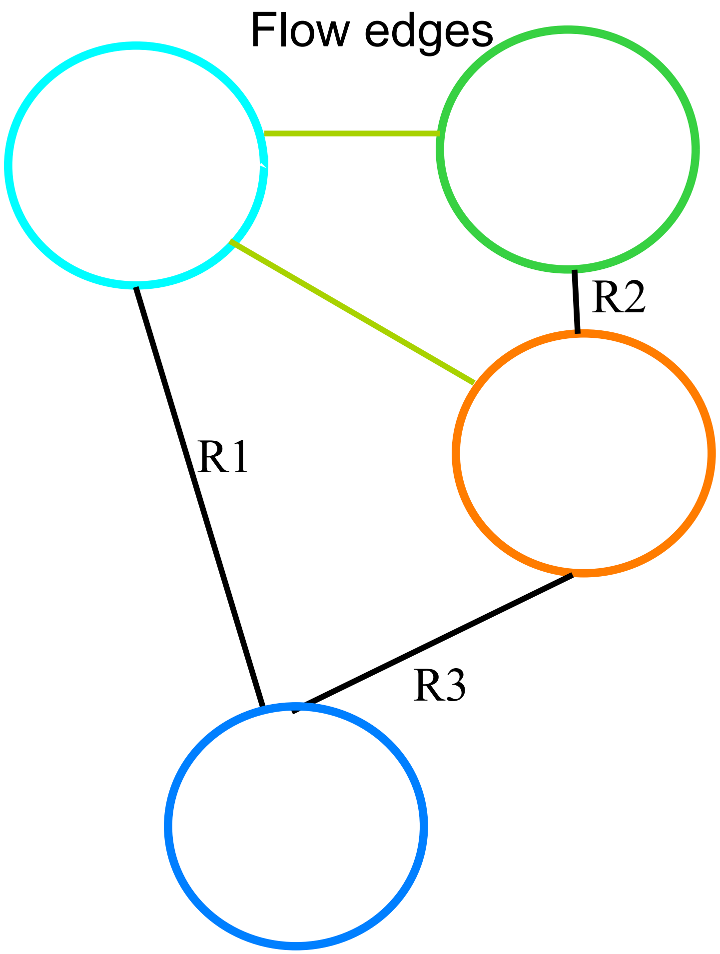

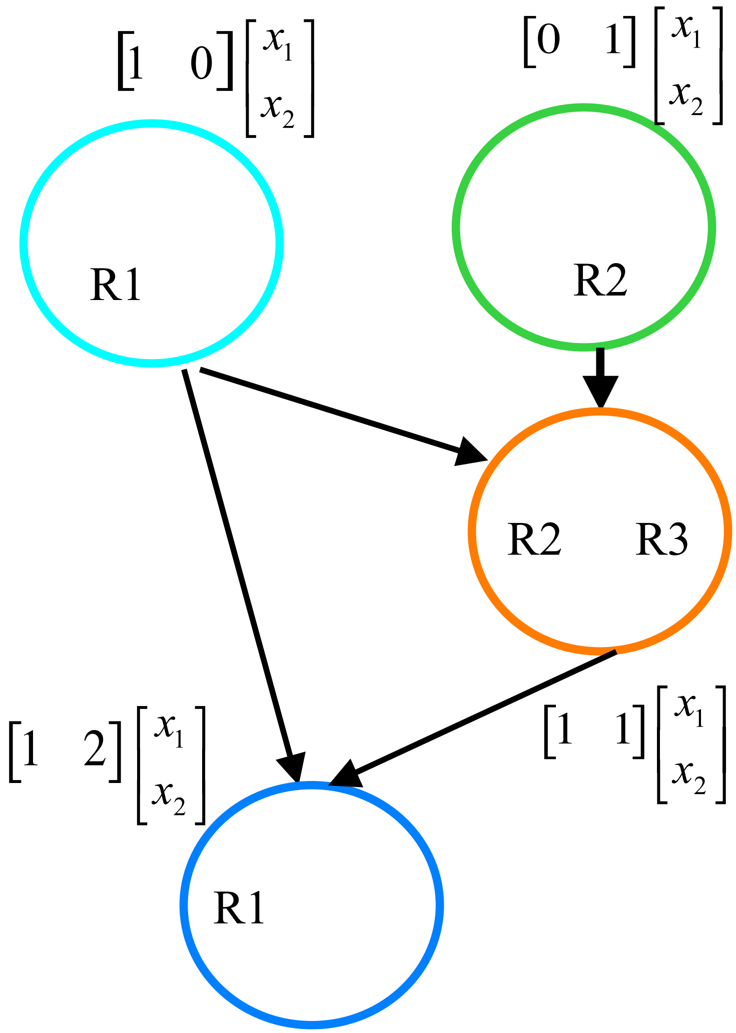

Coloring problem

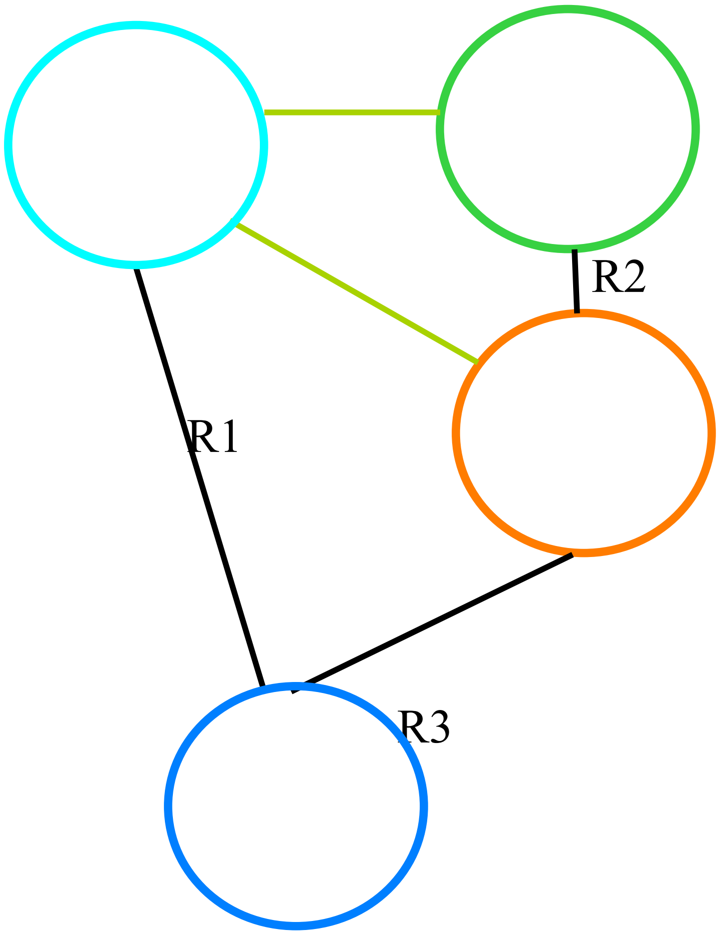

q+1 coding vectors over F_q:

$$[0\quad 1]\quad [1\quad 0]\quad [1\quad a]\quad [1\quad a^2]\quad \cdots\quad [1\quad a^{q-1}]$$

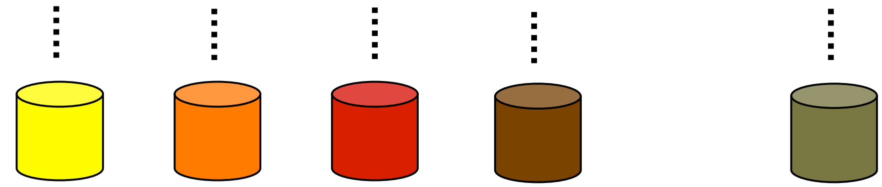

Elements of Proof:

1. If k colors are required:
-the graph has k vertices of degree k-1
-an alphabet of size q=k-1 is required

2. If we have N receivers there exist
-At most N receiver edges
-At most N-1 flow edges

Thus $q \geq \left\lfloor \sqrt{2N - \frac{7}{4}} + \frac{1}{2} \right\rfloor$

Throughput benefits

How much do we lose if we don’t use network coding?

What throughput we can get by only using routing?

Throughput benefits

How much do we lose if we don't use network coding?

What throughput we can get by only using routing?

Common throughput= 1
Average throughput=1.5

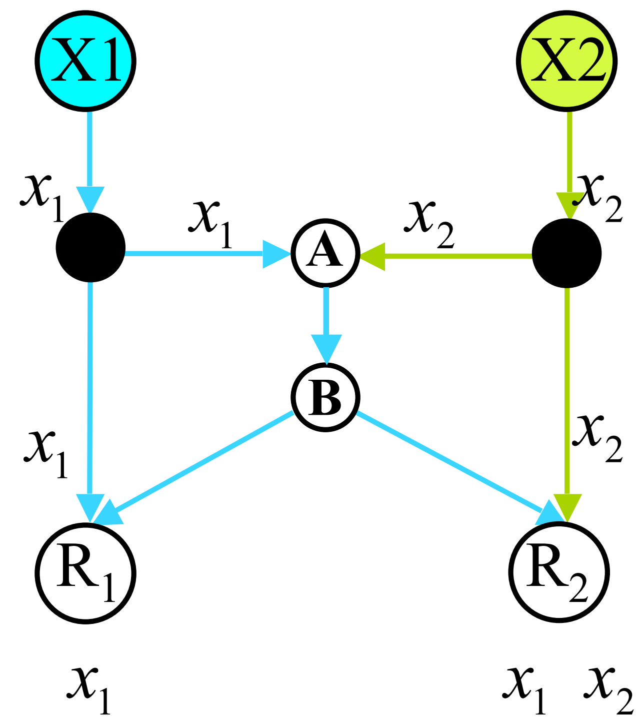

Common Throughput: Packing Steiner Trees

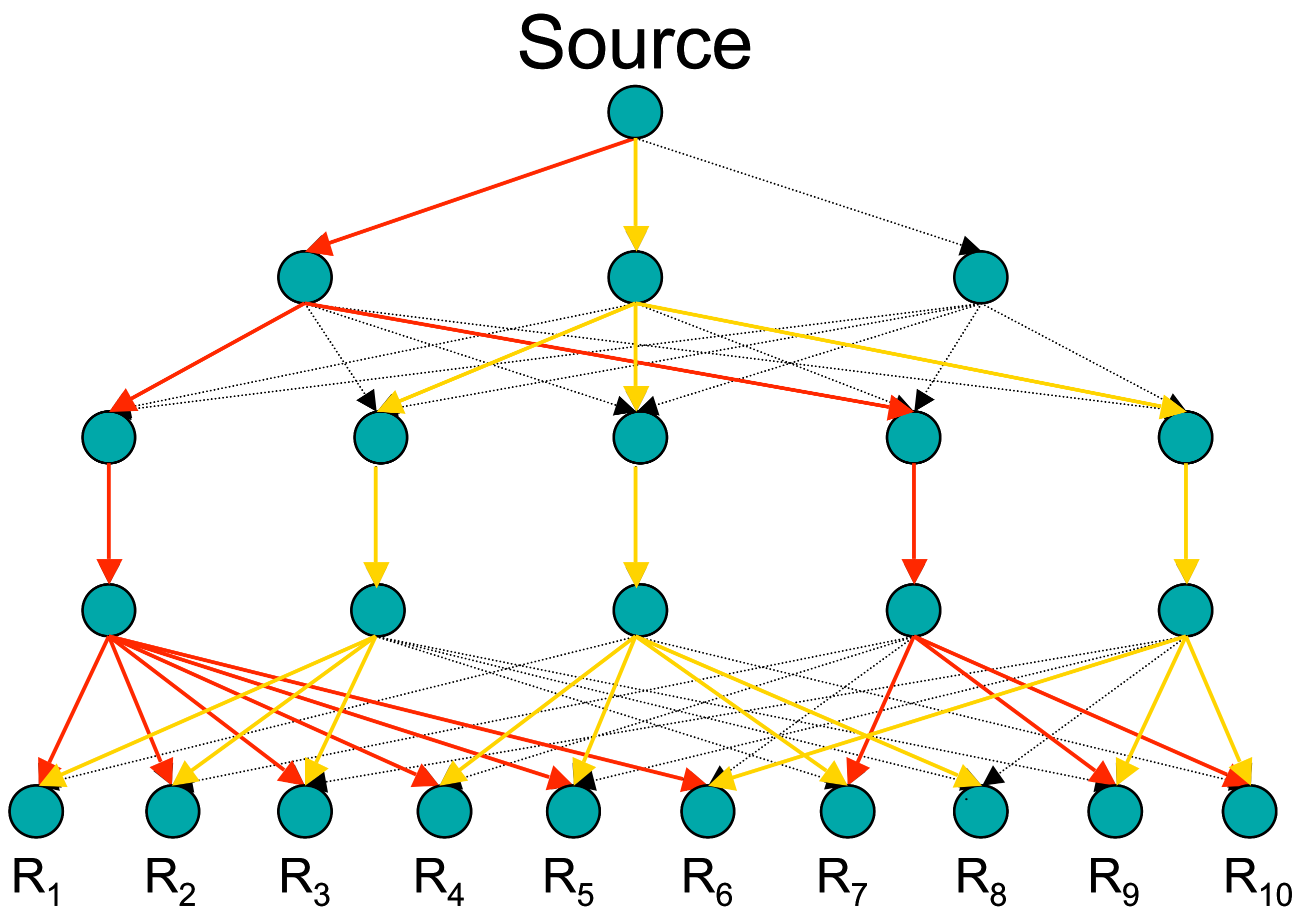

Common throughput benefits

There exist directed graphs where network coding offers throughput benefits as compared to the average throughput proportional to h where h is the number of sources.

(Sanders et al. 2002)

For $k>h^2$ there exist h edges that get allocated the same source.

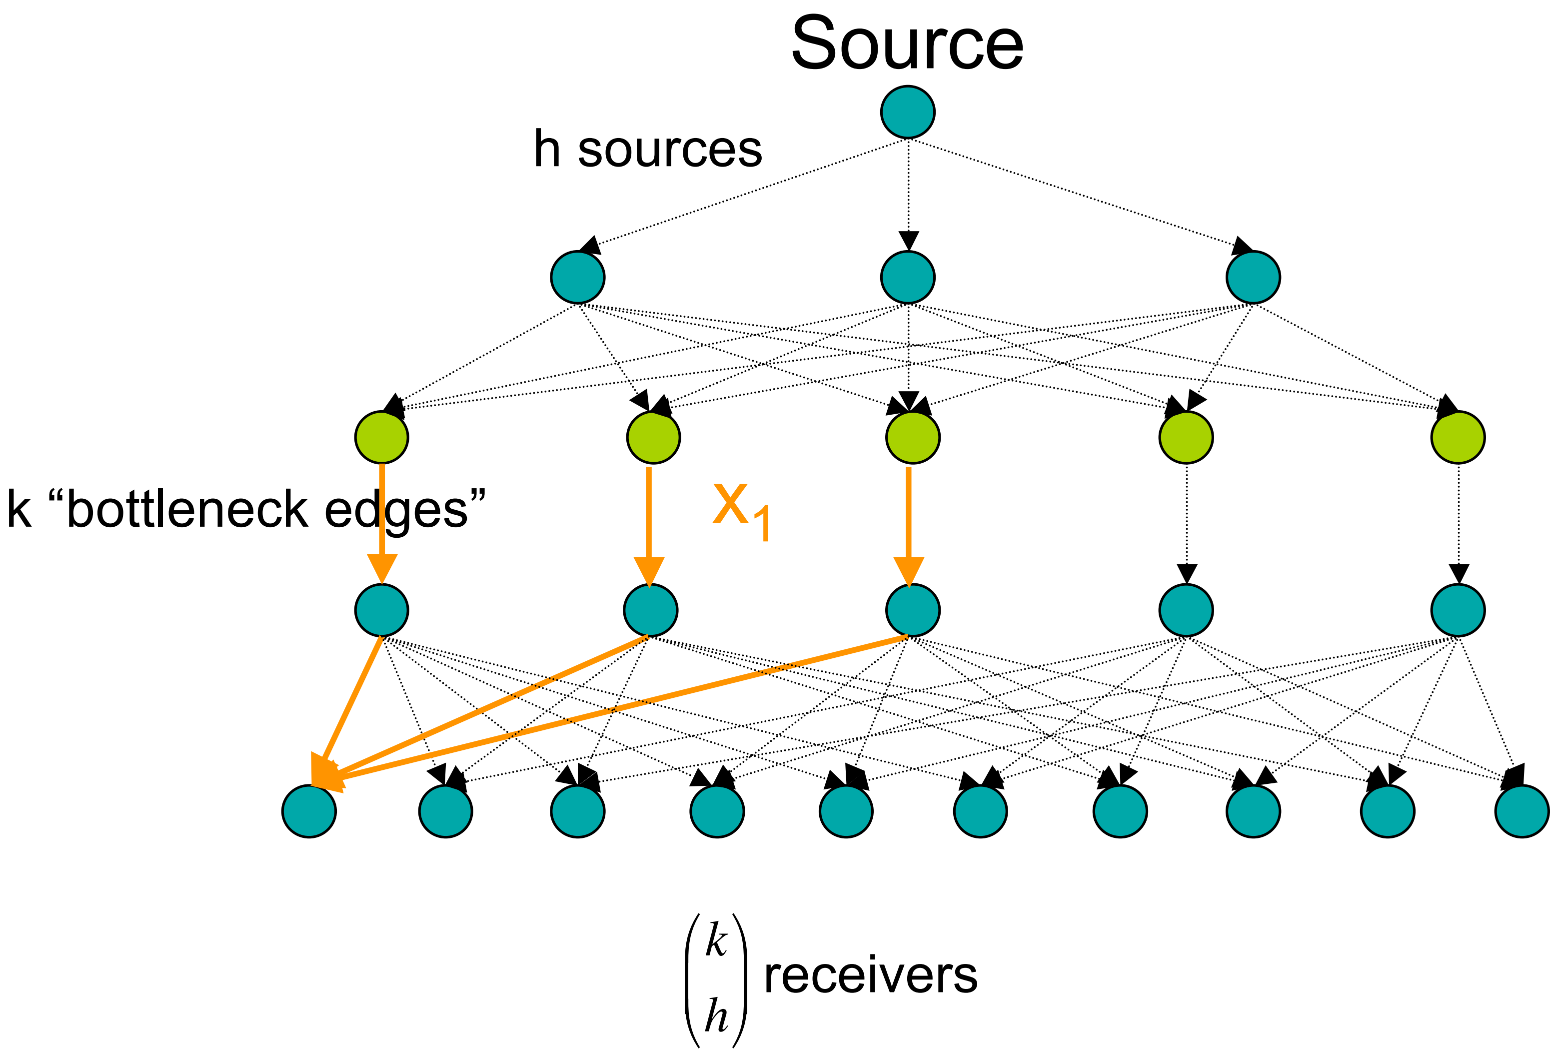

Common throughput benefits

There exist directed graphs where network coding offers throughput benefits as compared to the average throughput proportional to h where h is the number of sources.

(Sanders et al. 2002)

(Agarwal, Charikar 2004)

Theorem

Let a(G,S,R) be the integrality gap of the Steiner tree problem on a directed graph G, with source S and a set R of N receivers. Let b(G,S,R) denote the maximum ratio of network coding throughput versus common throughput. Then

$$b(G,S,R) = a(G,S,R)$$

Throughput benefits

How much do we lose if we don’t use network coding?

What throughput we can get by only using routing?

Common throughput= 1

Average throughput=1.5

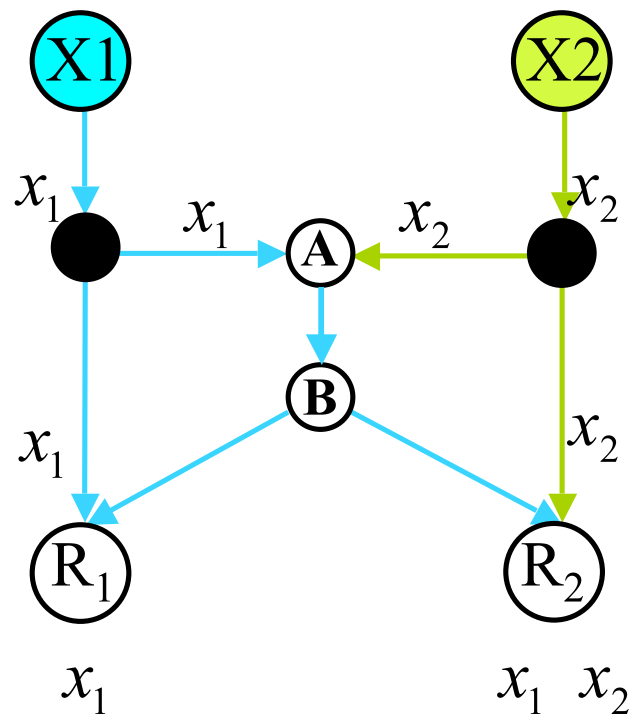

Average Throughput

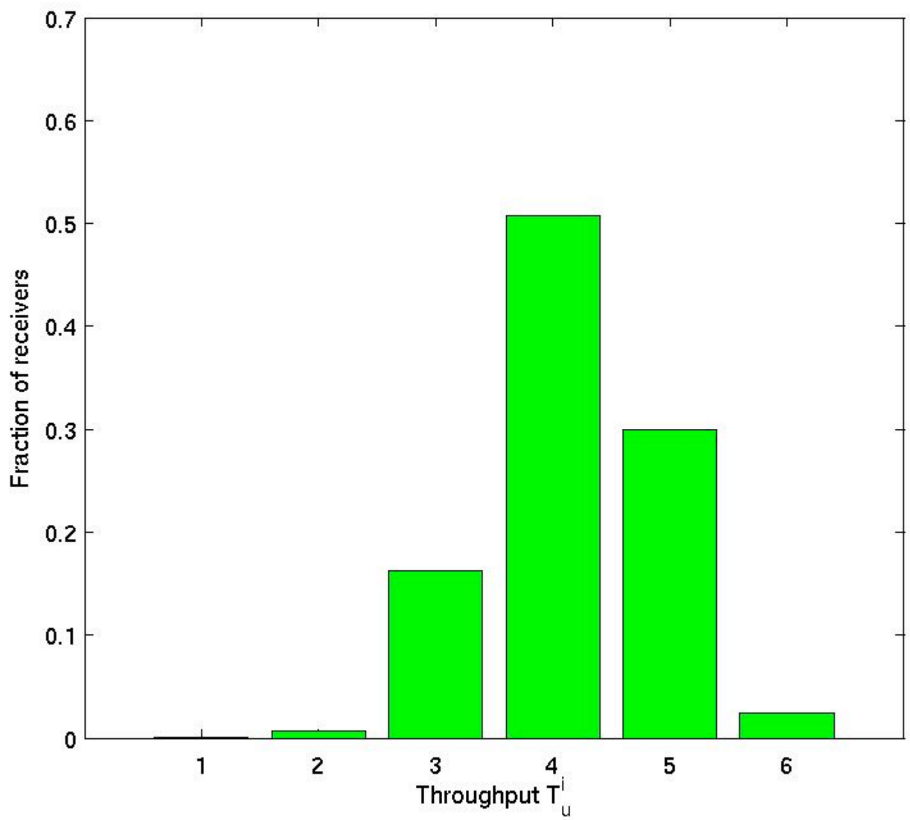

1. Average throughput much higher
2. Concetration around the average
3. “Transform” average to common throughput

Average Throughput: Packing Partial Steiner Trees

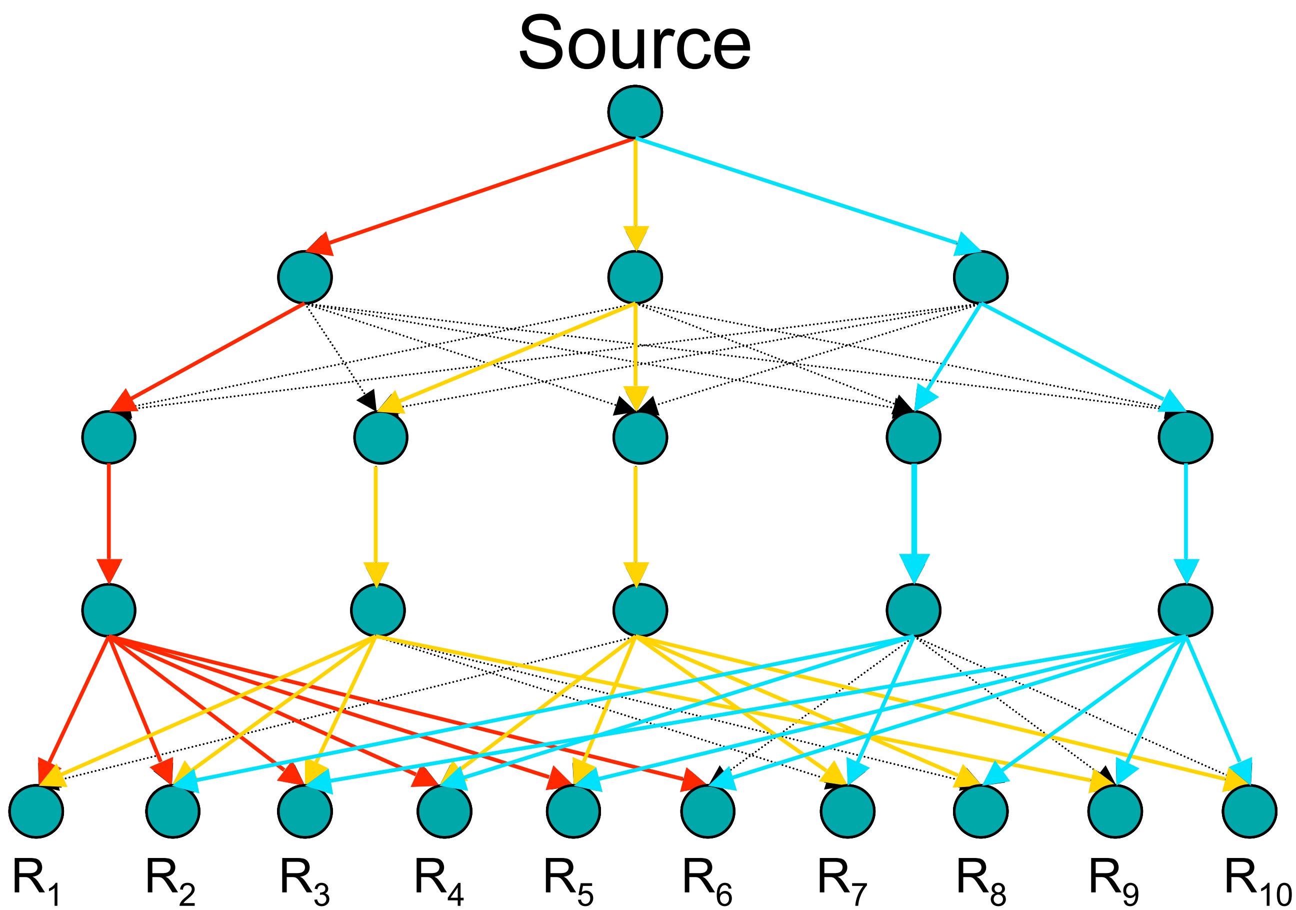

Average throughput benefits

[Chekuri, Fragouli, Soljanin 2005]

Theorem

Let a(G,S,R) be the integrality gap of the Steiner tree problem on a directed graph G, with source S and a set R of N receivers. Let b(G,S,R*) denote the maximum ratio of network coding throughput versus average throughput. Then

$$b(G,S,R^*) > a(G,S,R) / \log N$$

There exist directed graphs where network coding offers throughput benefits as compared to the average throughput proportional to $\sqrt{N}$ where N is the number of receivers.

Network Multicast

Routing

Nodes in the network are only allowed to forward the incoming information flows

Problem of Packing Steiner Trees

- NP-hard
- We do not always achieve rate h to each receiver.

Network Coding

Nodes in the network are allowed to process the incoming information flows

There exist polynomial time algorithms that achieve rate h to each receiver.

Outline

1. Main Theorem in Multicasting

2. Benefits and Requirements

3. Network Code Design

Network Code Design

- Polynomial time algorithms

  (Sanders, Egner, Tolhuizen, Jaggi, Chou, Effros 2003)
- Randomized Algorithms

  (Ho, Medard, Shi, Koetter, Karger 2003)
- Deterministic decentralized algorithms

  (Fragouli, Soljanin 2004)
- Subspace coding

  (Koetter, Kschischang 2007)

Acyclic Networks

In acyclic networks, we can impose a partial order on the edges, so that no edge is visited before all its incoming edges.

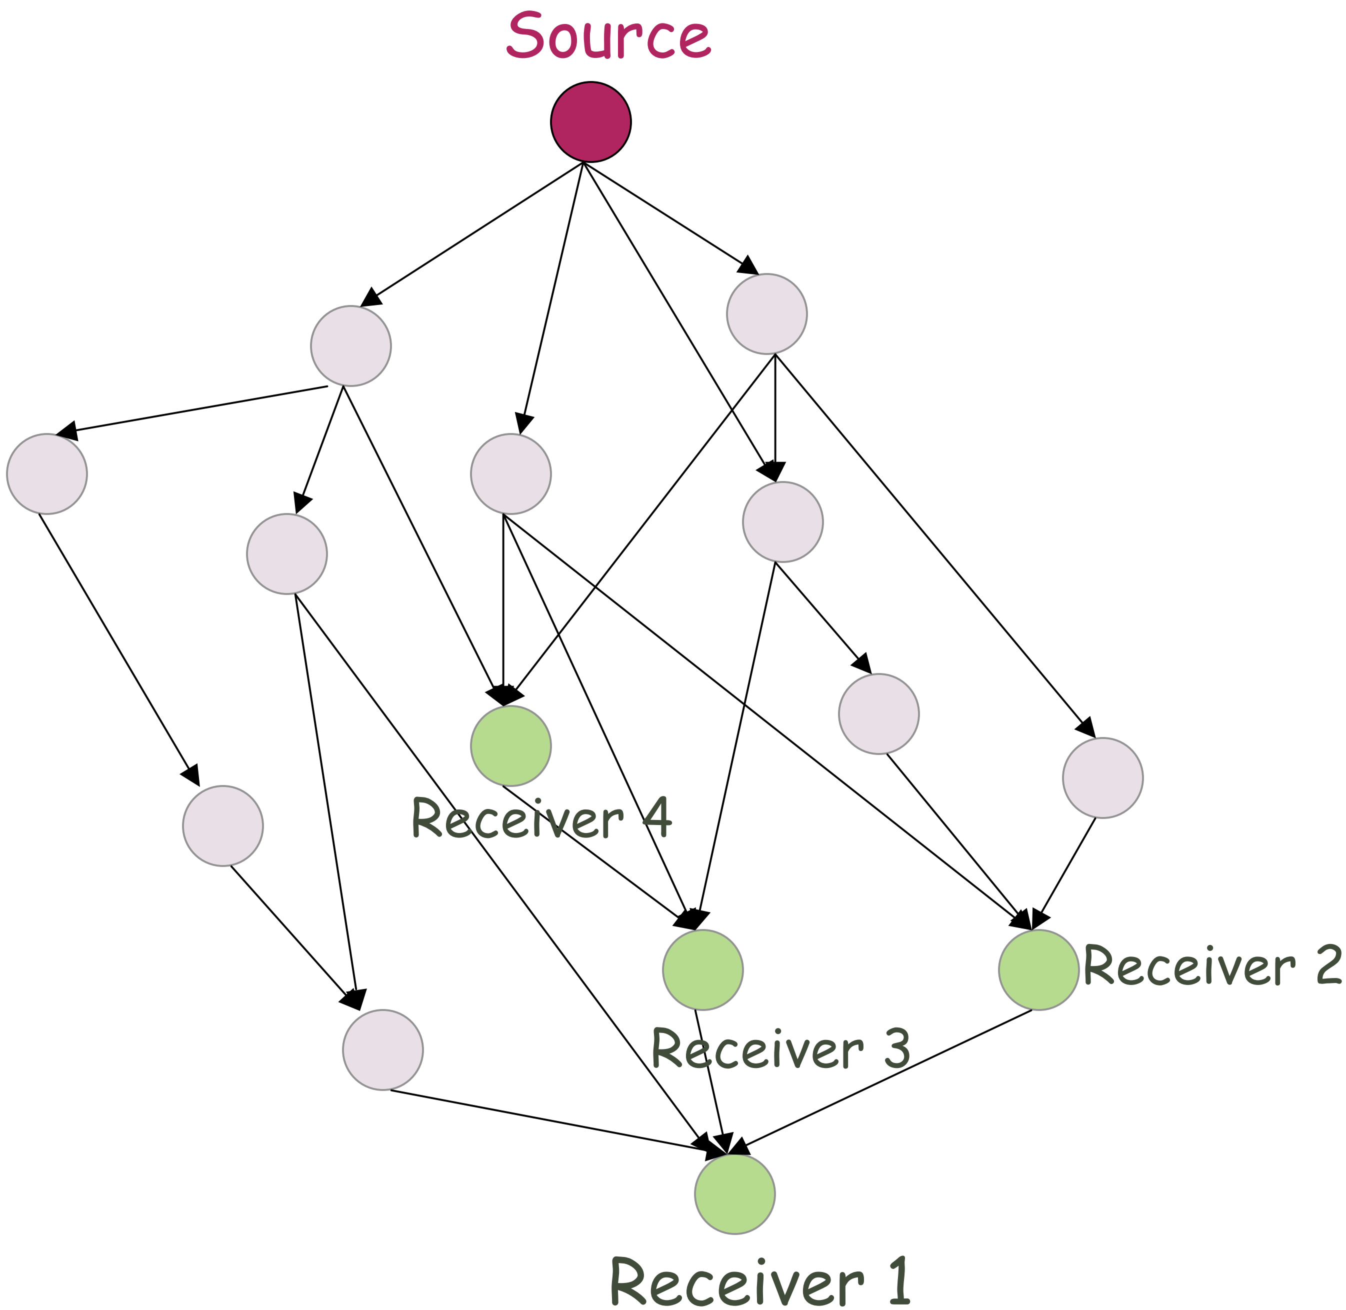
Source
Receiver 4
Receiver 2
Receiver 3
Receiver 1

Acyclic Networks

In acyclic networks, we can impose a partial order on the edges, so that no edge is visited before all its incoming edges.

First common step of all algorithms:

Find paths from the source to each receiver

Paths towards different receivers may overlap

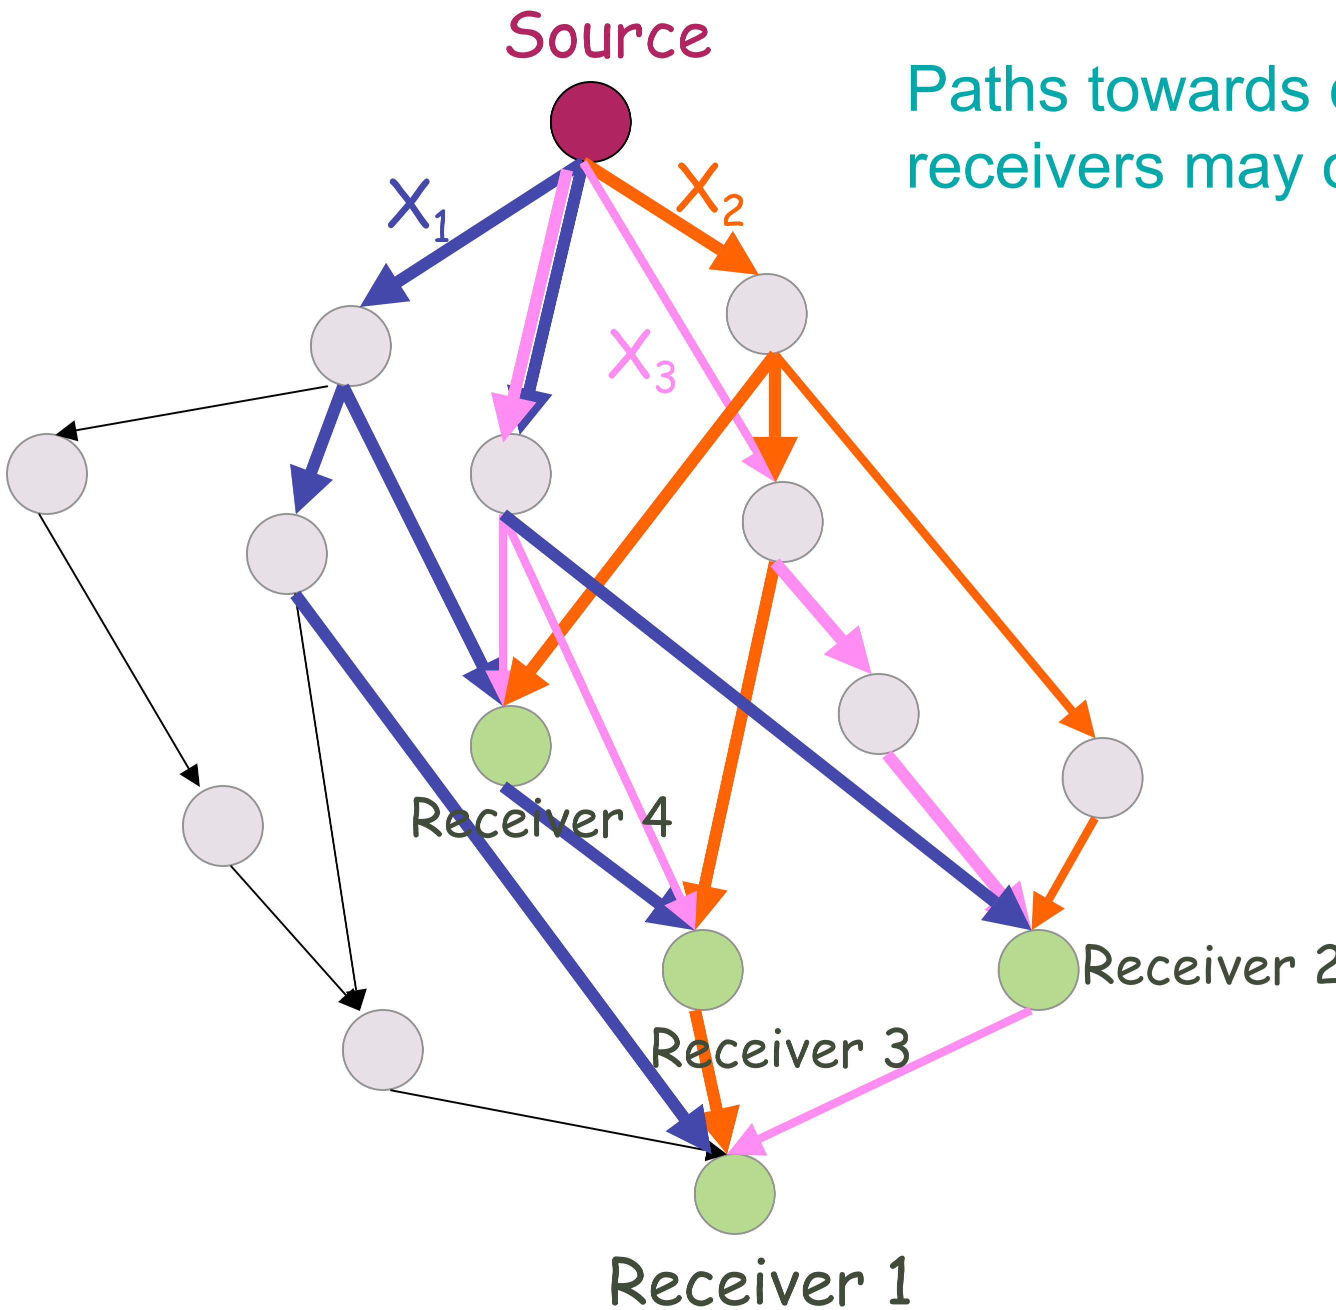

Linear Information Flow (LIF) Algorithm

- Consider an acyclic multicast configuration G=(V,E) with h sources and N receivers.
- Find paths from the source to each receiver.
- Keep for each receiver a set (matrix) of h coding vectors (initially these vectors form the identity matrix) corresponding to the most recently visited edge, on the paths from the source to the receiver.

Initially for each receiver we have the matrix:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

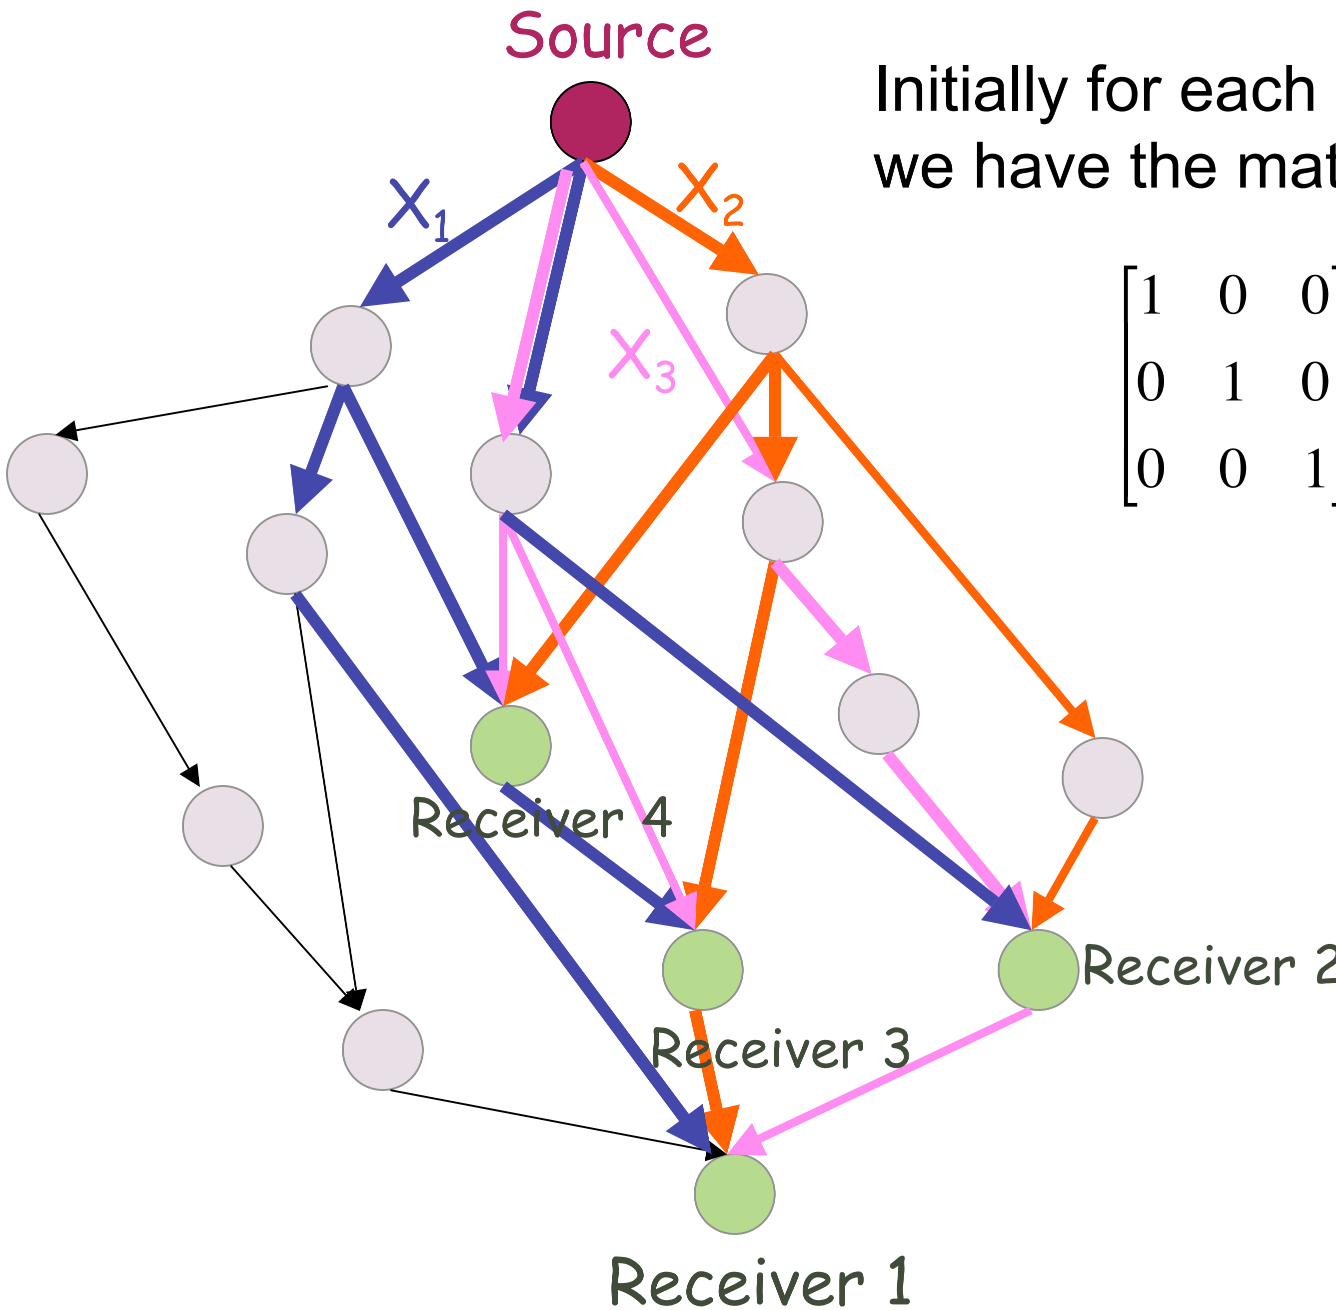

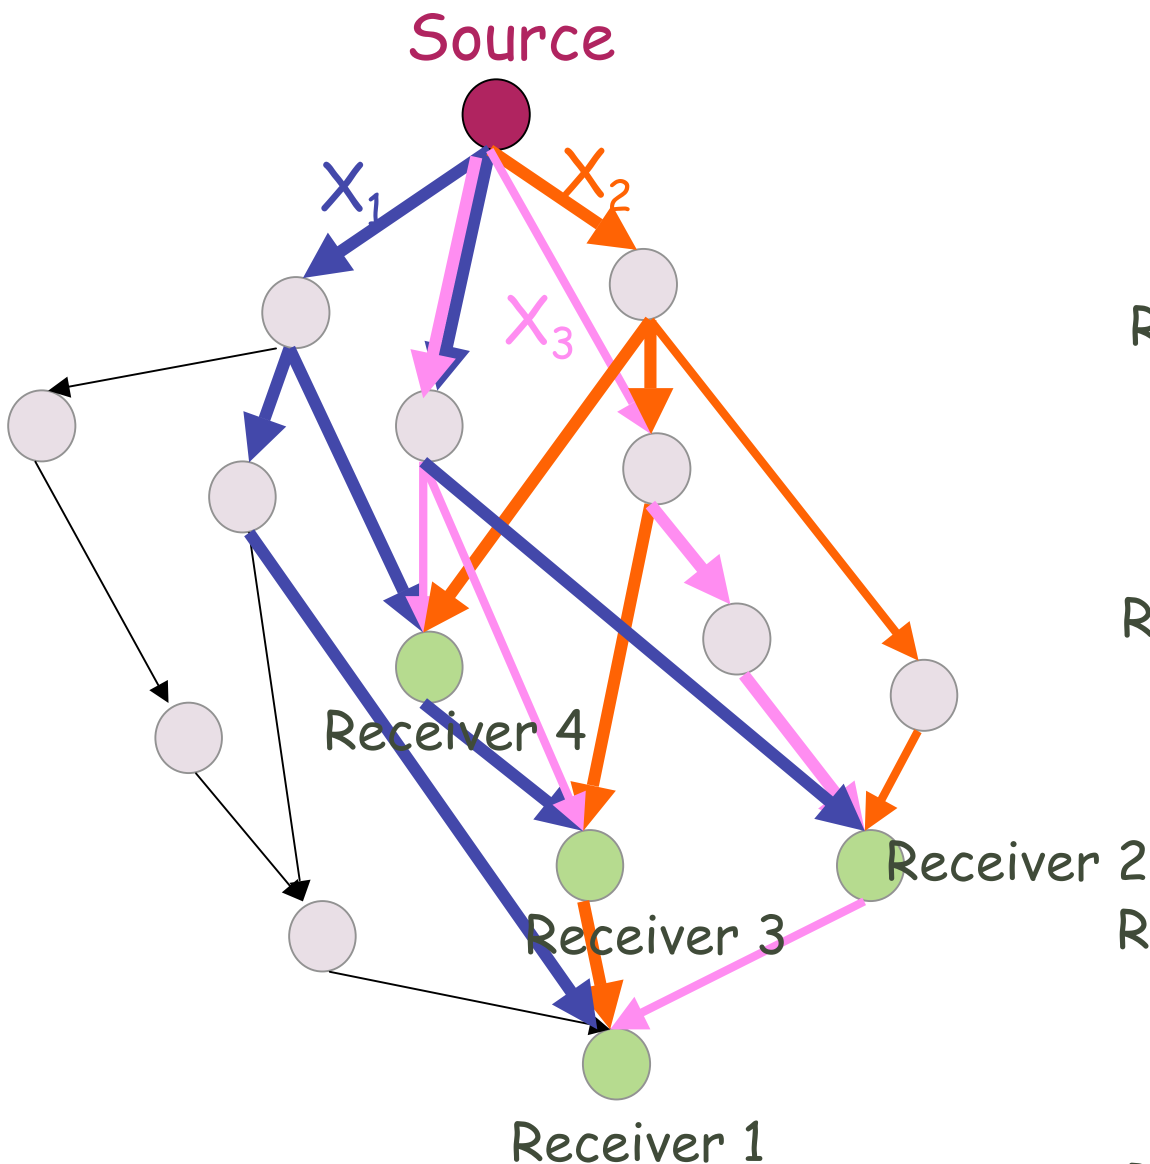

Receiver 1: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 2: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 3: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 4: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

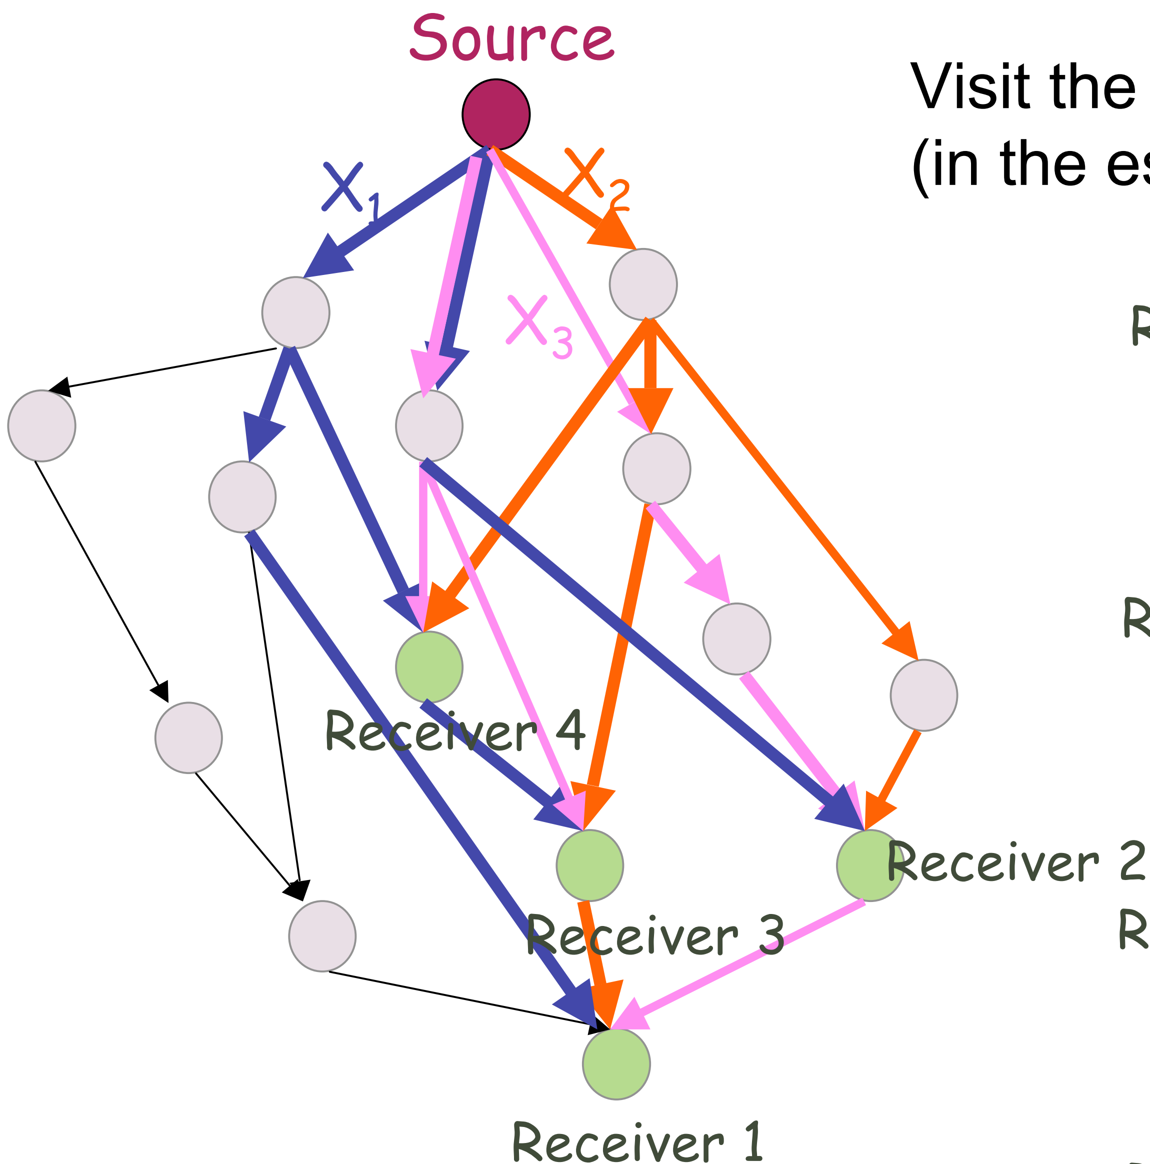

Visit the edges of the graph
(in the established partial order)

Receiver 1: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 2: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 3: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 4: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

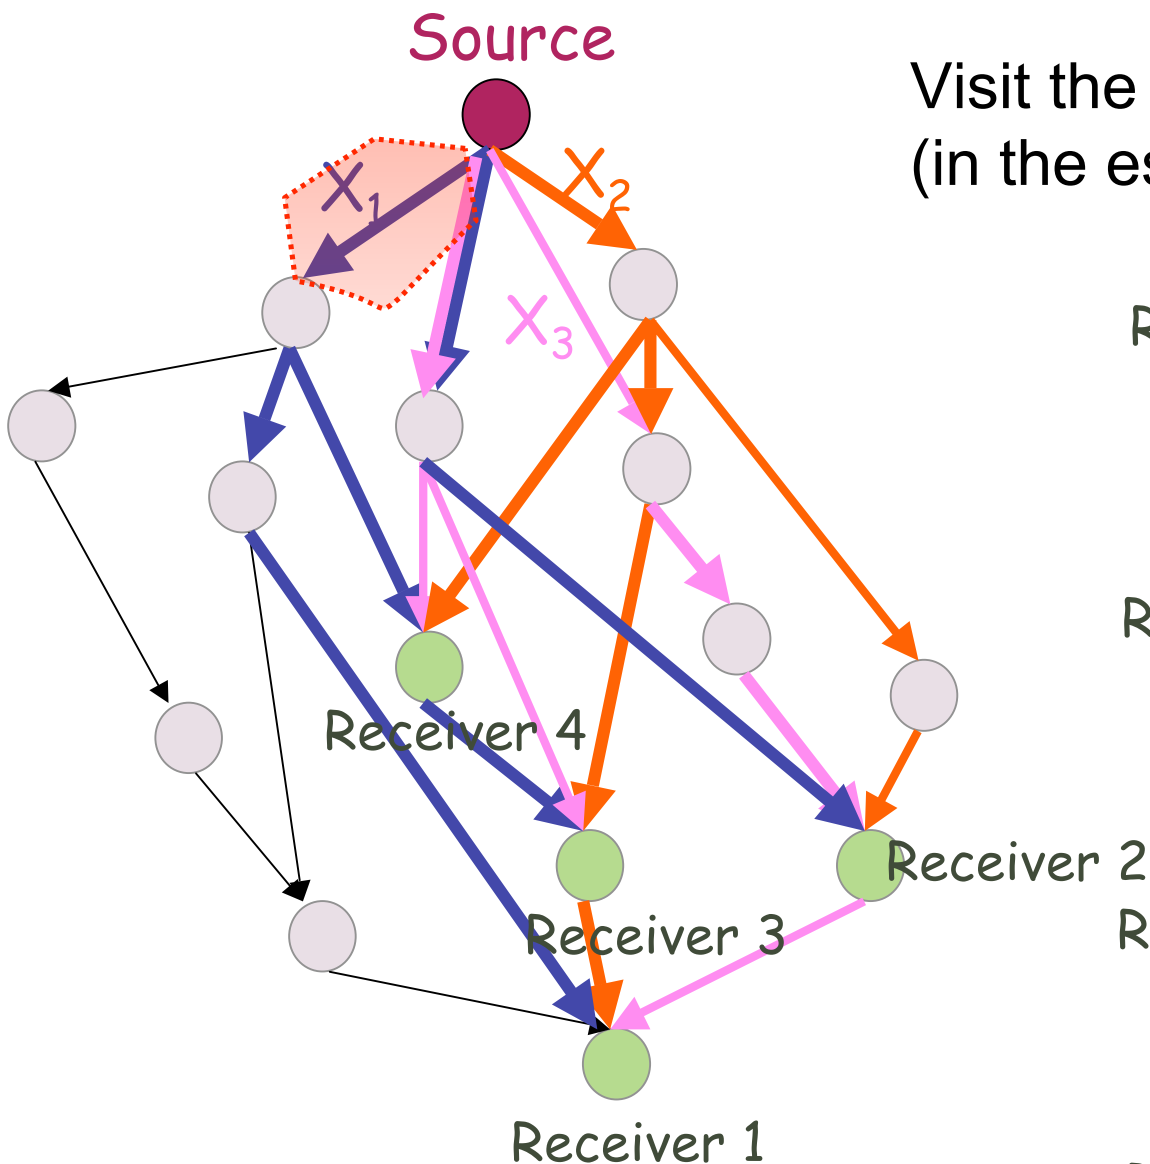

Visit the edges of the graph
(in the established partial order)

Receiver 1: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 2: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 3: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 4: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

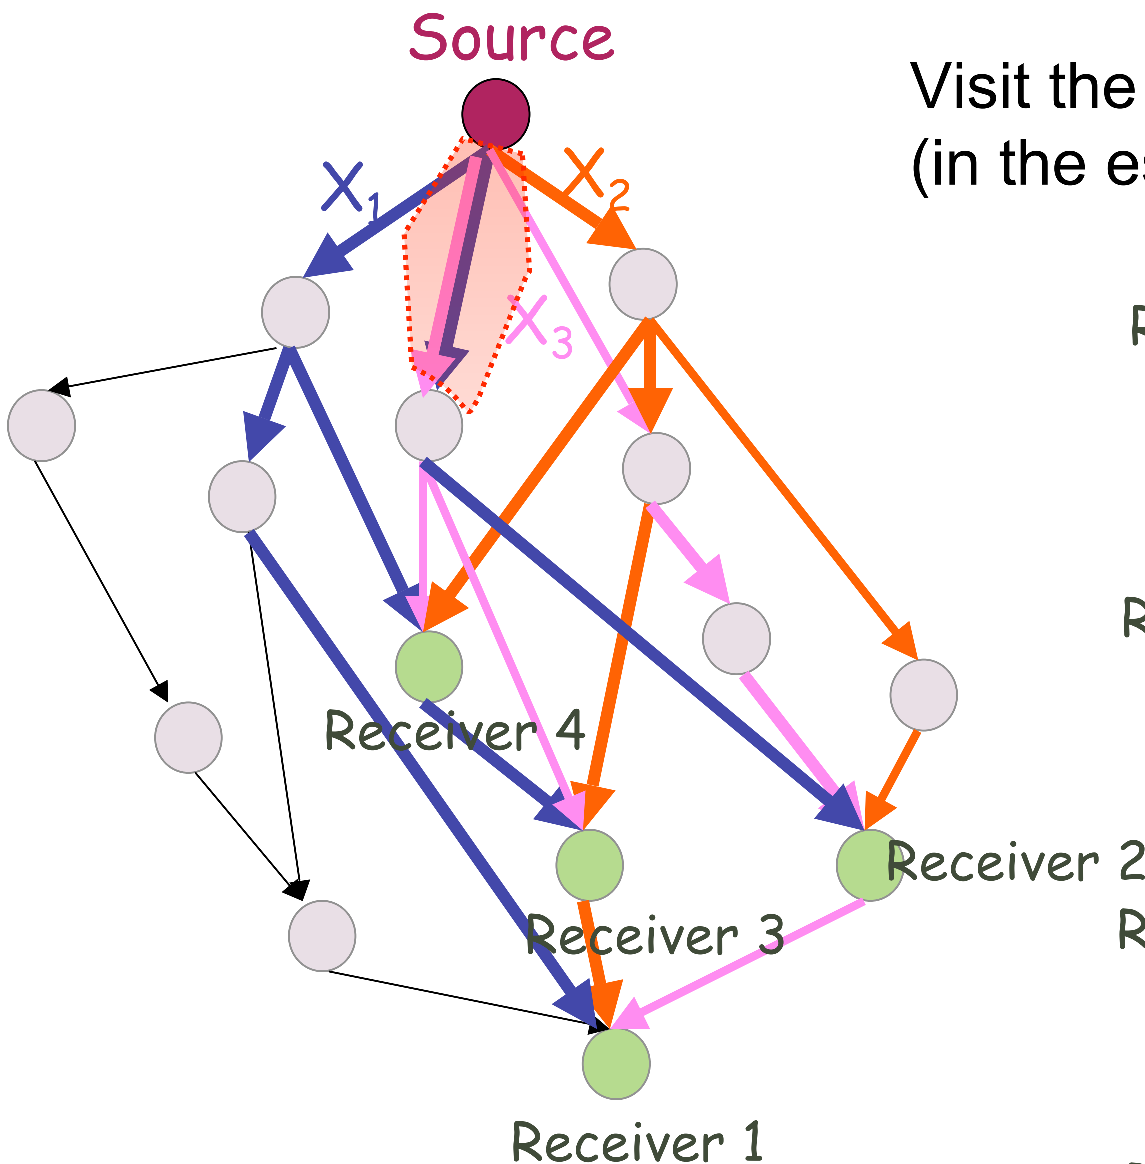

Visit the edges of the graph
(in the established partial order)

Receiver 1: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 2: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 3: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 4: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

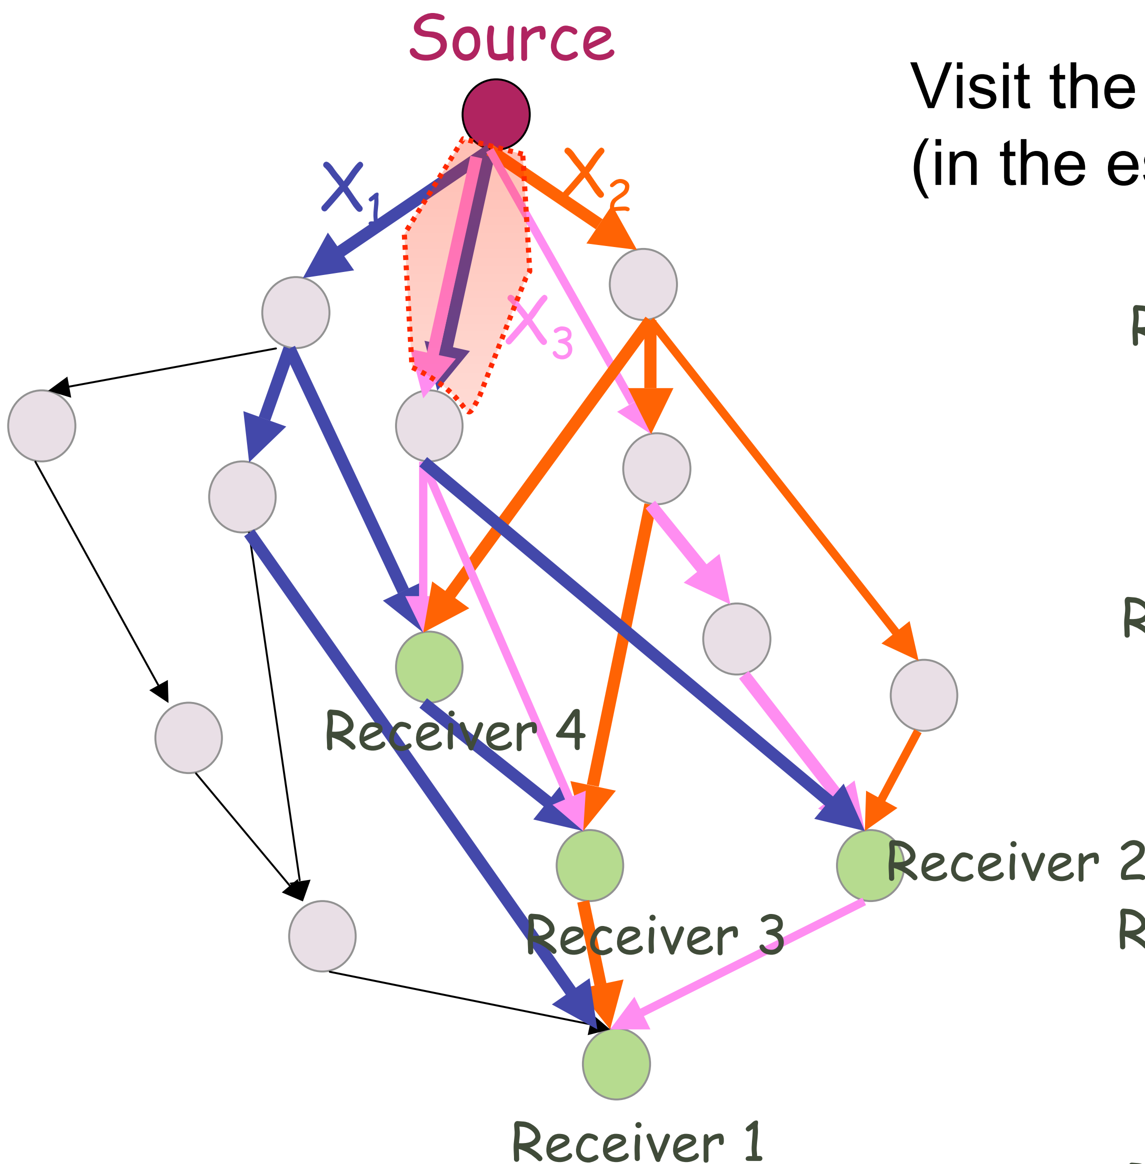

Visit the edges of the graph
(in the established partial order)

Receiver 1: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 2: $\begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$

Receiver 3: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$

Receiver 4: $\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$

Linear Information Flow (LIF) Algorithm

- Consider an acyclic multicast configuration G=(V,E) with h sources and N receivers.
- Find paths from the source to each receiver.
- Keep for each receiver a set (matrix) of h coding vectors (initially these vectors form the identity matrix) corresponding to the most recently visited edge, on the paths from the source to the receiver.
- Sequentially visit the edges of the graph. For each edge e, select a coding vector c(e) such that, all receiver that use this edge in one of their paths, when replacing this vector in their matrix, the matrix remains full rank.
- Such a coding vector exists, provided that the alphabet size is greater than N.

Randomized Algorithms

- Consider an acyclic multicast configuration G=(V,E) with h sources and N receivers.
- Find paths from the source to each receiver.
- At each edge, choose a uniform at random linear combination of the incoming symbols.
- The probability of error goes to zero as the alphabet size increases.

Networks with Cycles

In networks with cycles, we need to introduce delay to guarantee causality:

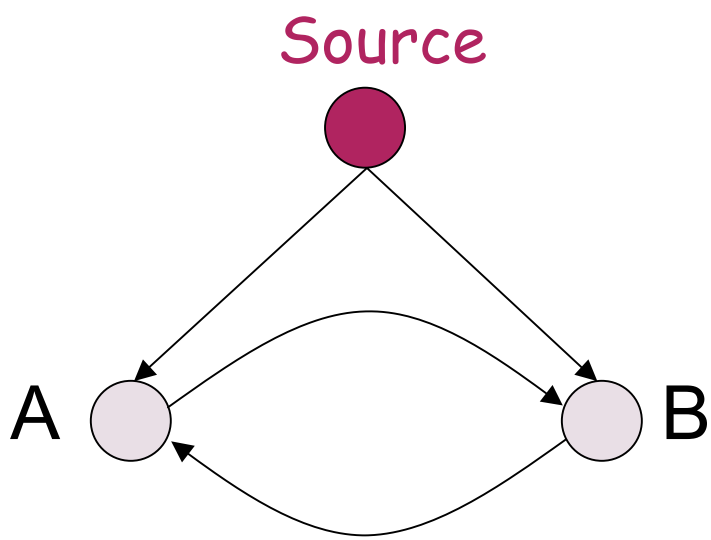

Networks with Cycles

We distinguish between

simple cycles

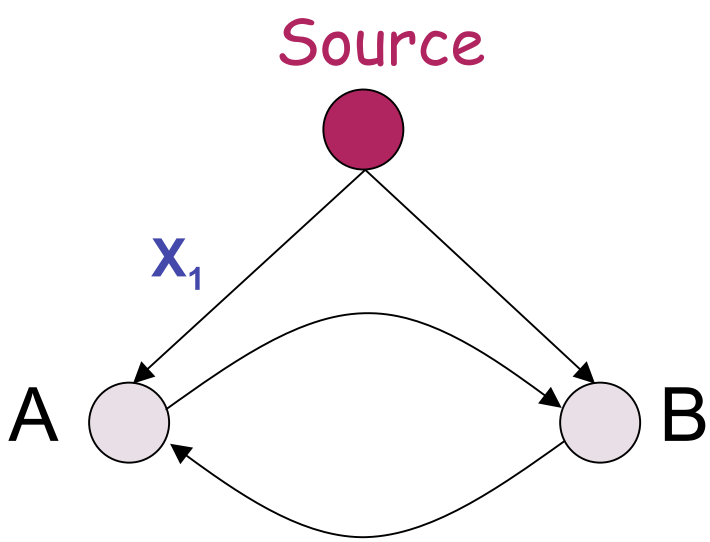

knots

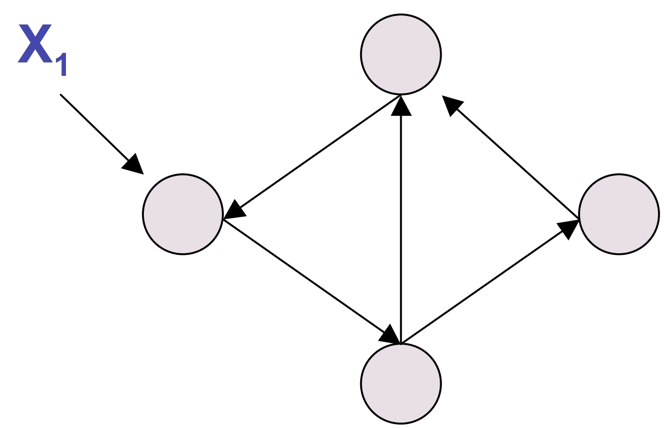

Outline

1. Main Theorem in Multicasting

2. Benefits and Requirements

3. Network Code Design

4. Applications

Applications of Network Coding

- Ad-hoc wireless networks
- Content delivery in P2P networks
- Network tomography
- Sensor networks
- Security
- Chip design
- …..

Applications of Network Coding

- Ad-hoc wireless networks
- Content delivery in P2P networks
- Network tomography
- Sensor networks
- Security
- Chip design
- .....

Wireless Networks

Benefits: energy efficiency, delay, wireless bandwidth

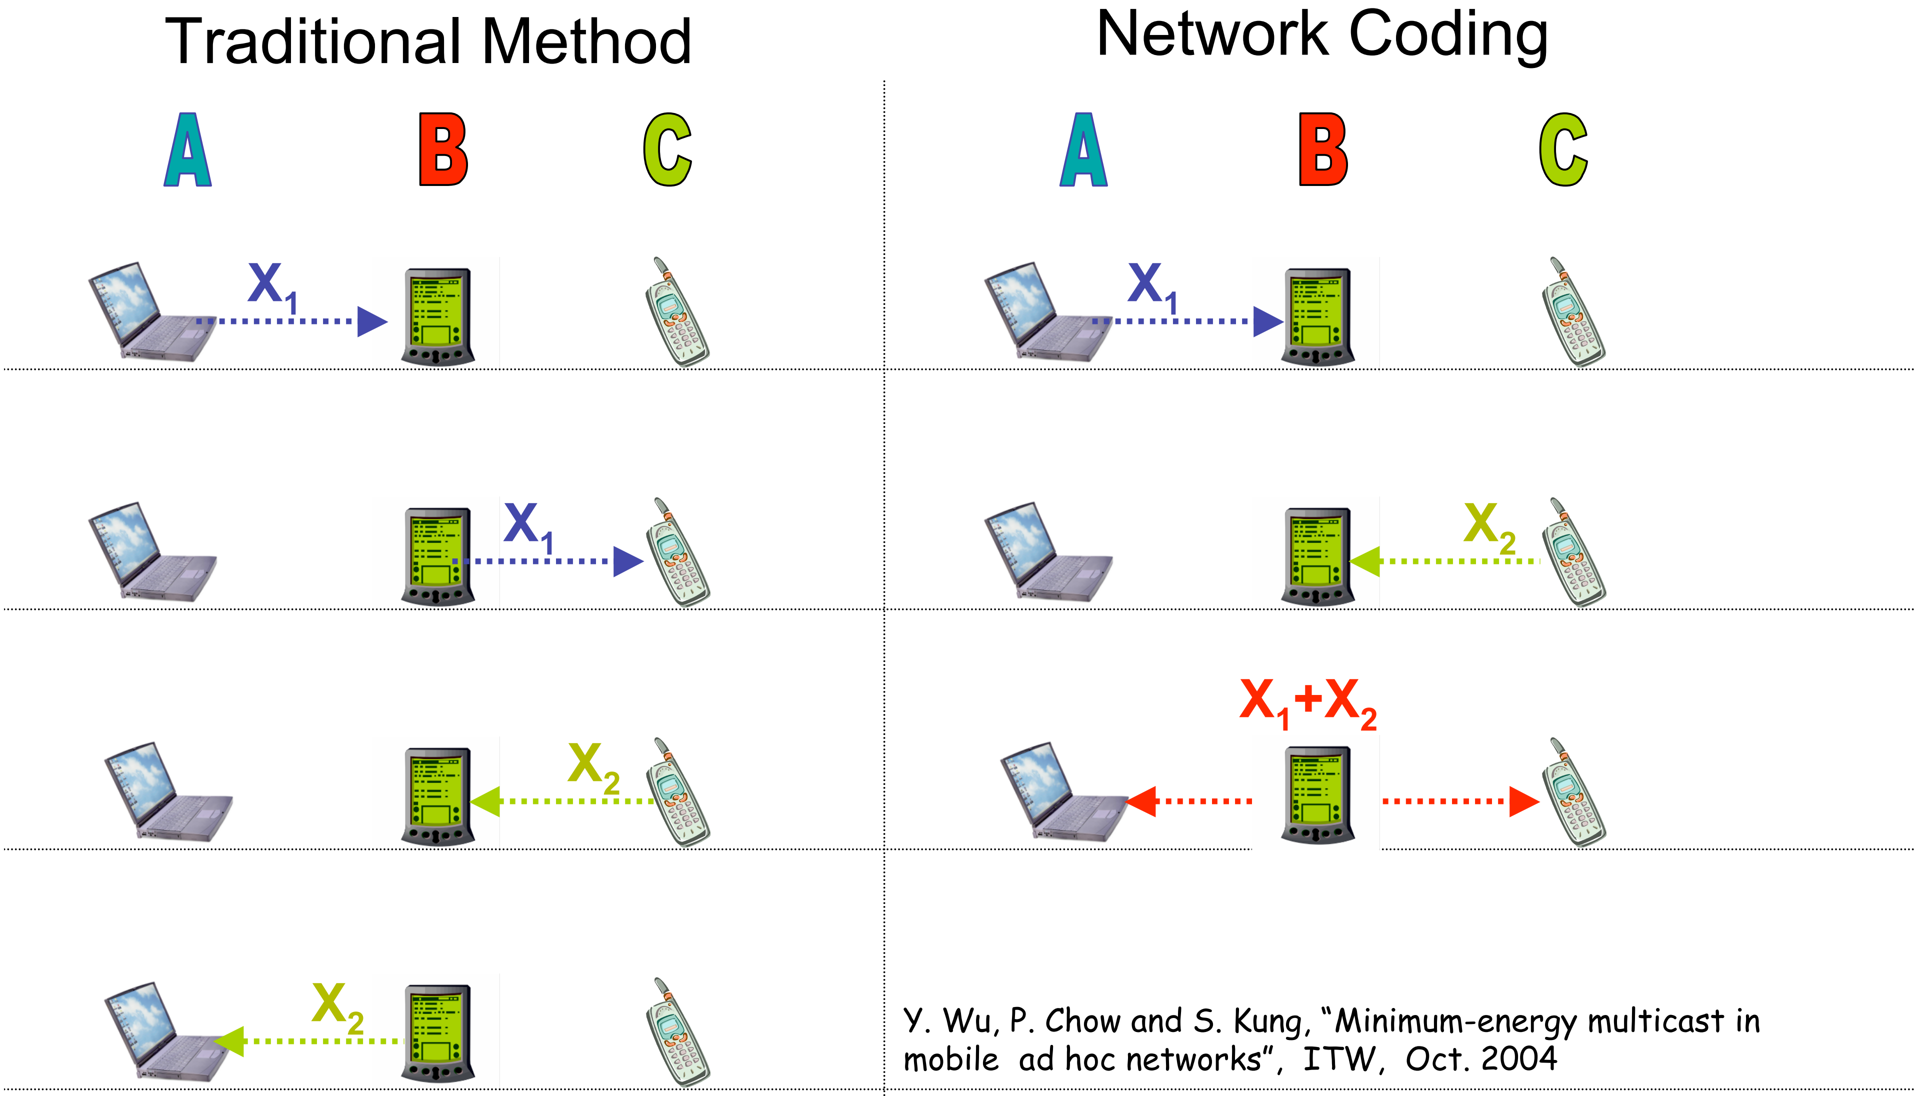

Y. Wu, P. Chow and S. Kung, "Minimum-energy multicast in mobile ad hoc networks", ITW, Oct. 2004

Energy-Efficient Broadcasting in Wireless Ad-hoc Networks

(Widmer, Fragouli, Le Boudec 2005)

Consider an ad-hoc wireless network, where

- all nodes are sources,
- every node needs to receive all sources, and
- each node can broadcast information to its closest neighbors.

Application: Discovery mechanisms at the network or application layer

Energy-Efficient Broadcasting in Wireless Ad-hoc Networks

Circular Network

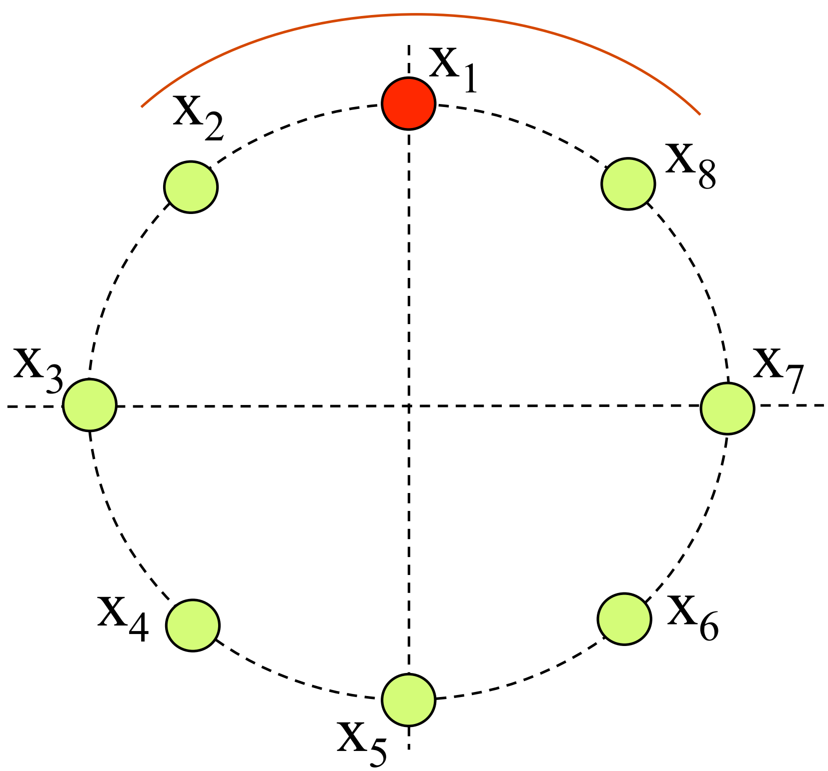

Square Grid

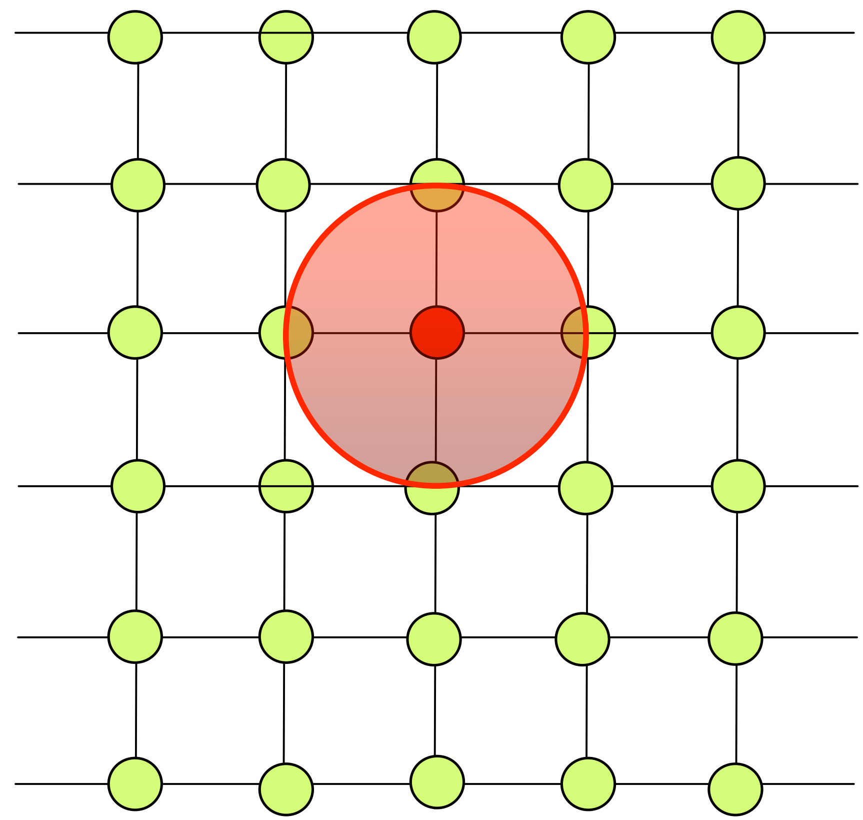

Theorem

(Widmer, Fragouli, Le Boudec 2005)

Let Nc be the total number of transmissions per information unit required with network coding and R the total number of transmissions required with routing. Then:

Network coding uses the smallest possible number of transmissions. Moreover,

For the circular network

$$R \geq 2Nc$$

For the square grid

$$R \geq \frac{4}{3}Nc$$

Elements of the Proof

1. Circular Network: without network coding

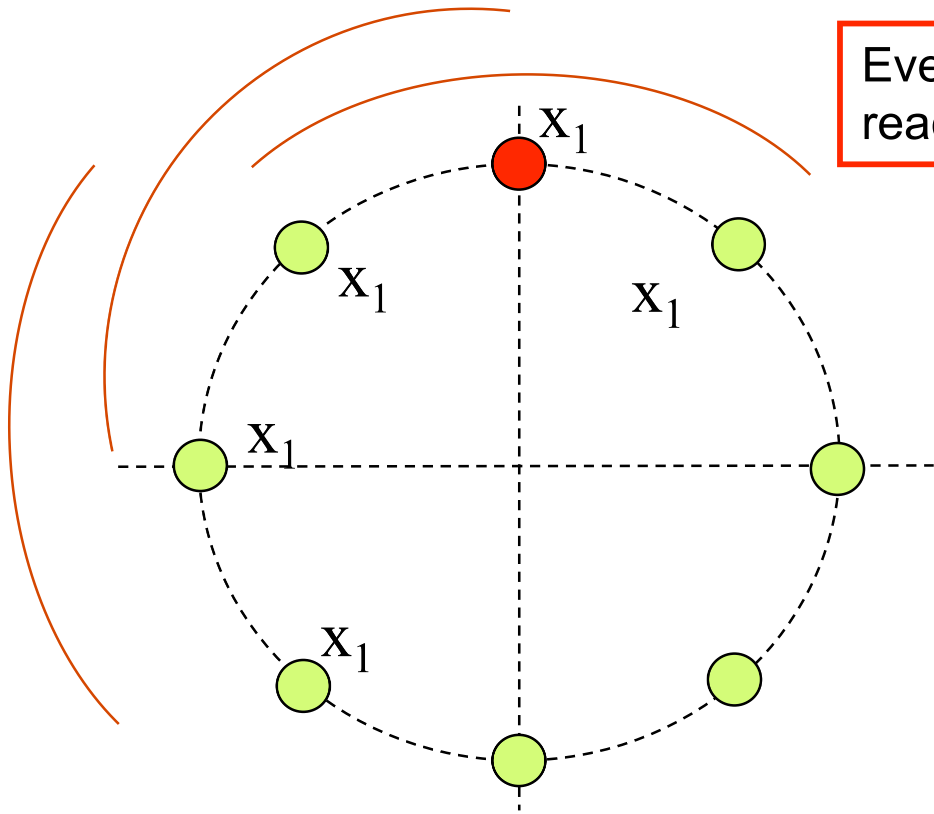

Every transmission reaches one new neighbor

2. Circular Network with network coding

Every transmission reaches two new neighbors

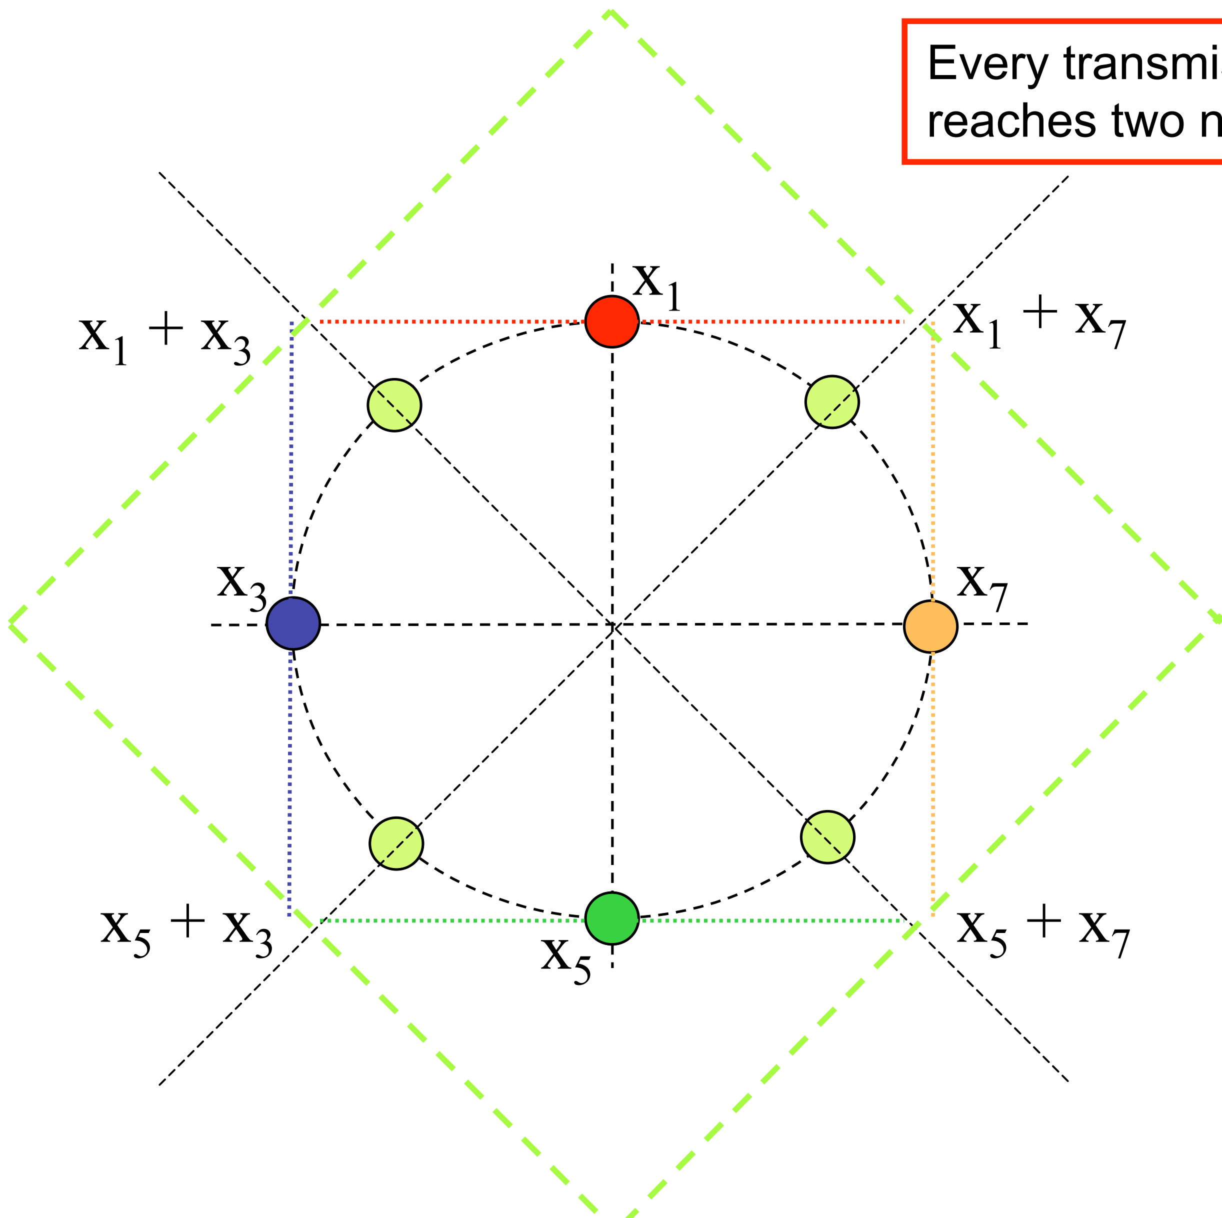

3. Square Grid: **R=4 N /3**

Forwarding: every transmission reaches three new neighbors

Network Coding: every transmission reaches four new neighbors

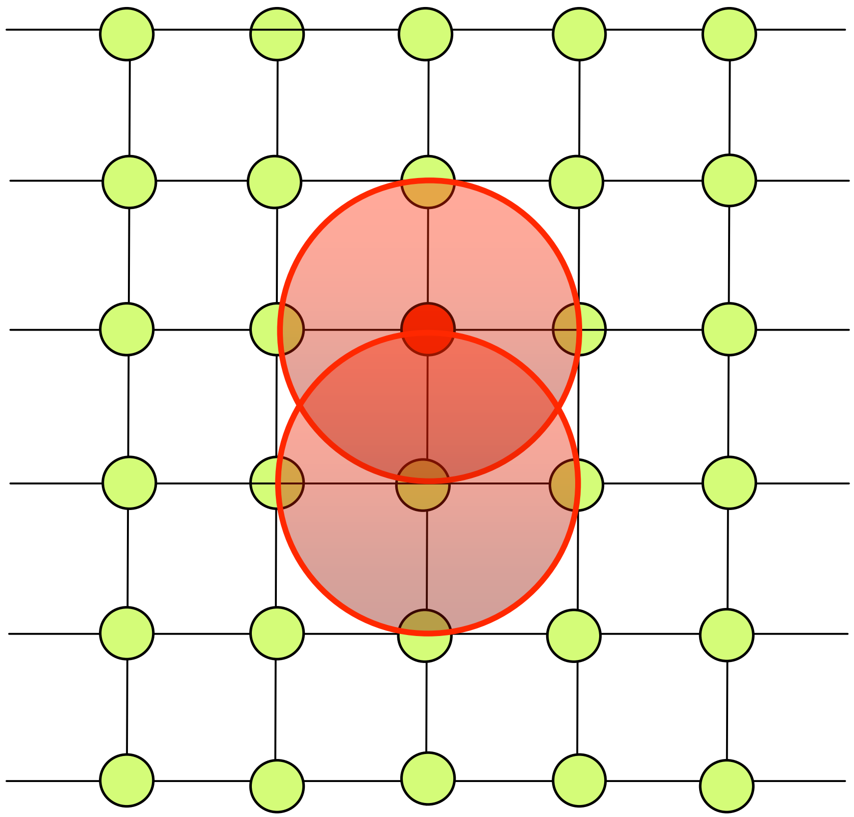

Decentralized Algorithms

Square Grid: Each broadcast transmission brings information to a constant number C of neighbors.

Each node broadcasts a packet (a random linear combination of whatever he has received in the past) as soon as he receives C new packets.

Random Network: Number of neighbors not constant.

Rebroadcast a new packet with probability d.

Simulation Results

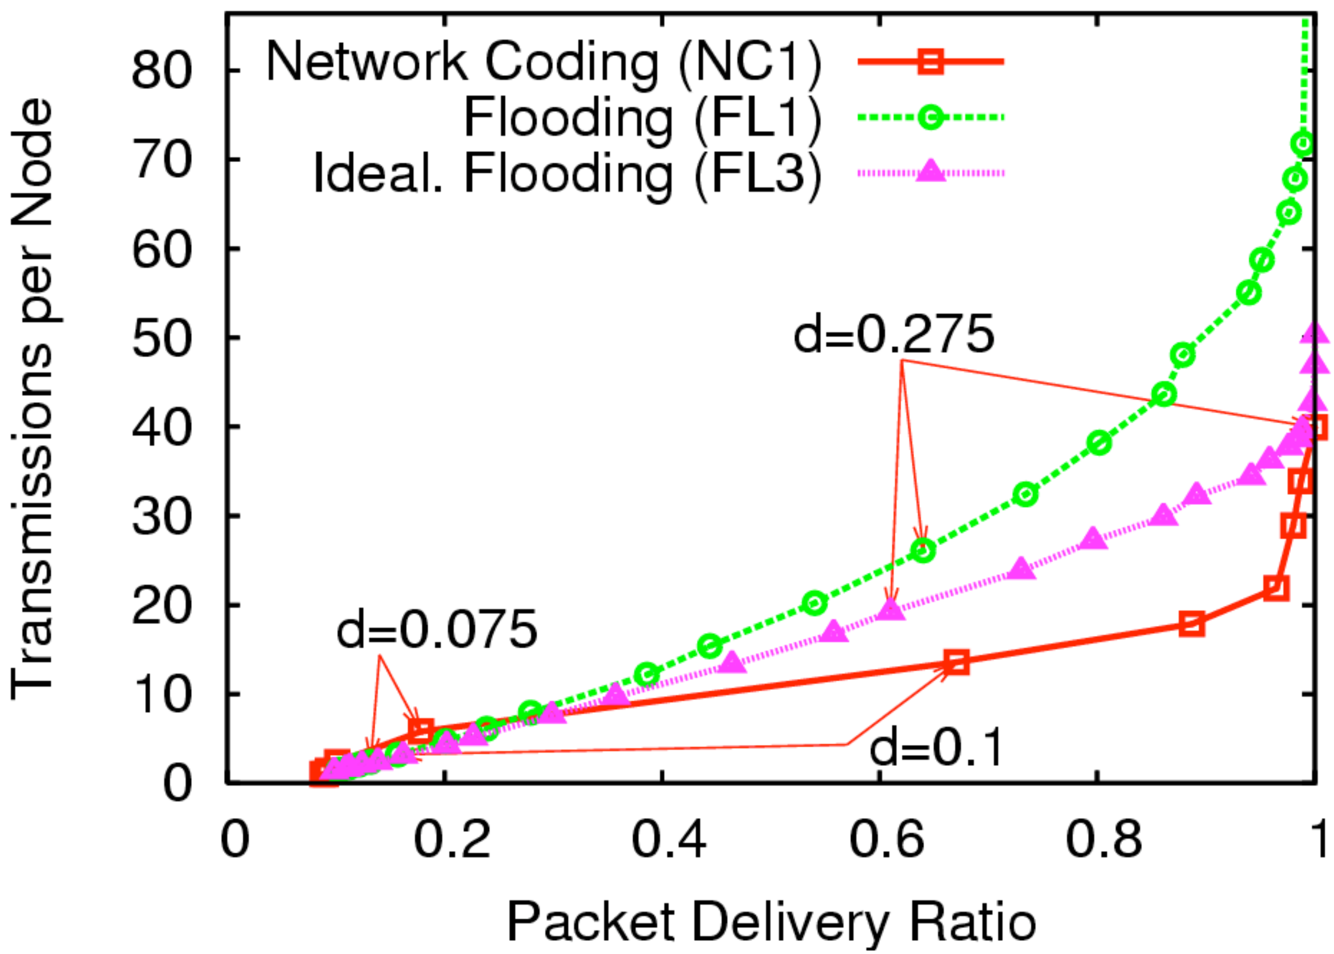

We saw that:

Network coding offers a constant factor of benefits in terms of energy efficiency over fixed wireless networks.

Other benefits?

S. Zhang, S. Liew and P. Lan, "Physical Layer Network Coding", ITW, Oct. 2006 and MobiCom 2006

Network Coding

Physical Layer Network Coding

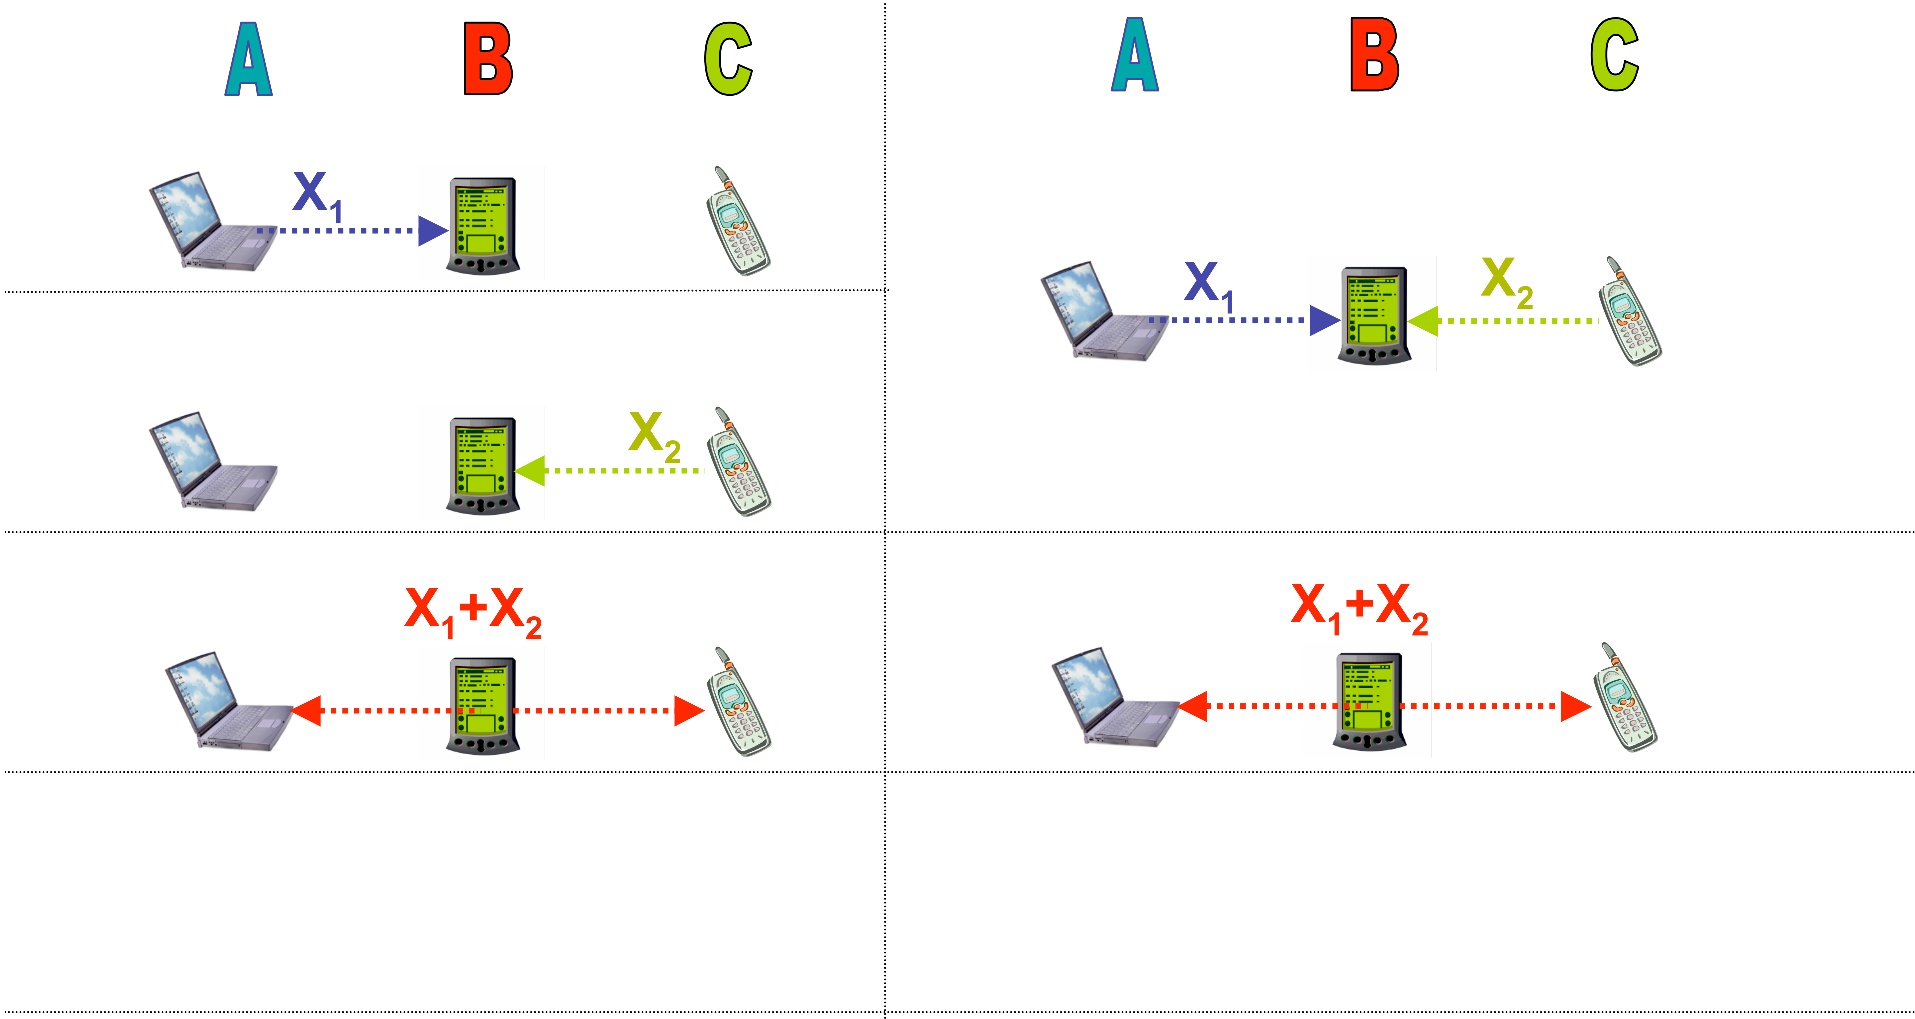

Challenges

(similar to distributed space-time coding)

- Quantization errors
- Synchronization
- ….

Physical Layer Network Coding

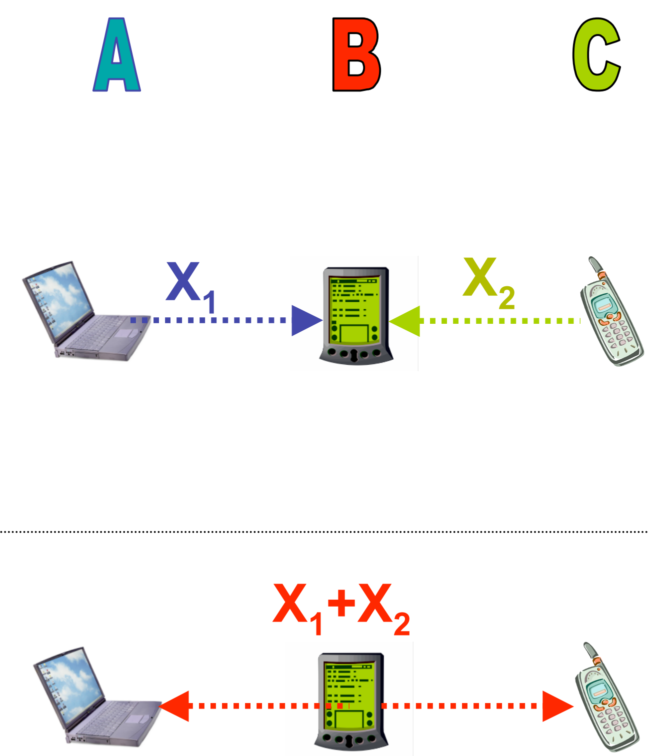

Network coding offers benefits in terms of

1) energy efficiency

Other benefits?

COPE

Katti, Rahul, Hu, Katabi, Medard, Crowcroft, SigComm 2006

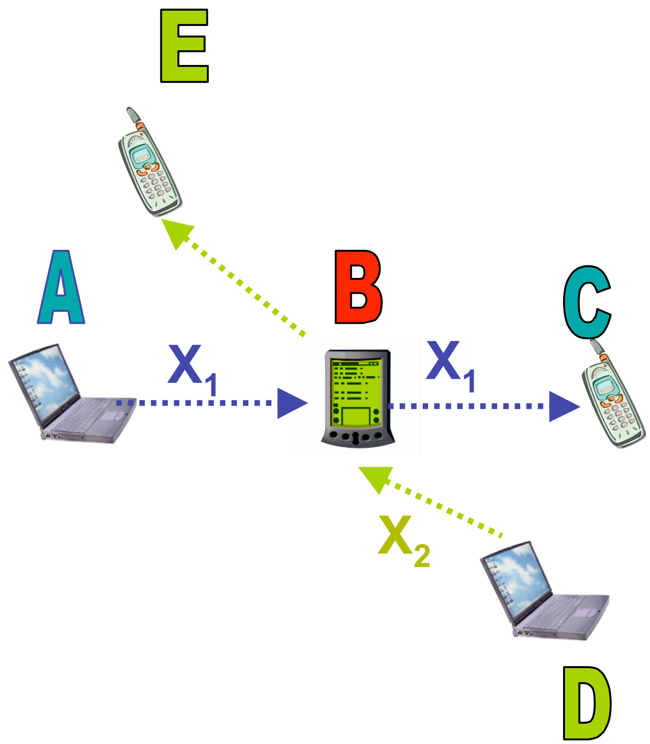

COPE

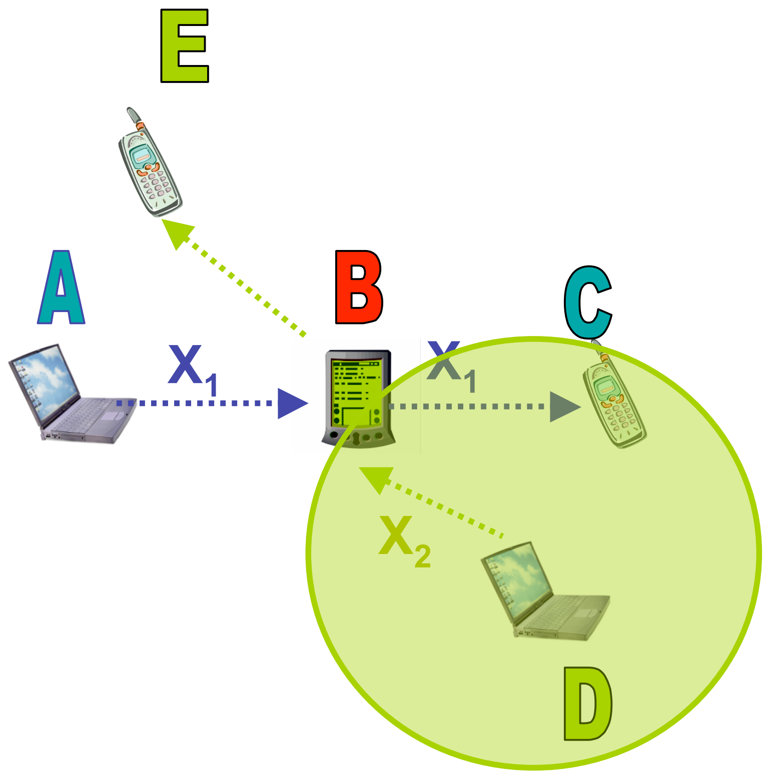

COPE

SigCom 2006

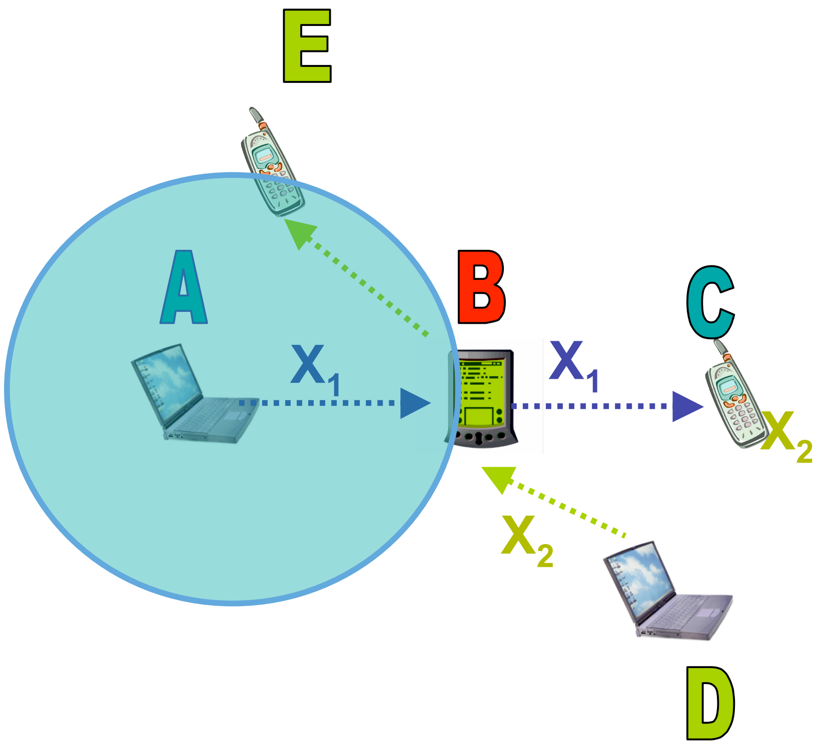

COPE

Katti, Rahul, Hu, Katabi, Medard, Crowcroft, SigComm 2006

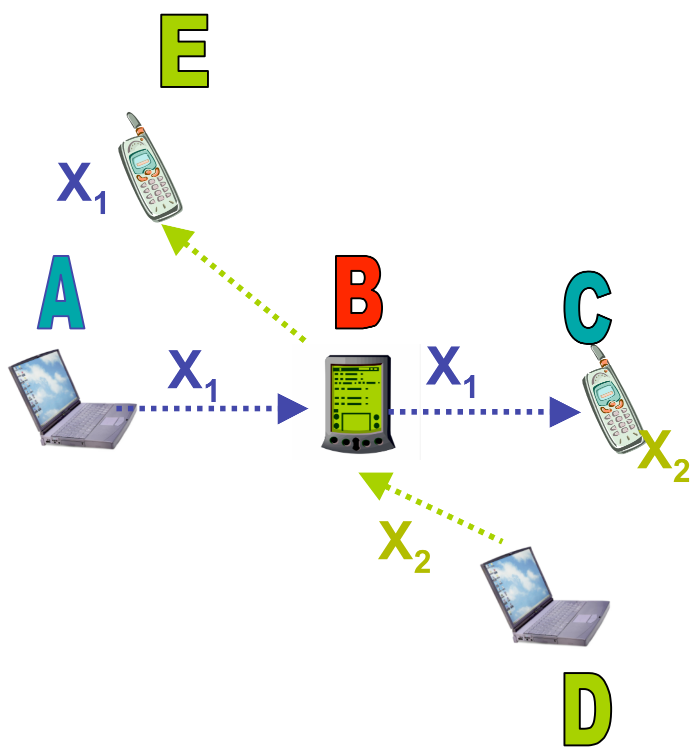

COPE

Katti, Rahul, Hu, Katabi, Medard, Crowcroft, SigComm 2006

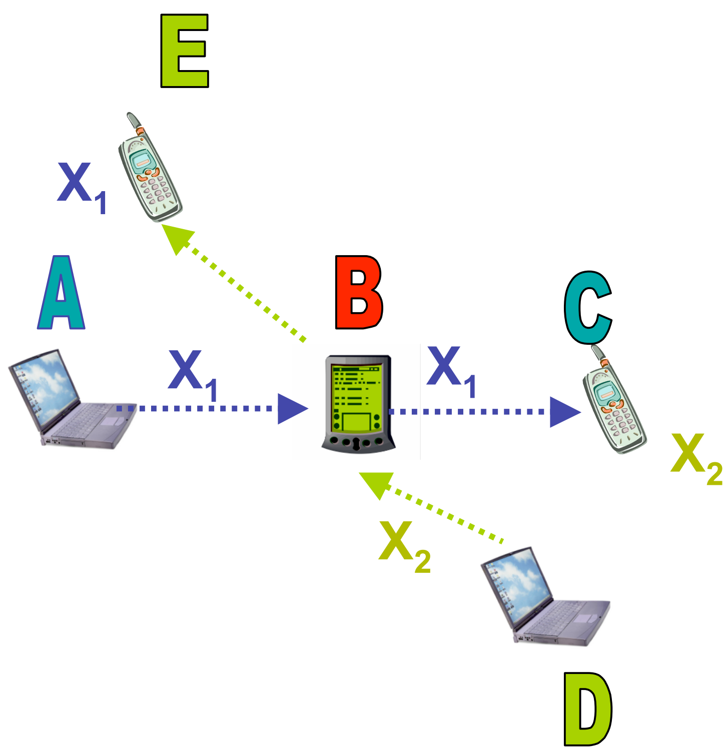

Having node B transmit once instead of twice, makes traffic more uniform.

Network coding offers benefits in terms of

1) energy efficiency
2) making traffic more uniform

Other benefits?

Network coding offers benefits in terms of

1) energy efficiency

2) making traffic more uniform

Significant benefits over dynamically changing environments

We can immediately see why from the very first proof of the main theorem in network coding!

(Ahlswede, Cai, Li, Yeung 2000)

Main Theorem in Network Coding

(Ahlswede, Cai, Li, Yeung 2000)

Main Elements:

- The network is represented as a directed graph G=(V,E) with unit capacity edges.
- A source produces information at a rate R.
- The min-cut from the source to each receiver is h.

It is possible to reliably send to each receiver rate
R<h
provided intermediate network nodes are allowed to combine their incoming information flows

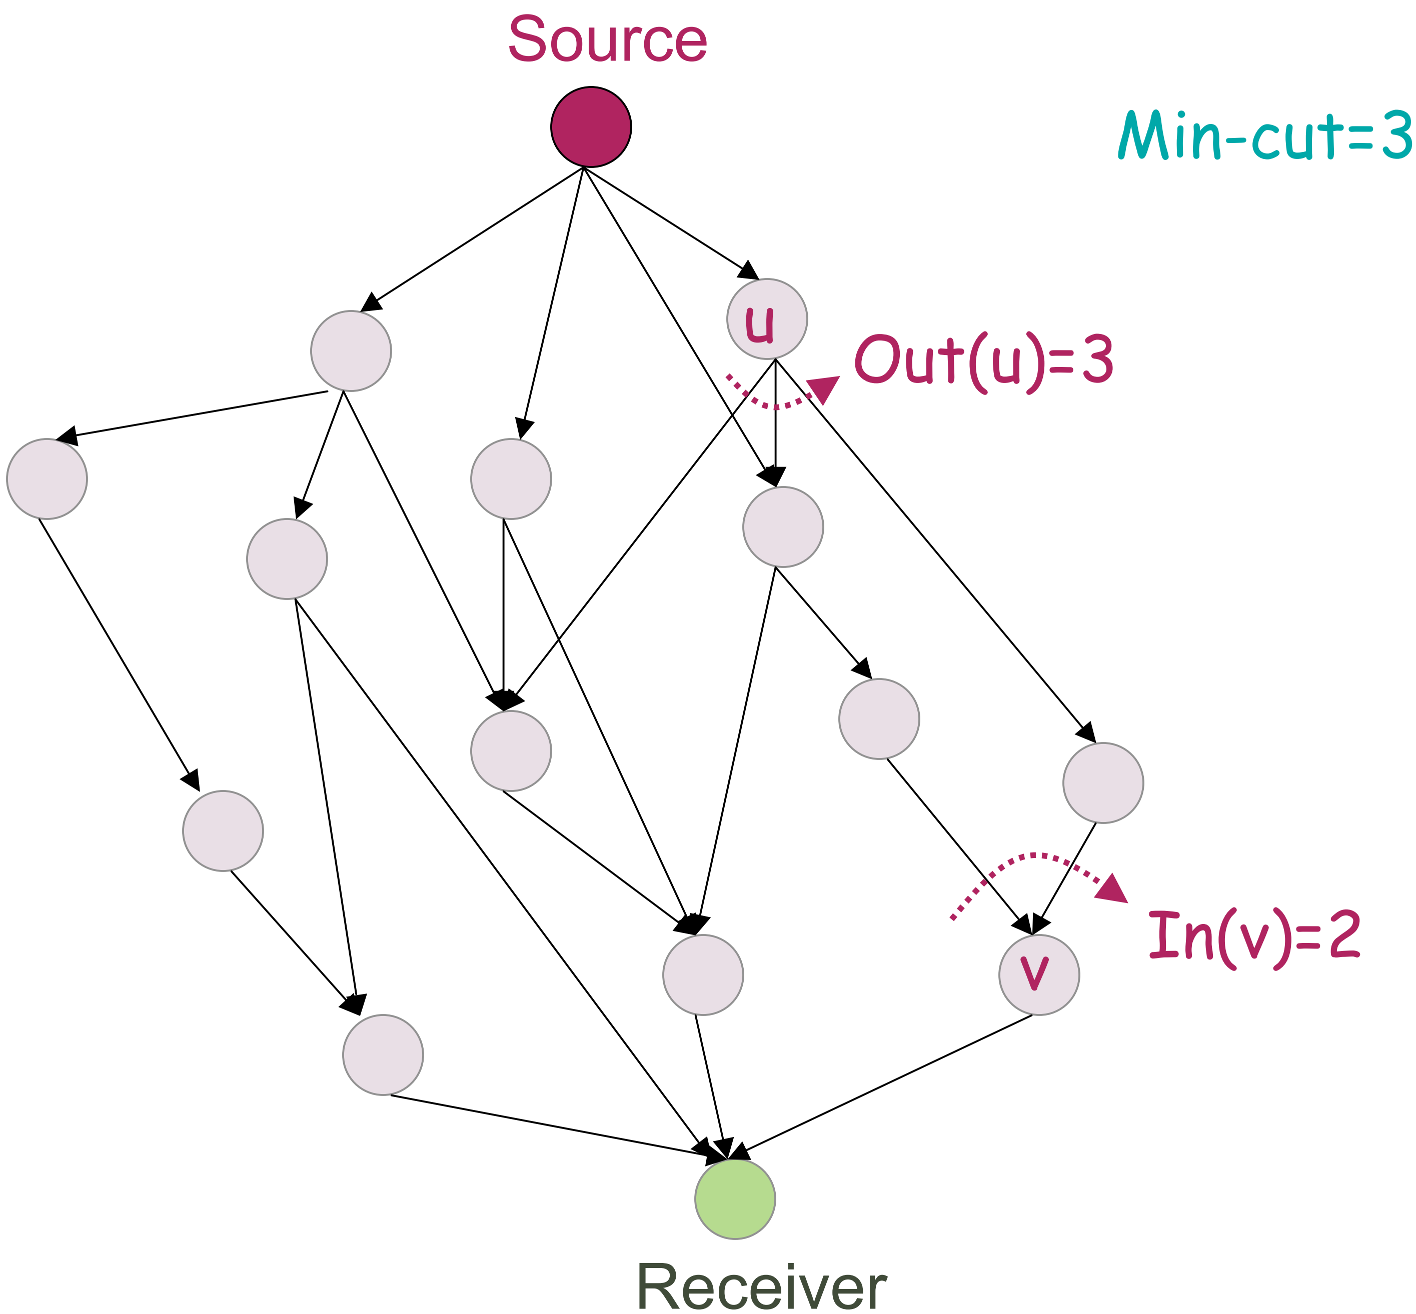
Source
Min-cut=3
u
Out(u)=3
v
In(v)=2
Receiver

Network Operation

The source produces B packets, $m_1, m_2, \ldots, m_B$

Each packet contains nR information bits.

Through every edge of the network we will send packets of length n bits.

Network Operation

The source selects uniformly at random Out(S) functions, one for each outgoing edge e:

$$f_e:\ 2^{nR} \longrightarrow 2^{n}$$

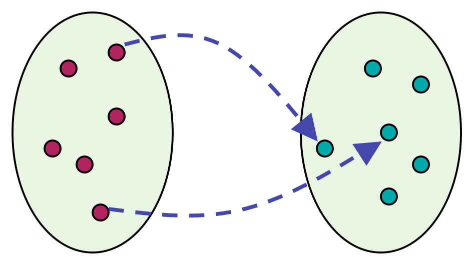

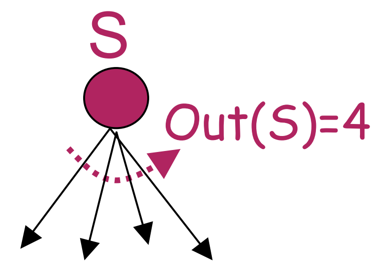

Each vertex v also selects Out(v) functions:

$$f_e:\ 2^{nIn(v)} \longrightarrow 2^{n}$$

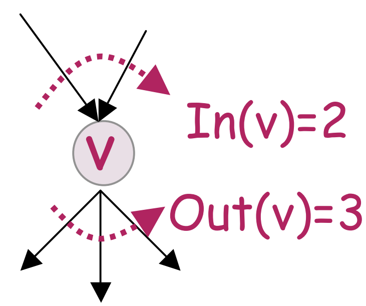

Timeslot 1

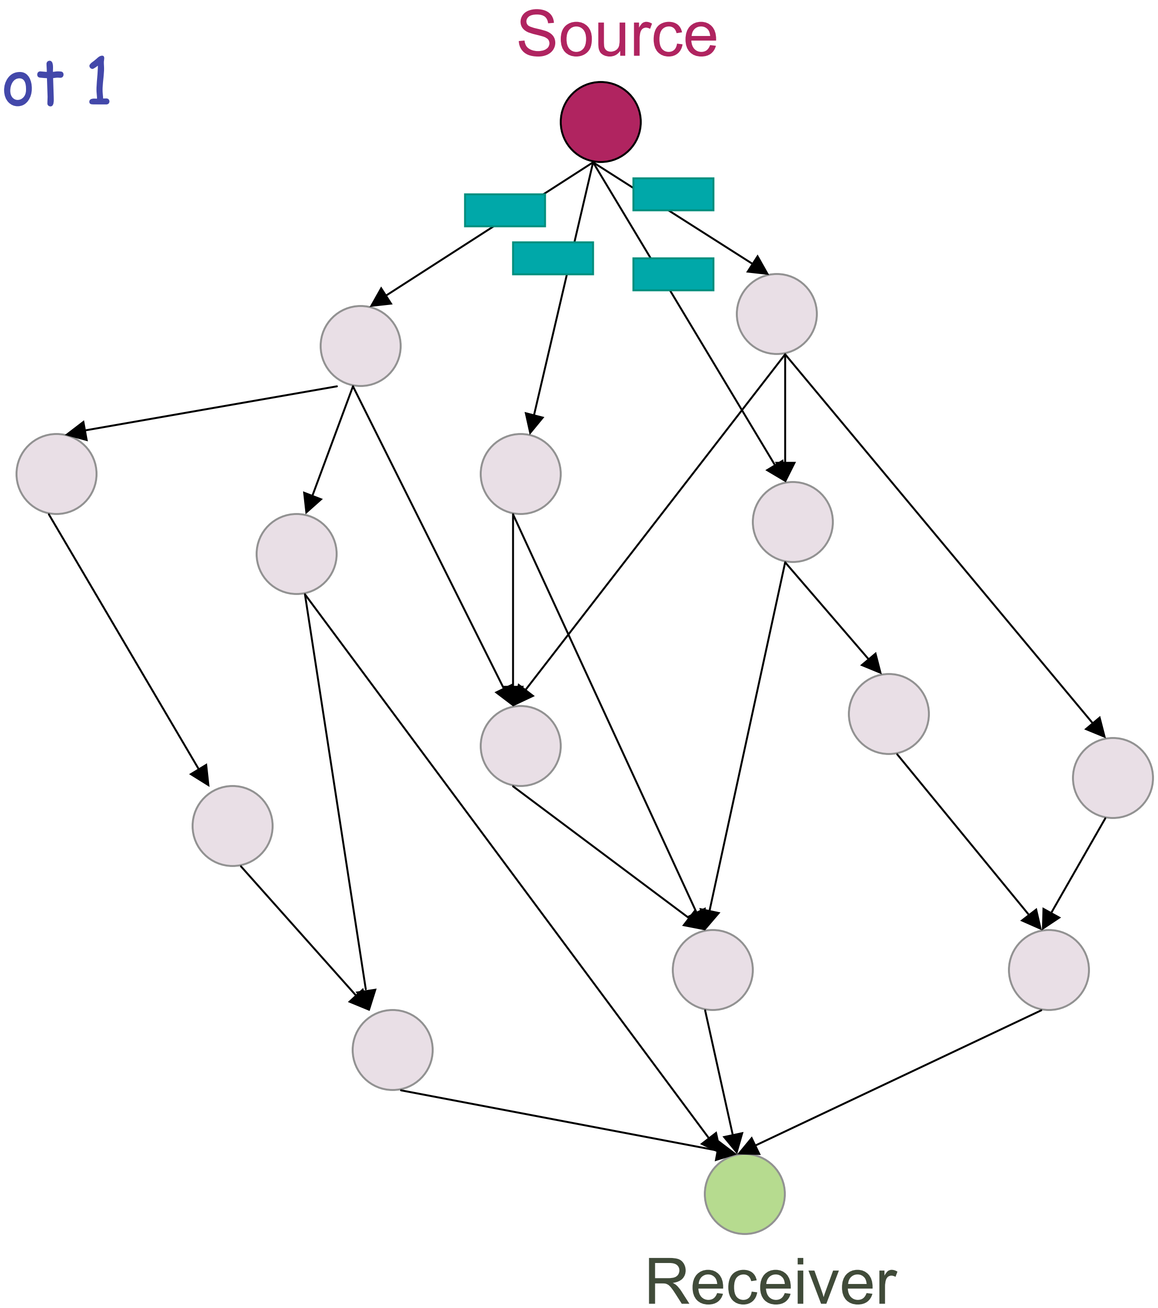

At the end of timeslot 1

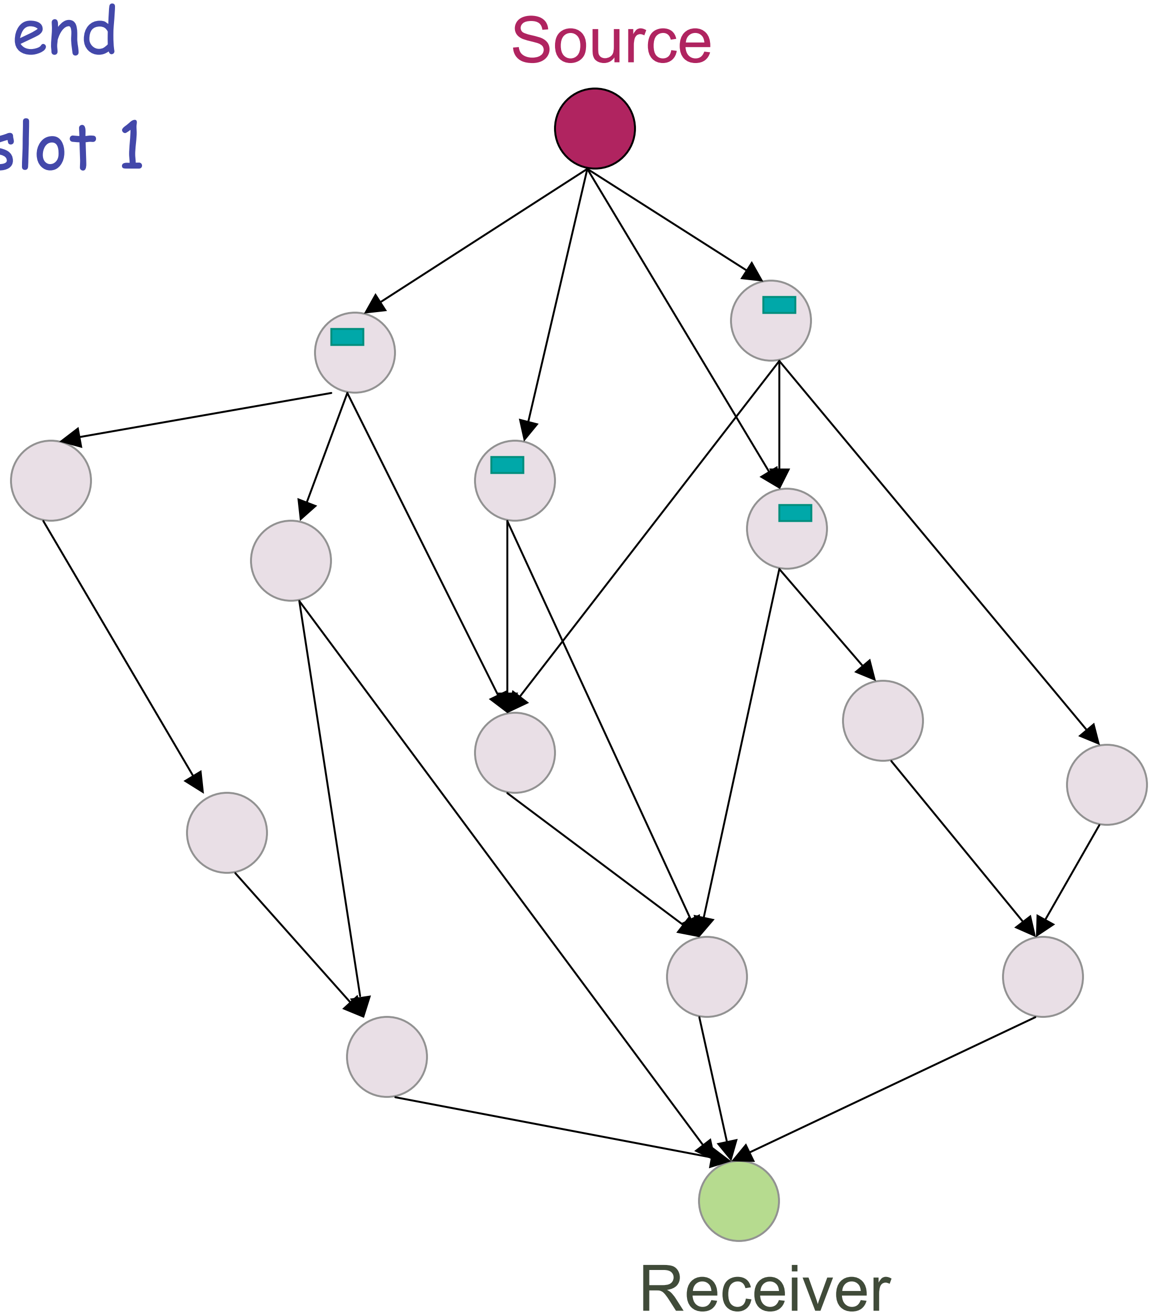

Timeslot 5

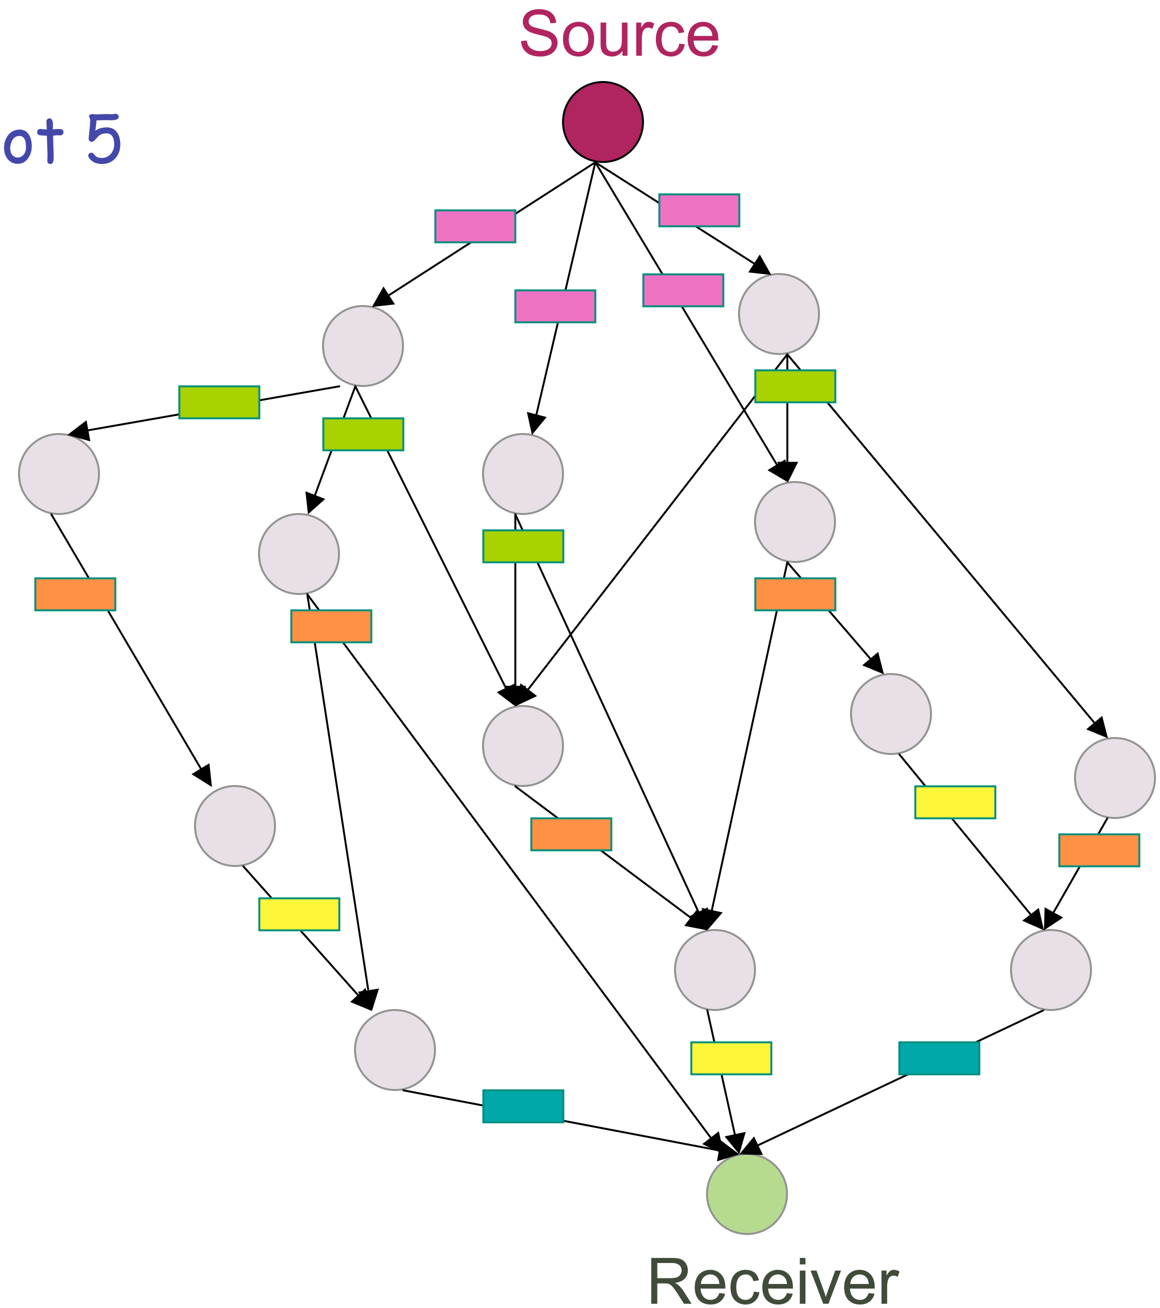

Network Operation

The network is clocked. At time-slot k, the source produces the packet m_k, and maps this packet to packets that it sends through its outgoing edges.

Each vertex v waits to collect In(v) packets that only depend on the source packet m_k and then maps these packets to packets it sends through its outgoing edges.

The receiver uses the packets it receives that depend on the source packet m_k and the knowledge of the network operation to decode packet m_k.

Why this network operation “works”

We will calculate the pairwise probability of error:

$P(m_1,m_2)$= probability that the receiver cannot distinguish between the messages m_1 and m_2 that the source sends

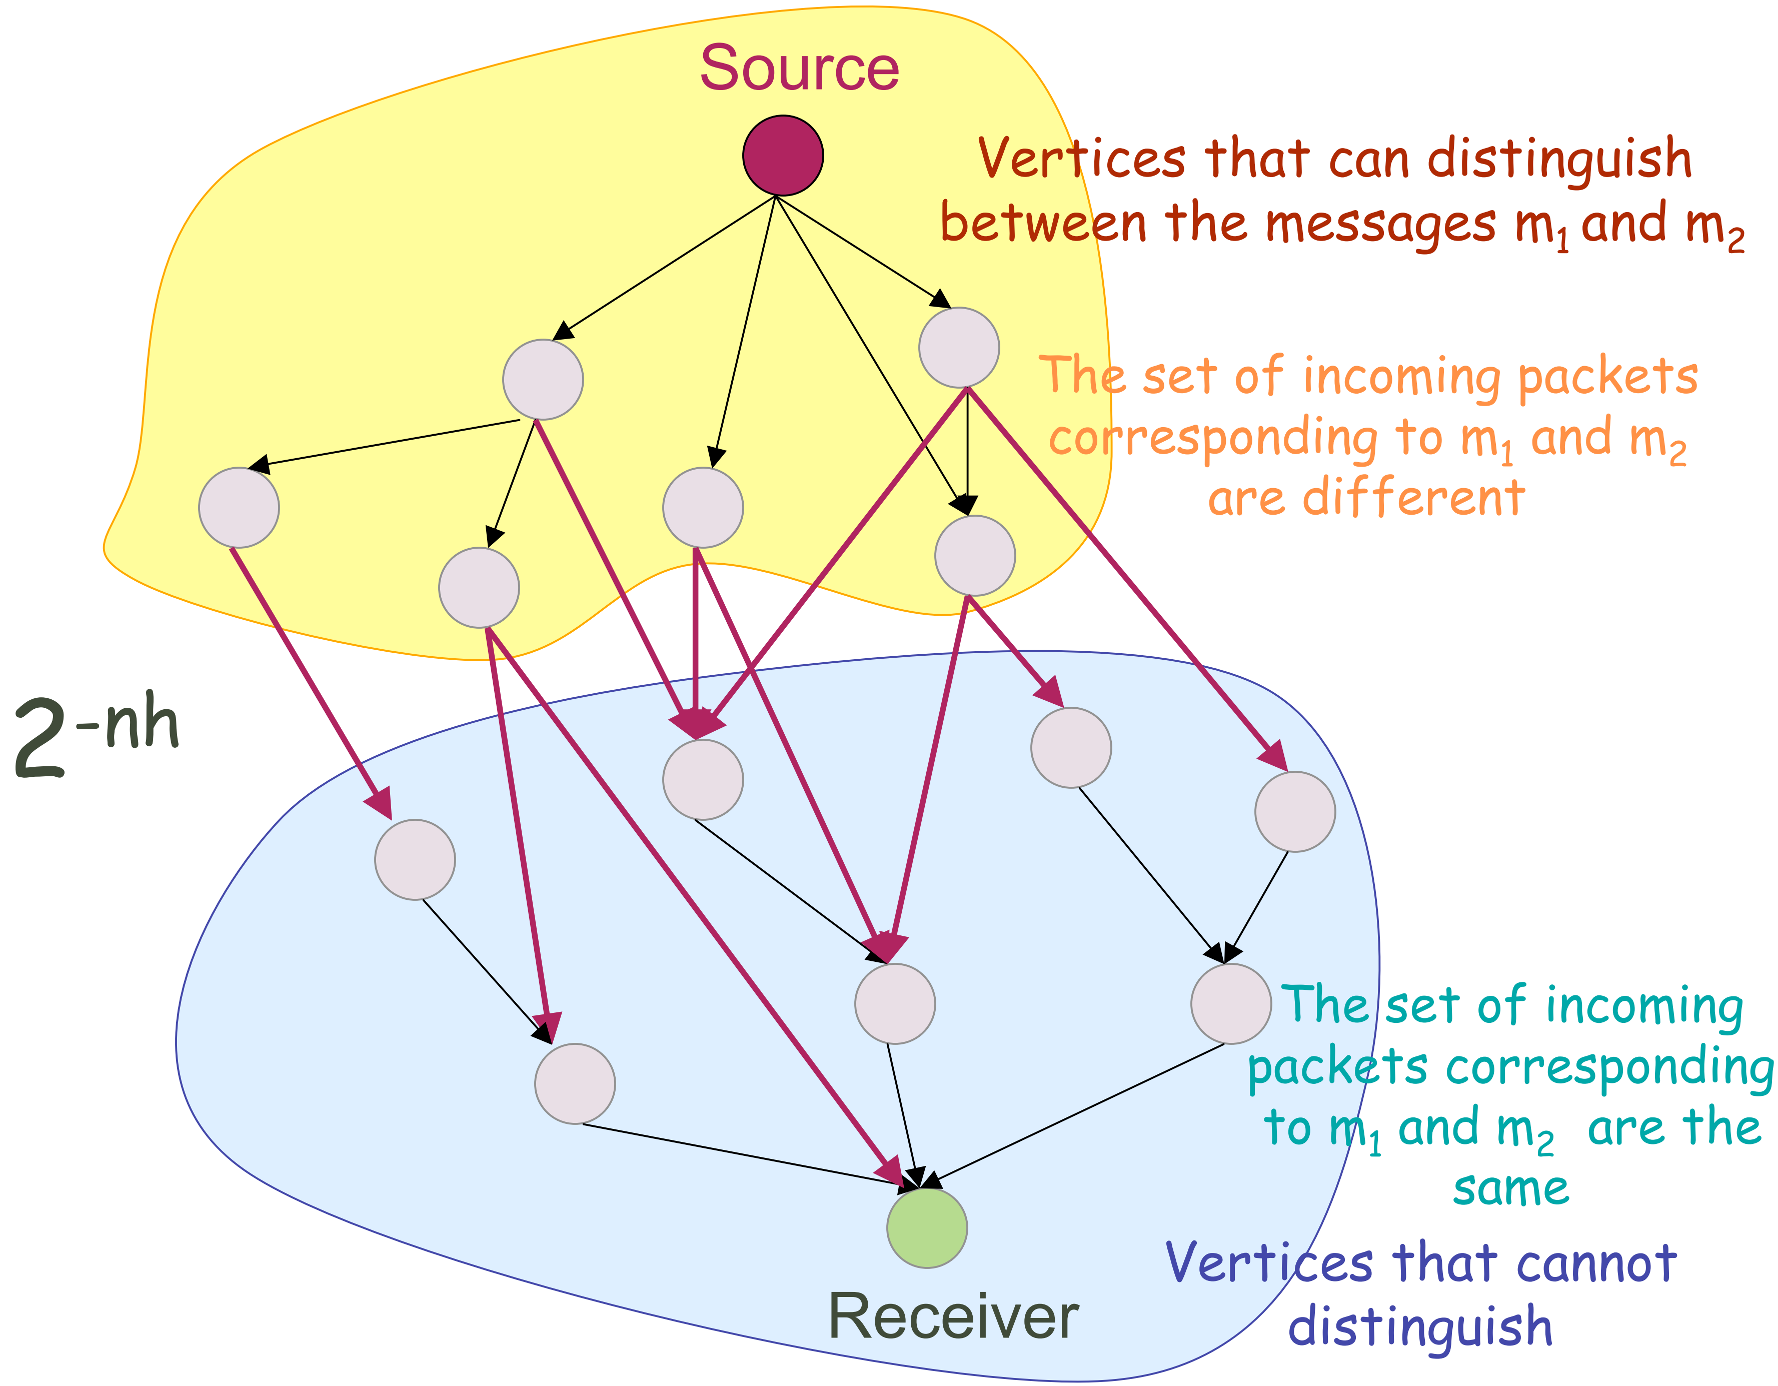
Source
Vertices that can distinguish between the messages m_1 and m_2
The set of incoming packets corresponding to m_1 and m_2 are different
2-nh
The set of incoming packets corresponding to m_1 and m_2 are the same
Vertices that cannot distinguish
Receiver

Why this network operation "works"

Provided 2^{nR-nh} goes to zero, i.e.,

$$R<h$$

we can transmit rate R from the source to the receiver.

“Interesting” components

Each vertex choses its operation independently of:

-where it is situated in the network,

-what choice of operation the remaining vertices select.

Very simple and decentralized routing protocols

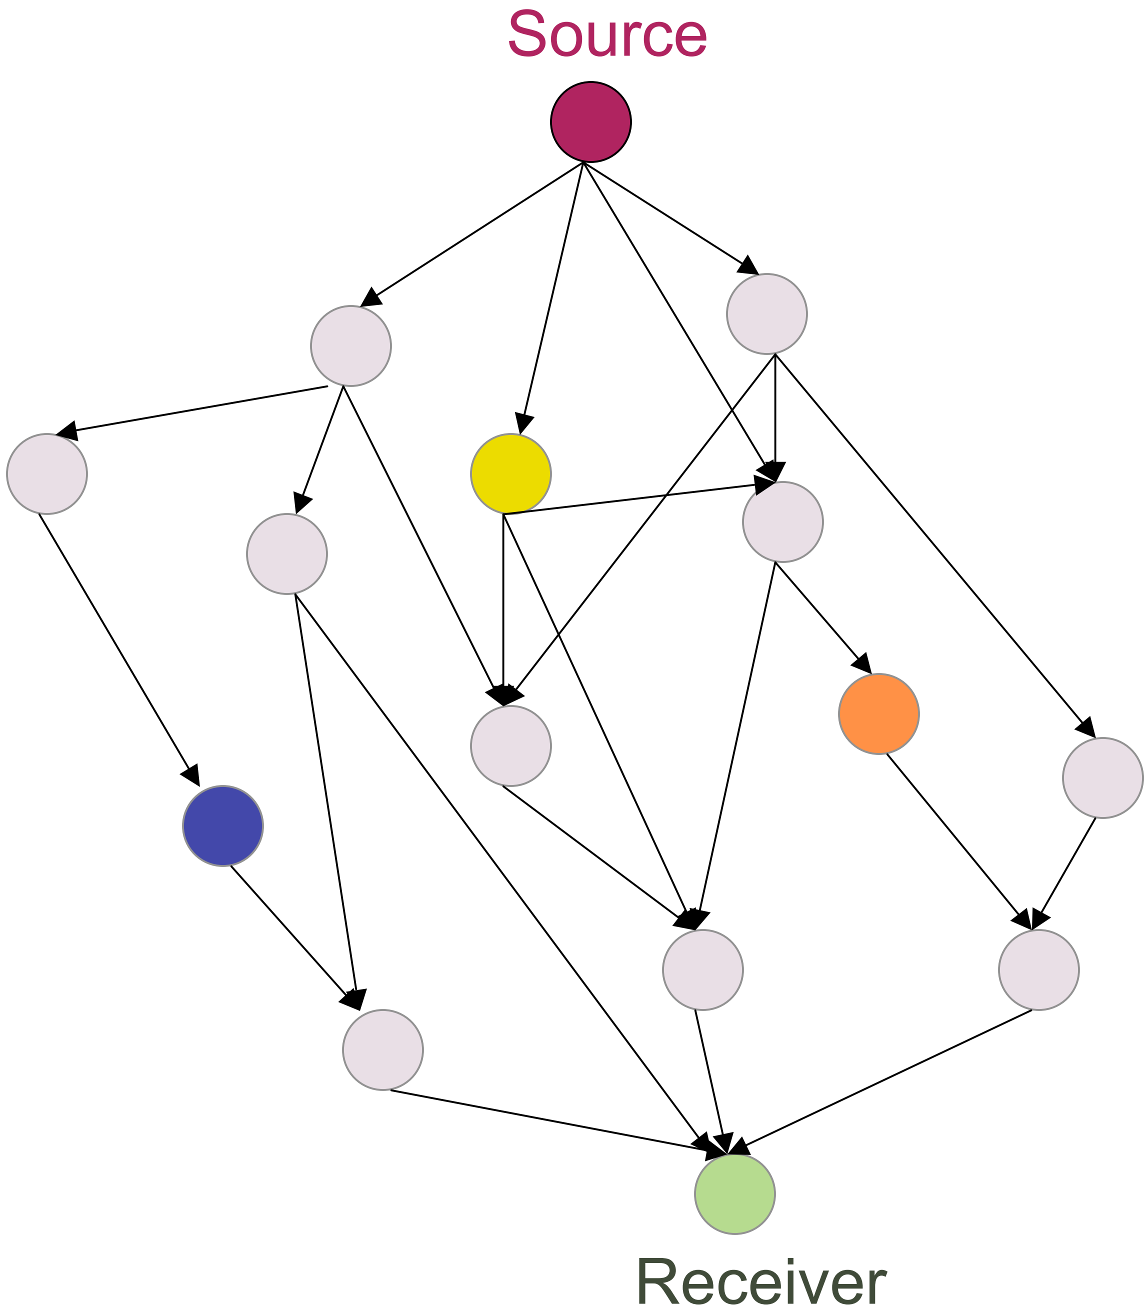
Source
Receiver

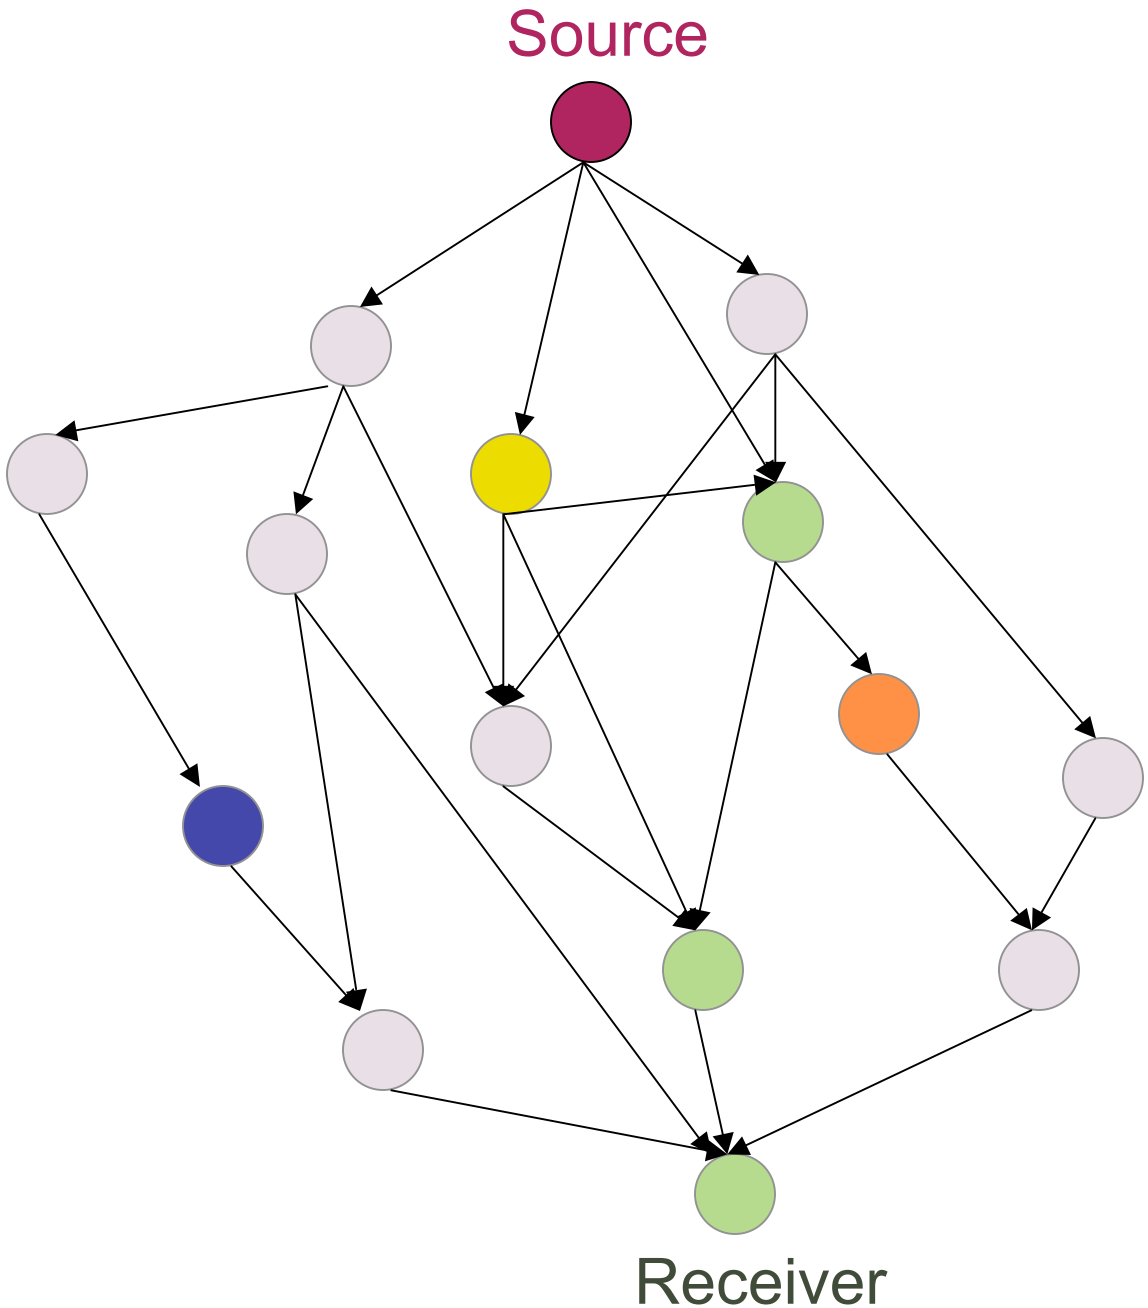
Source
Receiver

Outline

- Main Network Coding Theorem Proof
- Corollary: Coupons Collector Problem
- Application: Ad-hoc Wireless and Sensor Networks

Coupon Collector Problem

Traditional Approach

h coupons $x_1,x_2,x_3,\ldots,x_h$
are placed uniformly at random inside boxes

$\boxed{x_3}$ $\boxed{x_1}$ $\boxed{x_4}$ $\boxed{x_1}$ $\boxed{x_2}$ $\boxed{x_3}$ $\boxed{x_5}$ $\boxed{x_2}$

How many boxes do we need to buy on the average in order to collect all coupons?

$$\boxed{O(h \log h)}$$

Coupon Collector Problem

Traditional Approach

h coupons $x_1, x_2, x_3, \ldots, x_h$
are placed uniformly at random inside boxes

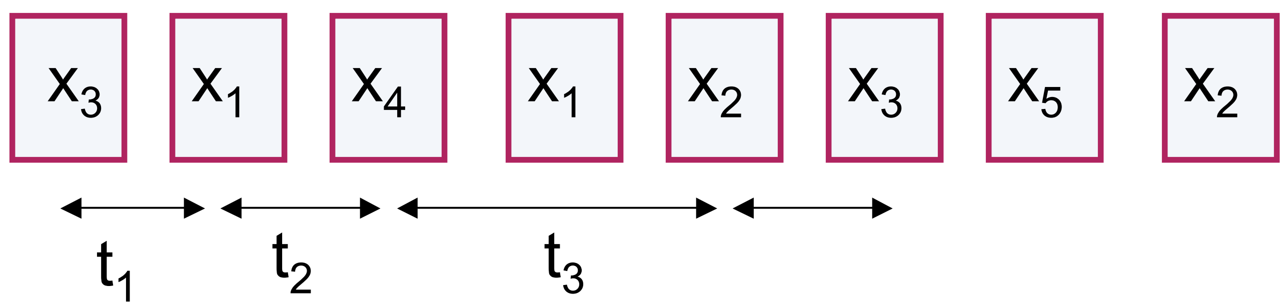

t_i=time to collect the i+1 coupon from the time we have collected i

Probability of success: $p_i=1-i/h$, thus $E(t_i)=1/p_i$

$p_h=1/h$

Coupon Collector Problem

using Network Coding

h coupons $x_1, x_2, x_3, \ldots, x_h$

Each box has a linear combination of the coupons

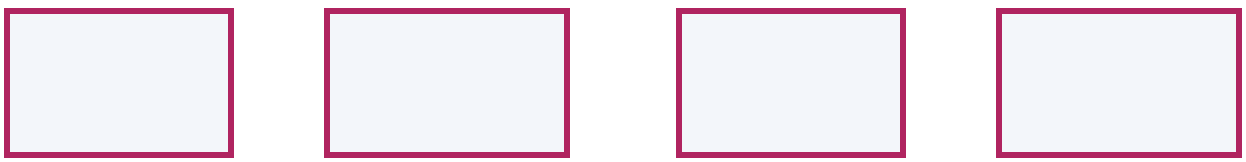

How many boxes do we need to buy on the average in order to collect all coupons?

Coupon Collector Problem

using Network Coding

h coupons $x_1, x_2, x_3, \ldots, x_h$

Each box has a linear combination of the coupons

x_1+x_2	x_3+x_4	x_1+x_5	x_3+x_2

How many boxes do we need to buy on the average in order to collect all coupons?

Coupon Collector Problem

using Network Coding

h coupons $x_1, x_2, x_3, \ldots, x_h$

Each box has a linear combination of the coupons

$\boxed{x_1+x_2}$ $\boxed{x_3+x_4}$ $\boxed{x_1+x_5}$ $\boxed{x_3+x_2}$

How many boxes do we need to buy on the average in order to collect all coupons?

(Deb and Medard 04) $\boxed{O(h)}$

Gain a factor of logh

Coupons Collector as a Network Problem

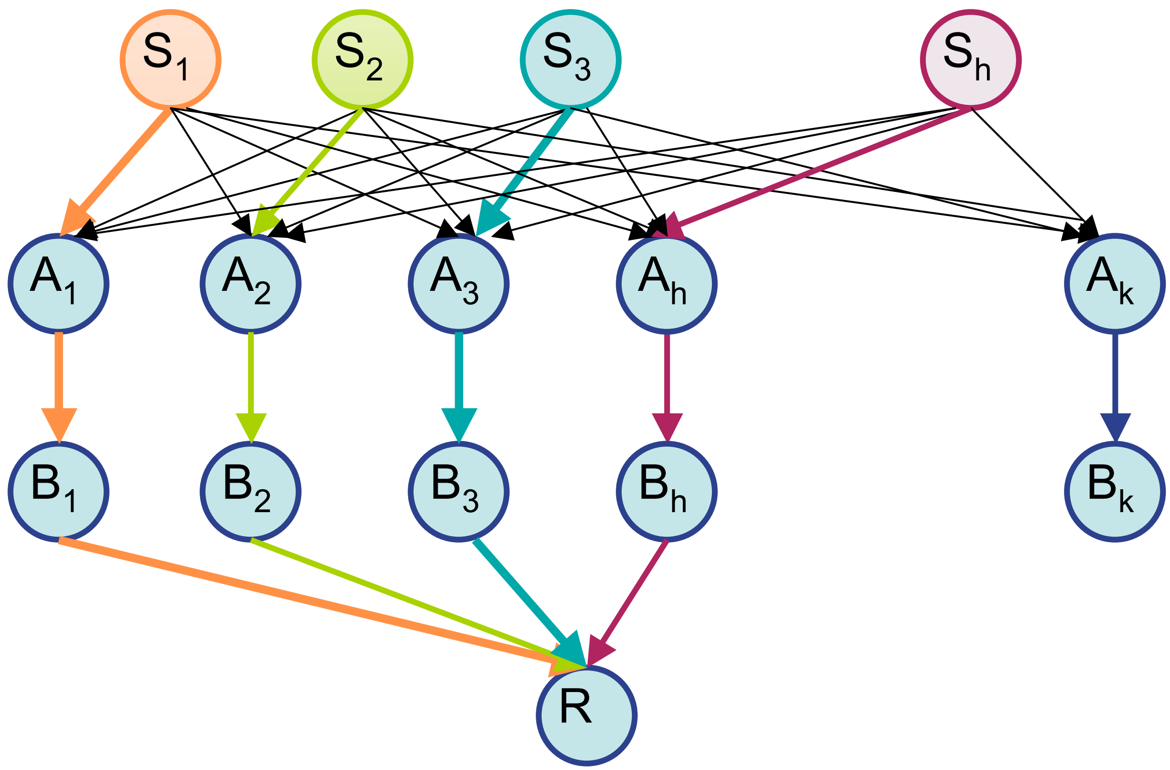

Coupons Collector as a Network Problem

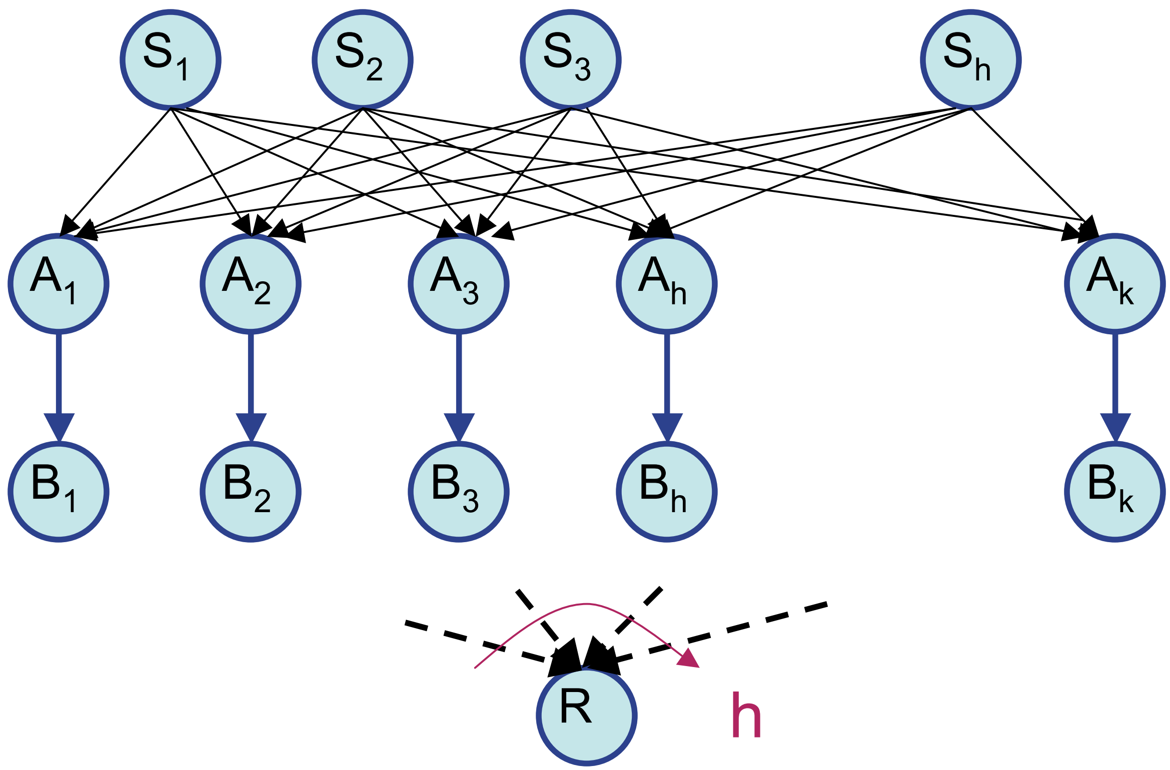

Broadcasting over a Square Grid

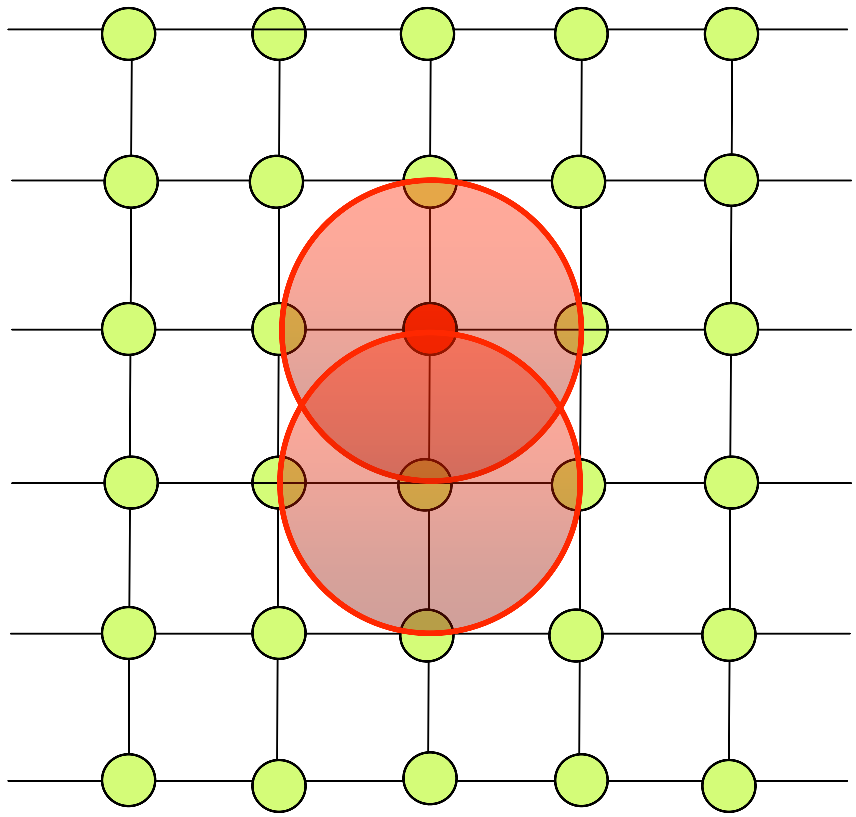

Static network:
We gain a factor
of 3/4

Forwarding: every transmission reaches three new neighbors

Network Coding: every transmission reaches four new neighbors

Forwarding routing protocol

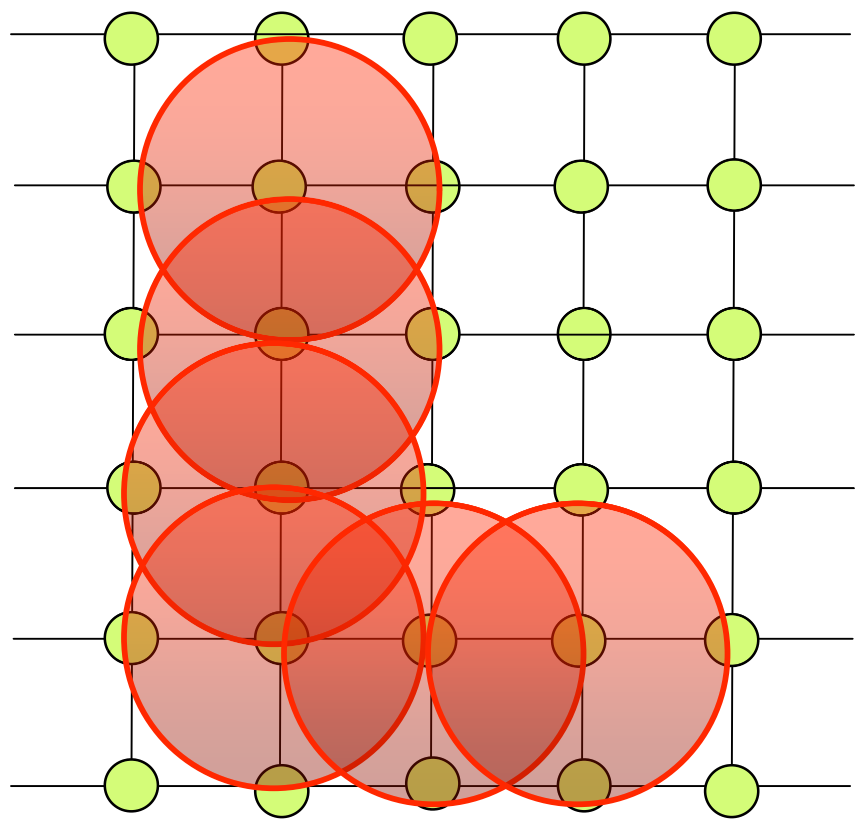

Network Coding Protocol

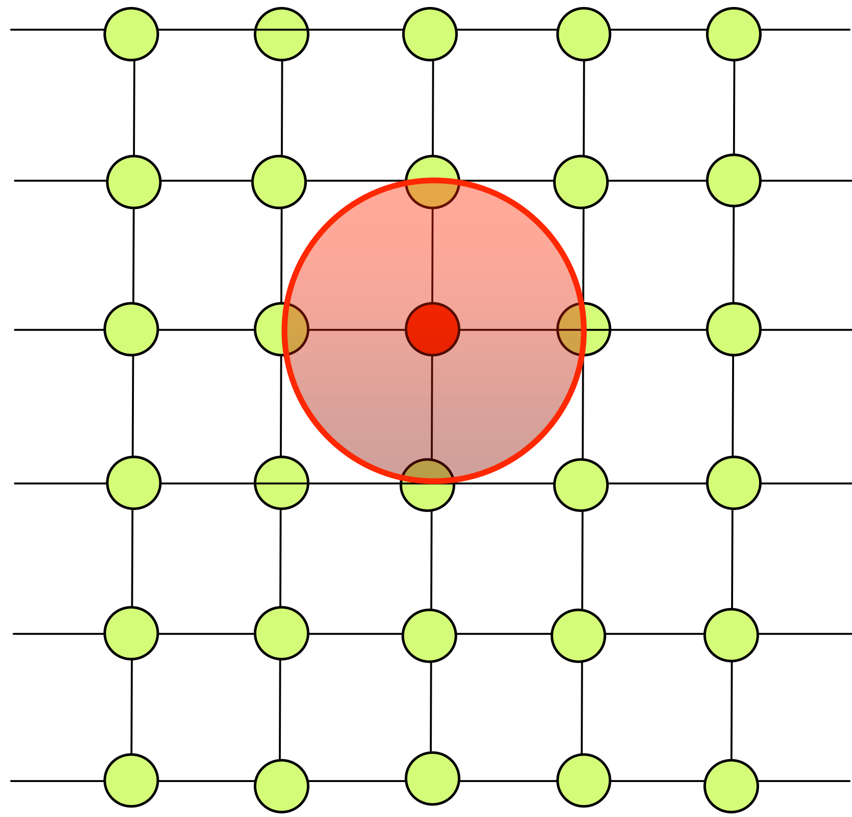

Each node
sends a random linear
combination of its
previously received
symbols

Network Coding Protocol

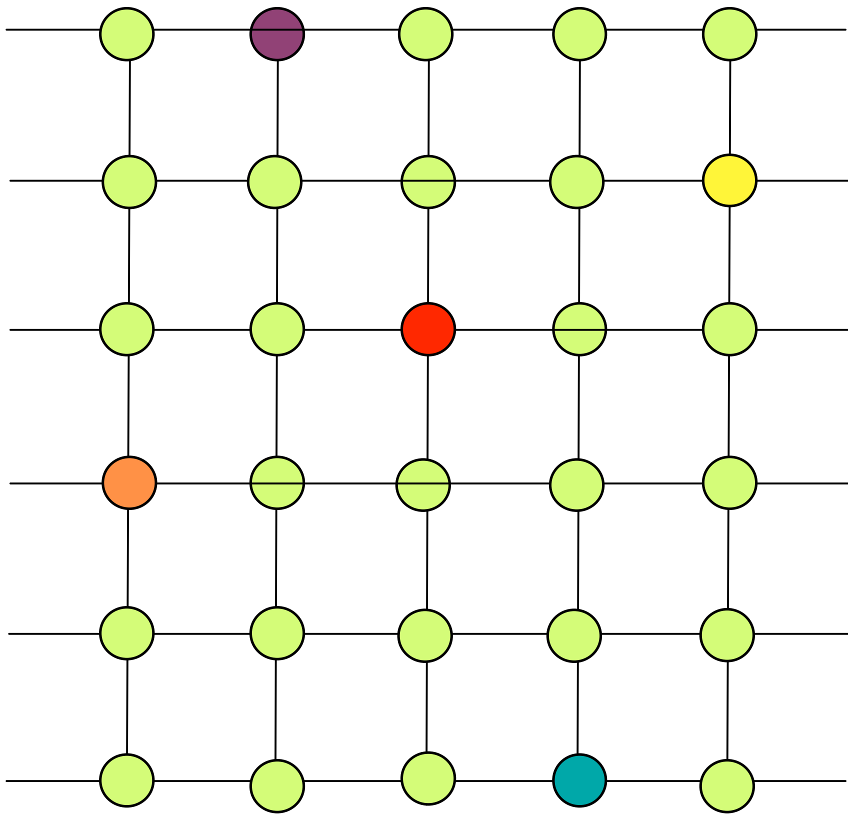

Each node
sends a random linear
combination of its
previously received
symbols

Theoretical Results

n nodes randomly placed on unit disk area.
Each node has transmission radius $r = \Theta(\frac{1}{\sqrt{n}})$.
Discrete time

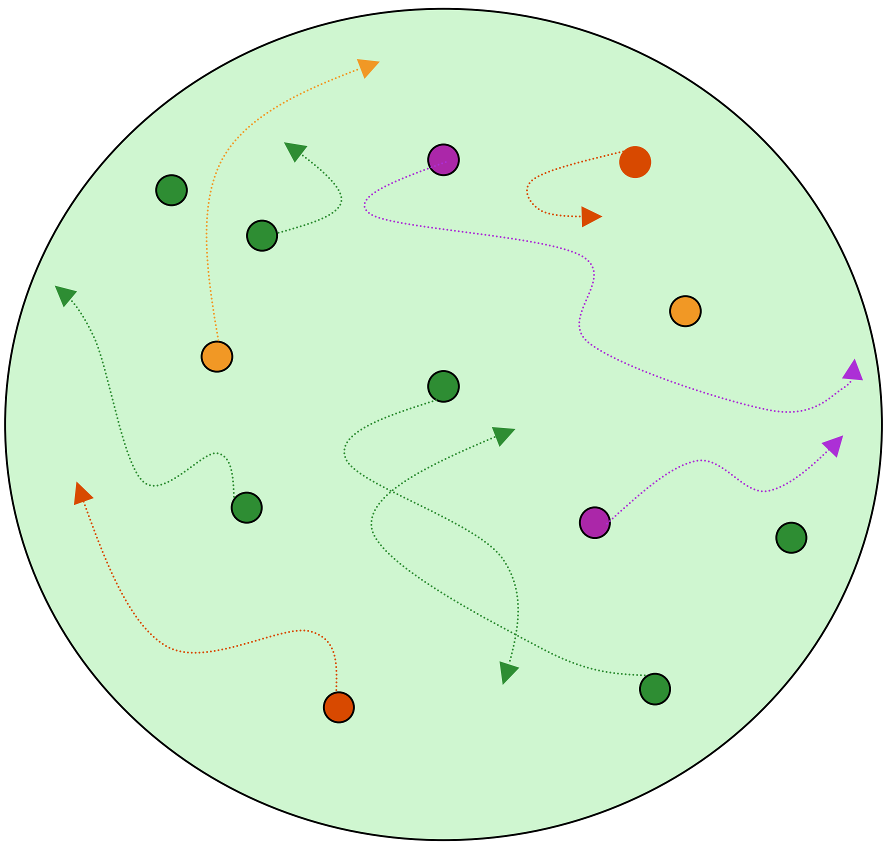

Nodes move uniformly at random on the unit area disk:
at each time slot, a node has on the average a constant number of neighbors chosen uniformly at random.

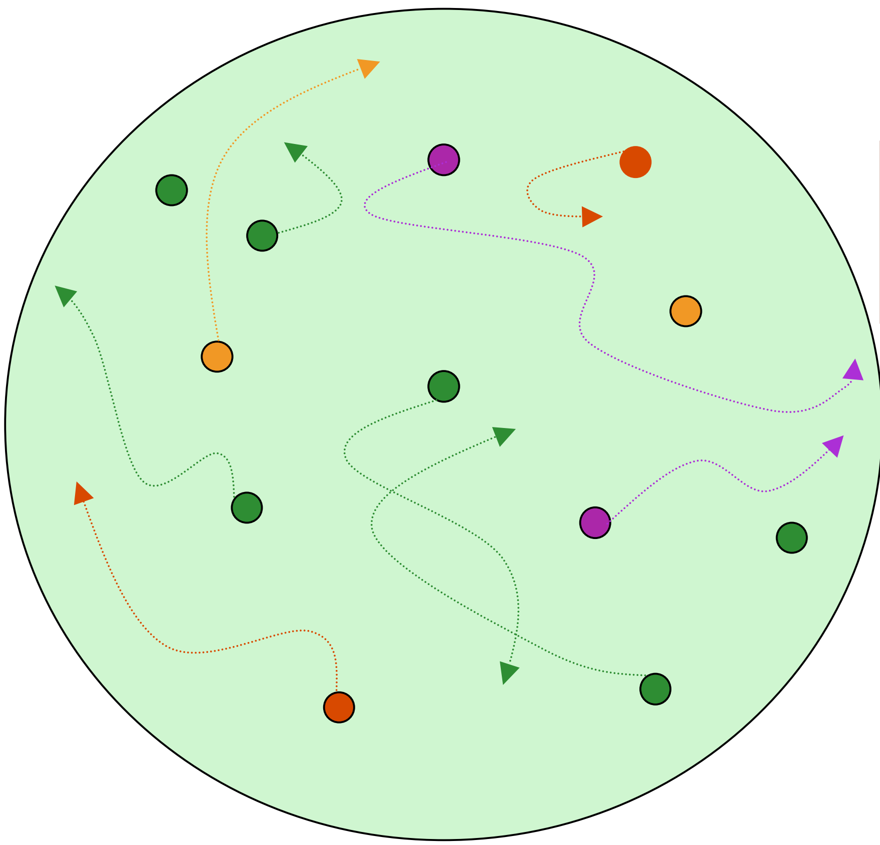

Transmission strategy

Forwarding: each node broadcasts its own symbol.

Network coding: each node transmits a random combination of the symbols it has received.

Network coding offers logn

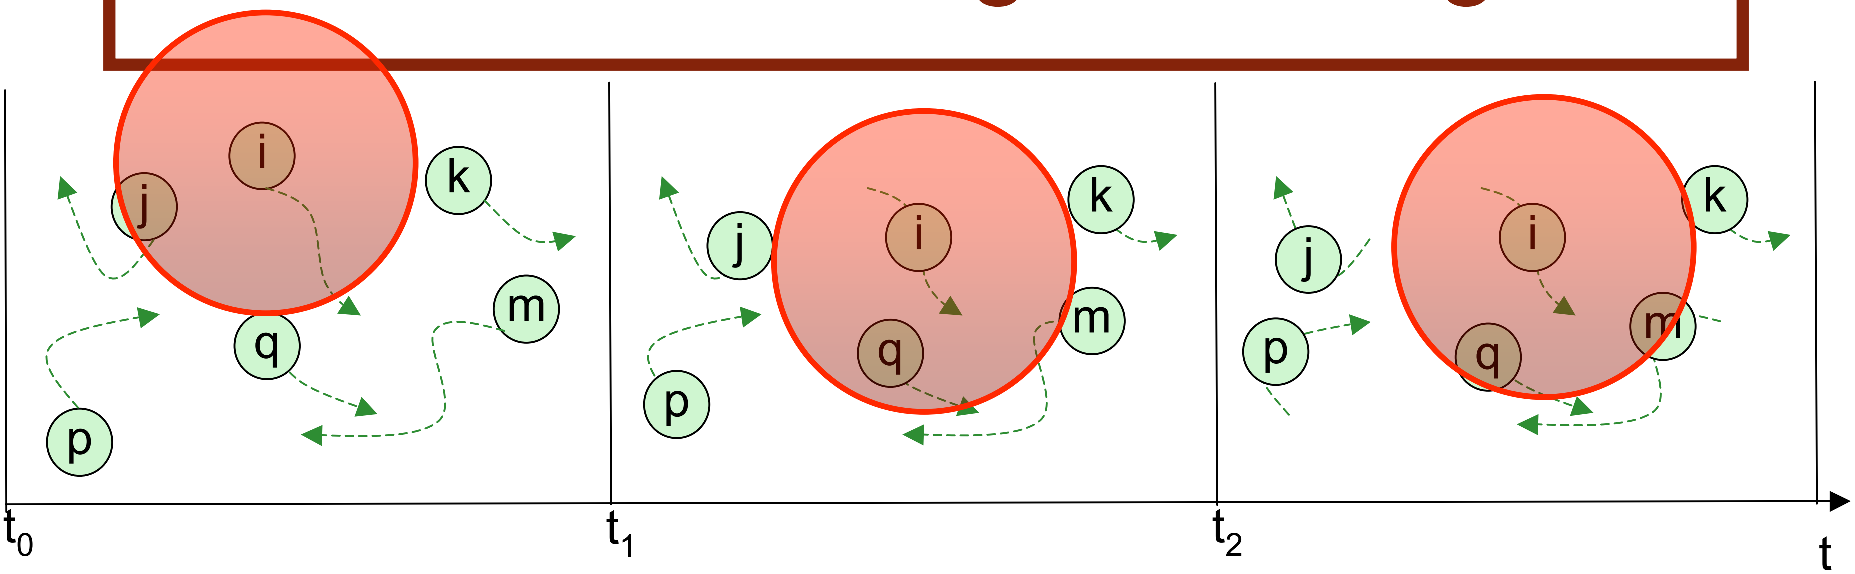

Reduce to a variation of the coupons collector problem:

Each node at each timeslot receives the information from a constant (on the average) number of neighbors.

Forwarding:

$\Theta(n \log n)$ timeslots

Network coding:

$\Theta(n)$ timeslots

Simulation Results

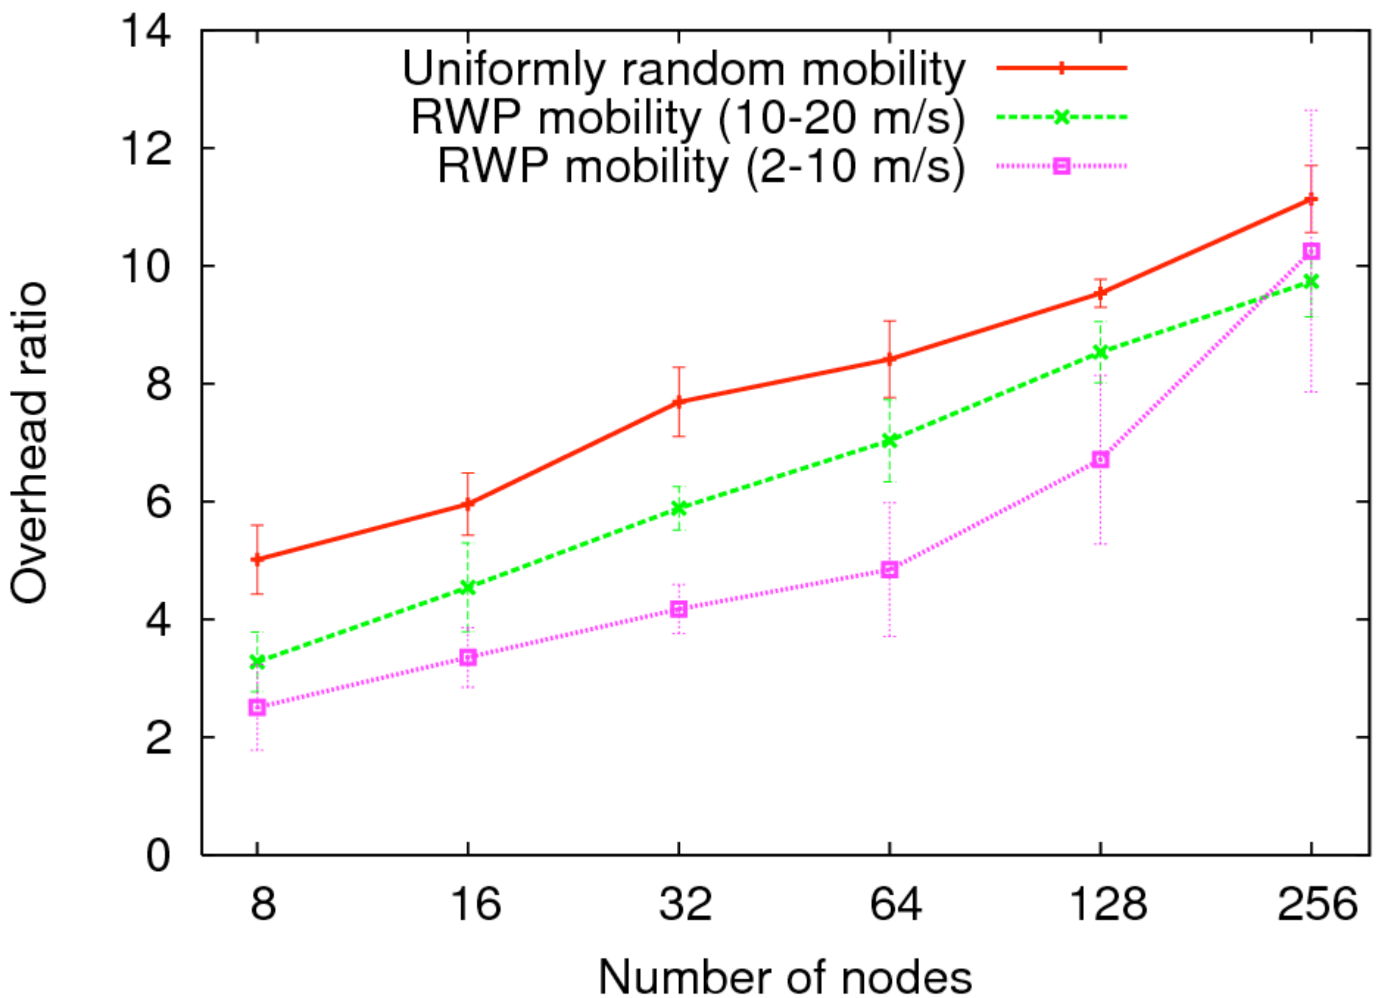

Other applications of the coupons collector problem

Vehicular Networks

Vehicles communicate with each other and with roadside infrastructure to increase safety and optimize traffic.

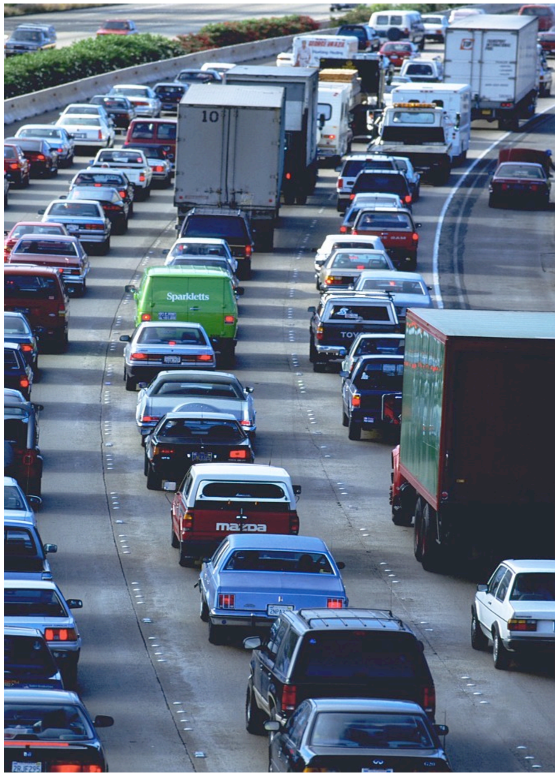

Goal: Distribute updates.

To get n messages, each car will have to receive:

- Without network coding: n logn broadcast transmissions
- With network coding: n broadcast transmissions

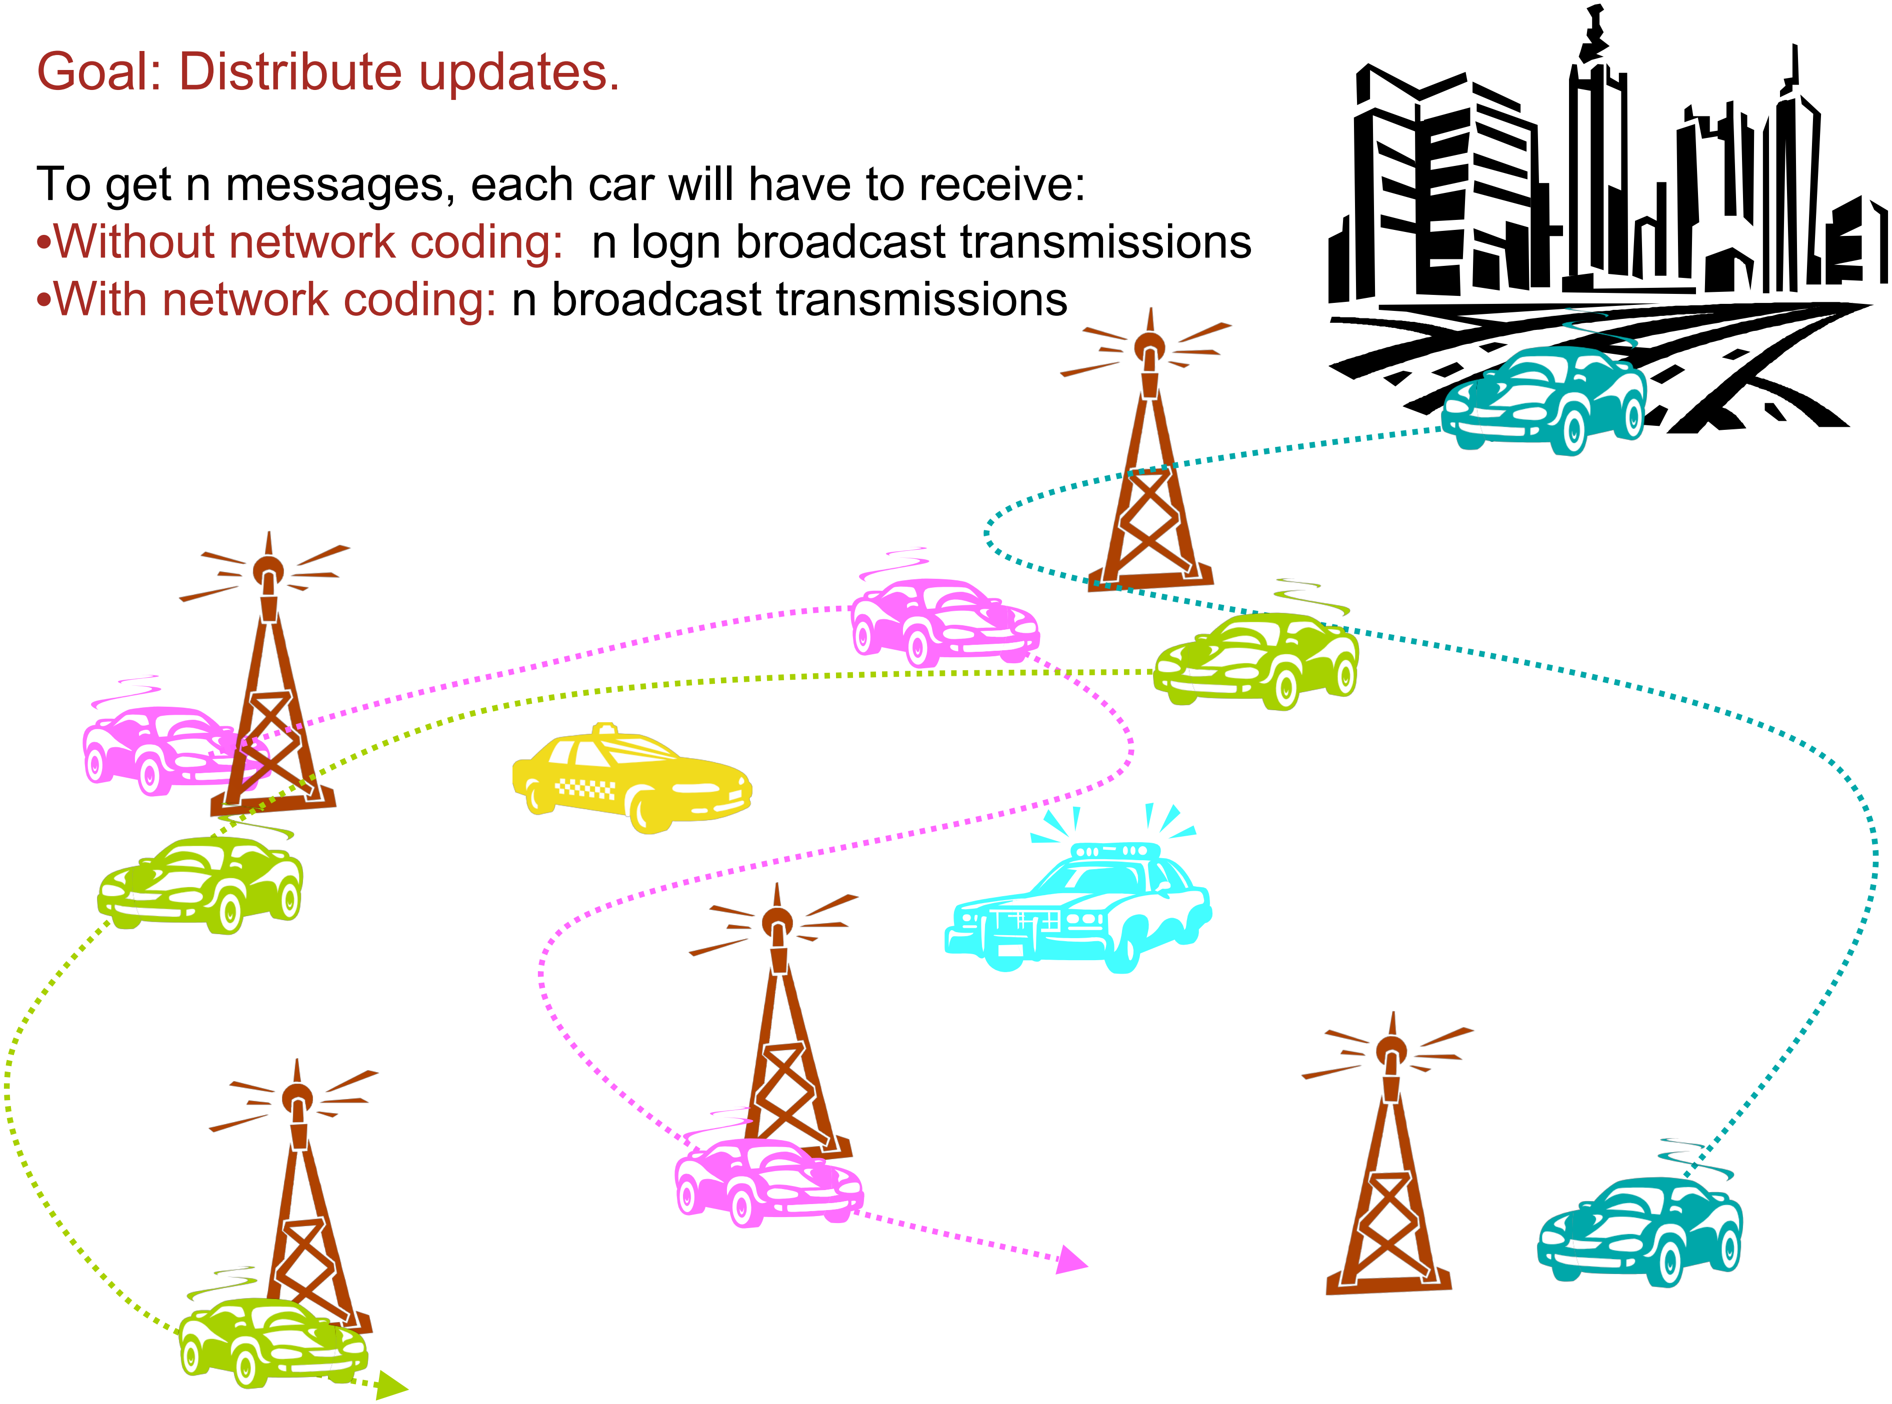

Applications of Network Coding

- Ad-hoc wireless networks
- Content distribution in P2P networks
- Network tomography
- Sensor networks
- Security
- Chip design
- .....

Content Distribution

Distribute content to millions of users, such as

- Software updates
- Music
- Films
- ….

Traditional Approach

Contect collected in servers, clients connect to servers to download the information.

Problems:

Not-scalable, expensive and slow (servers can crash)

P2P networks:

Capacity and computational power of the network increases with the number of users.

- File divided into n packets
- Peers collect and forward packets from and to their neighbors

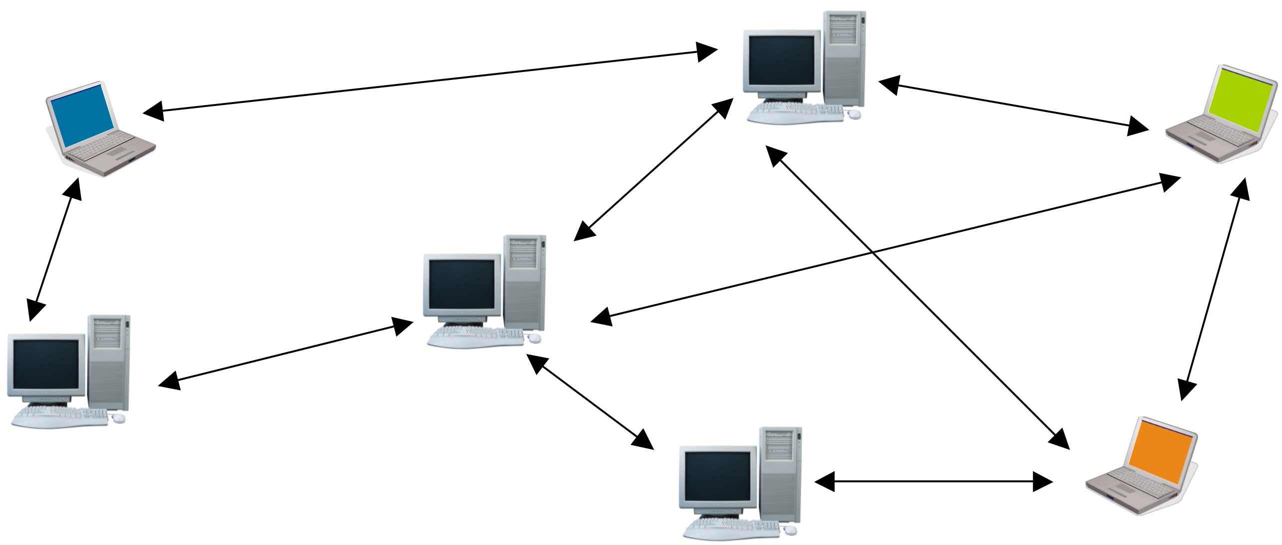

P2P networks:

Capacity and computational power of the network increases with the number of users.

Challenge:

how to optimally and securely route packets

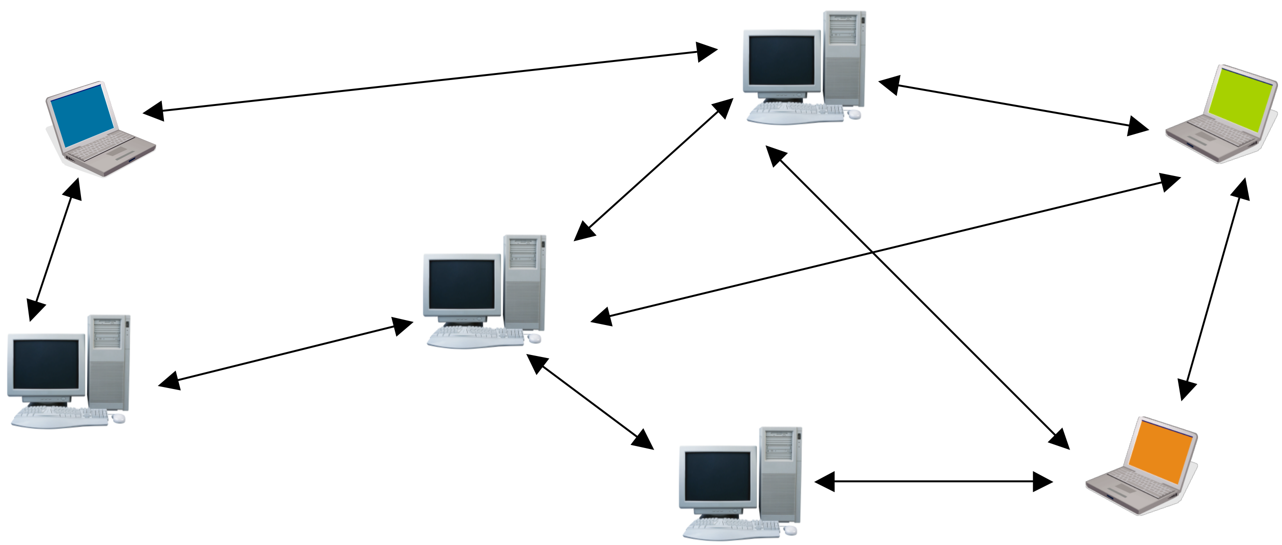

Challenge:

how to optimally route packets

- Same packets may be send several times over bottleneck links
- Some packets become rare (users leave)
- New users arriving slow down old users
- Tit-for-tat incentive mechanisms slow down new users

Avalanche

Microsoft Research Cambridge

P. Rodriguez, C. Gkantsidis

Peers exchange random linear combinations of their received data.

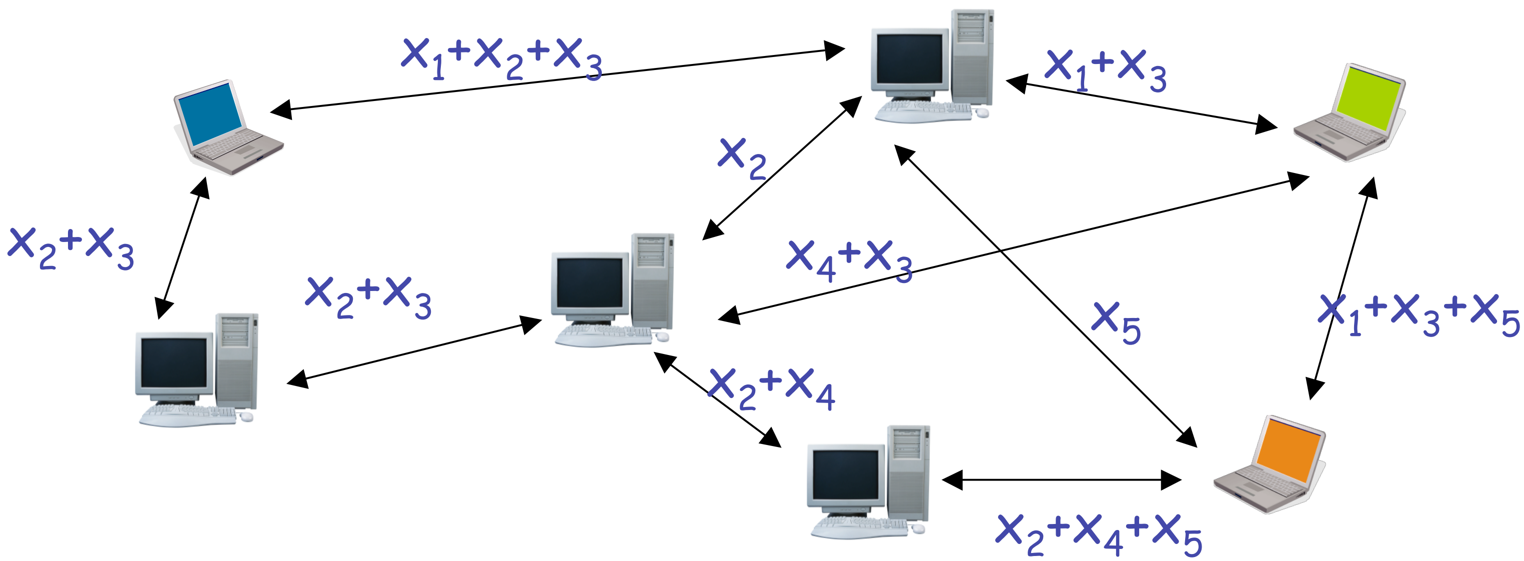

Avalanche Robustness

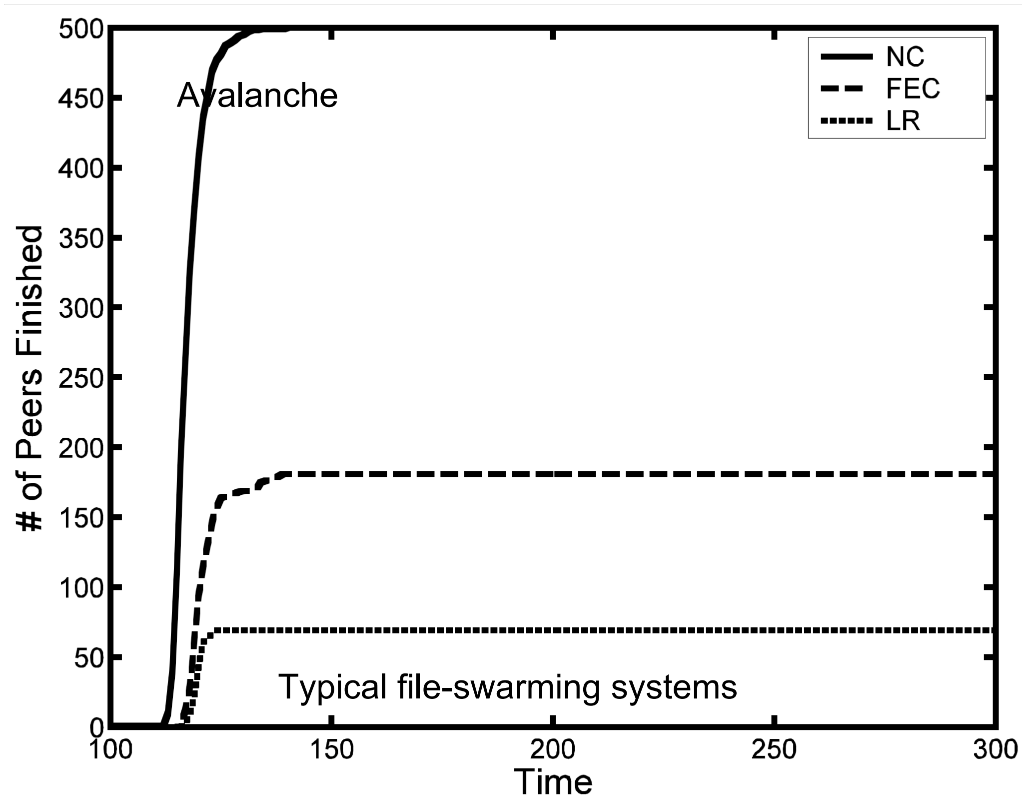

If source suddenly goes down (after serving the full file one), all Avalanche users are able to complete the download. Only 10% of users using typical file-swarming techniques are able to complete.

Plot provided courtecy of
P. Rodriguez and C. Gkantsidis

Microsoft® Research Cambridge

A new challenge:
Byzantine Attacks

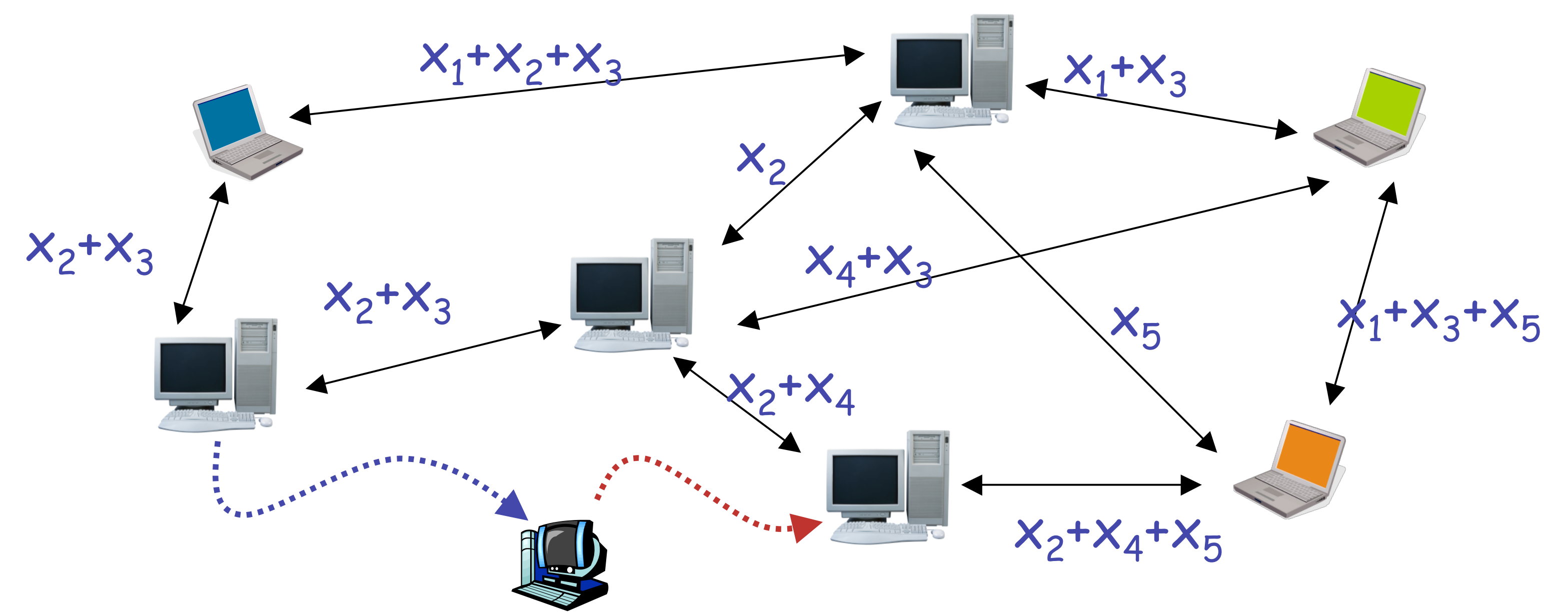

Byzantine Attacks

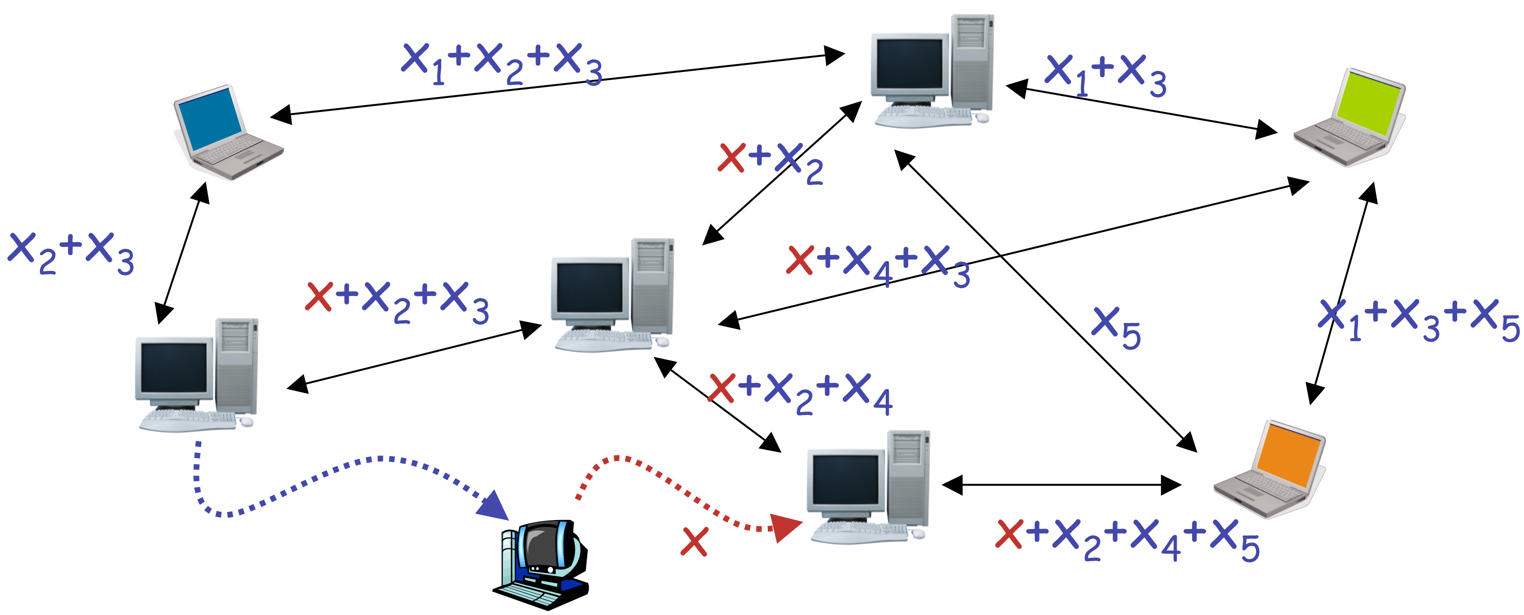

Applications of Network Coding

- Ad-hoc wireless networks
- Content delivery in P2P networks
- Network tomography
- Sensor networks
- Security
- Chip design
- …..

Network Tomography

Goal: Measure the Internet path characteristics such as loss and delay through active probing, to improve the robustness and reliability of the network

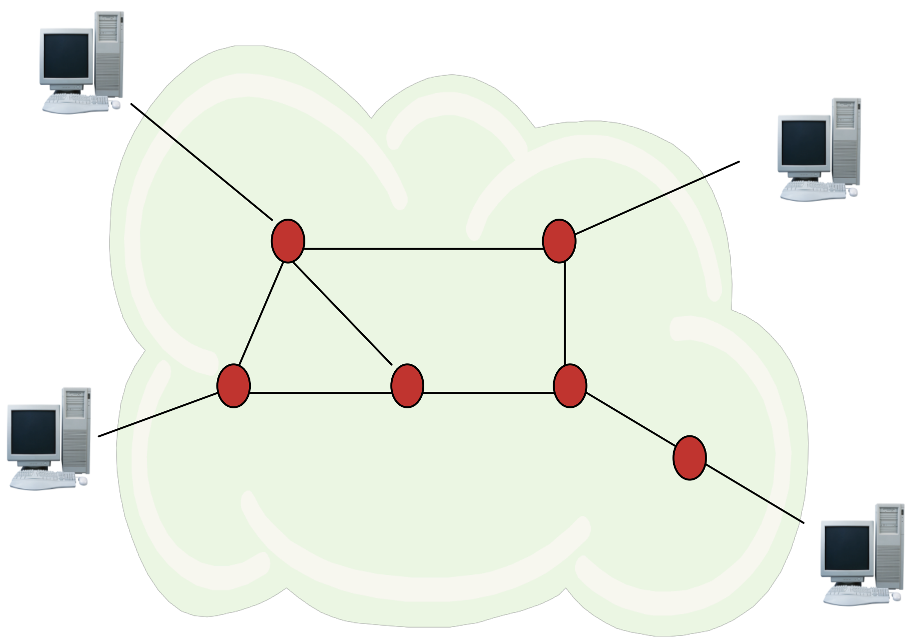

Approach:

Combine probe packets

Network coding benefits:

1) Bandwidth
2) Complexity
3) Identifiability

(Allerton 2005, IZS 2006)

Loss Inference w. Network Coding
Basic Example

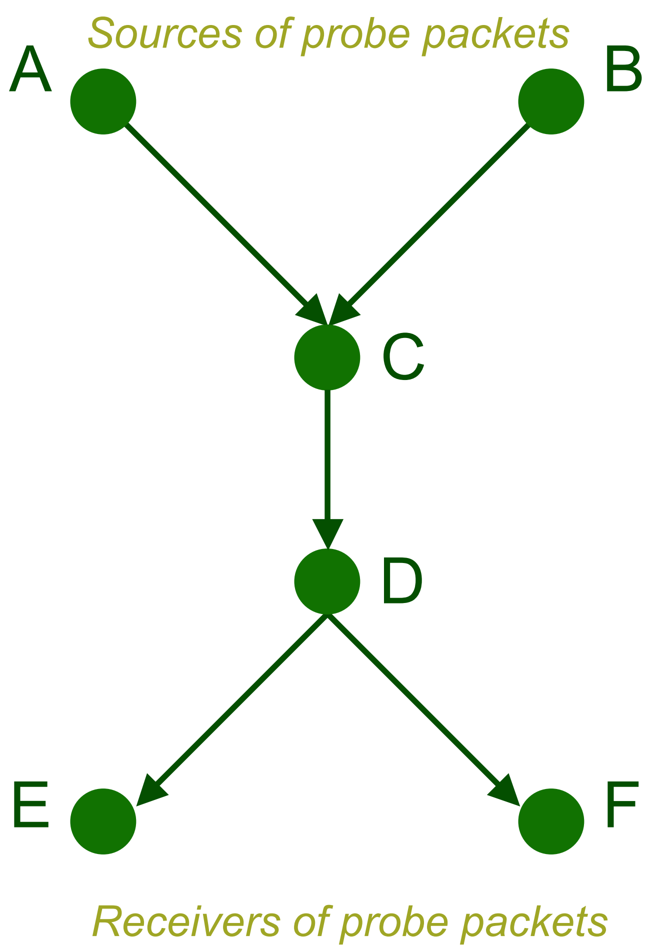

- We want to infer the link loss rates a_k on all links $k \in \{AB, AC, CD, DE, DF\}$
- using end-to-end probes from {A,B}, to {E,F}

Traditional Approach:
covering the graph with trees

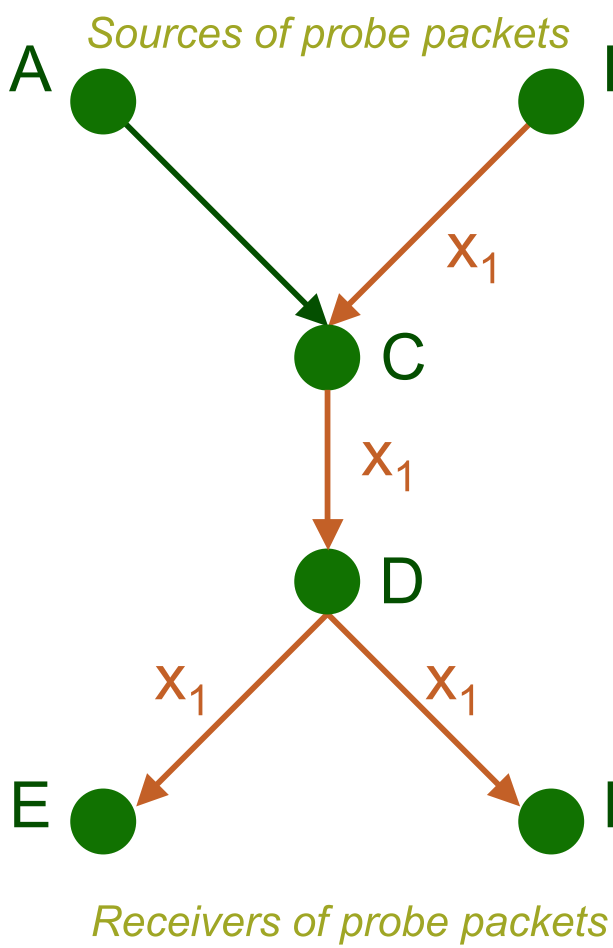

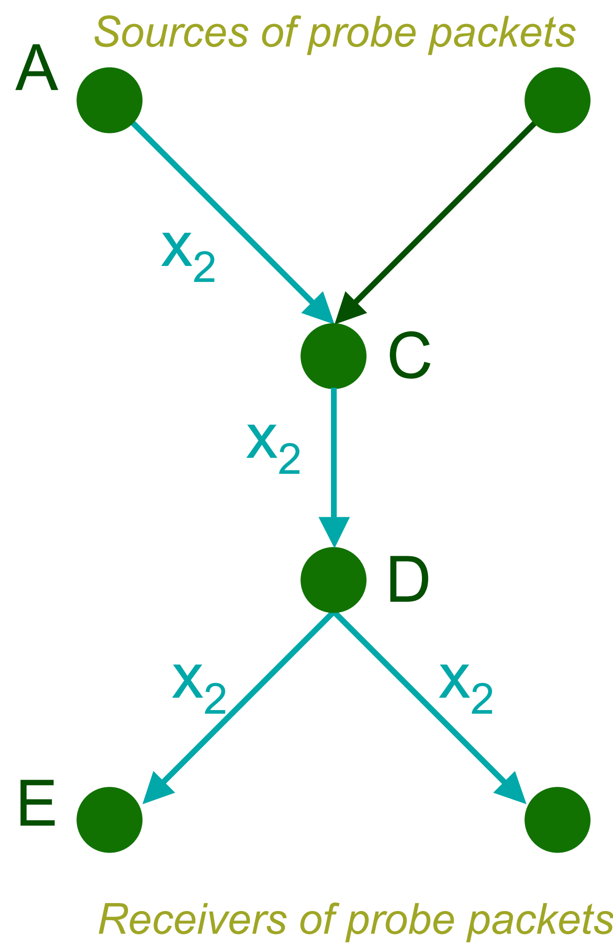

Drawbacks

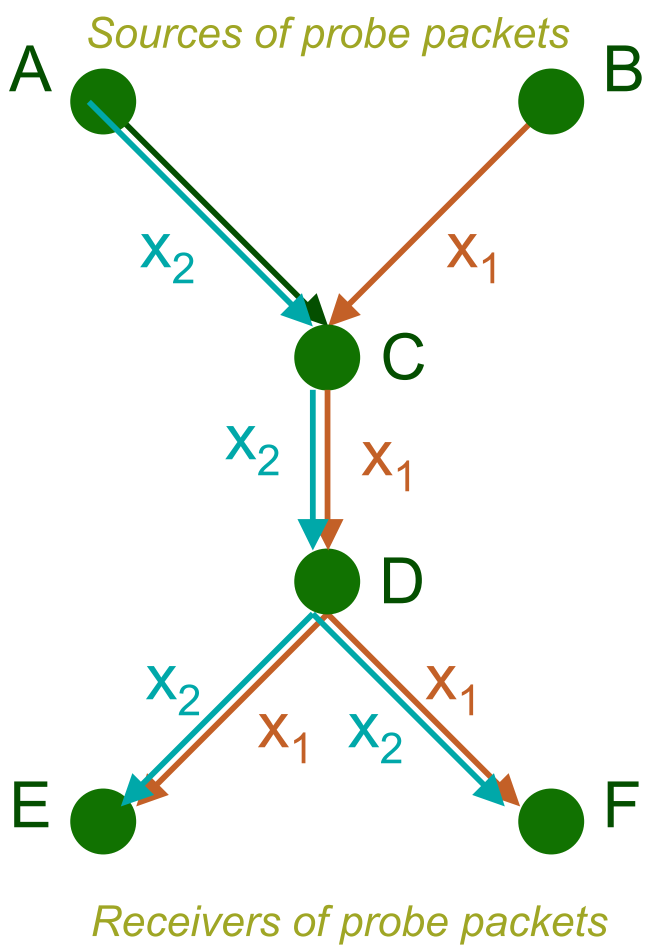

1. We cannot infer the loss rate for edge CD
2. Paths overlap from C and downstream
3. Minimum cost covering with multicast trees is NP-hard
4. Combining observations from 2 trees leads to suboptimal estimation

Network coding approach

[C.Fragouli, A. Markopoulou Allerton 05]

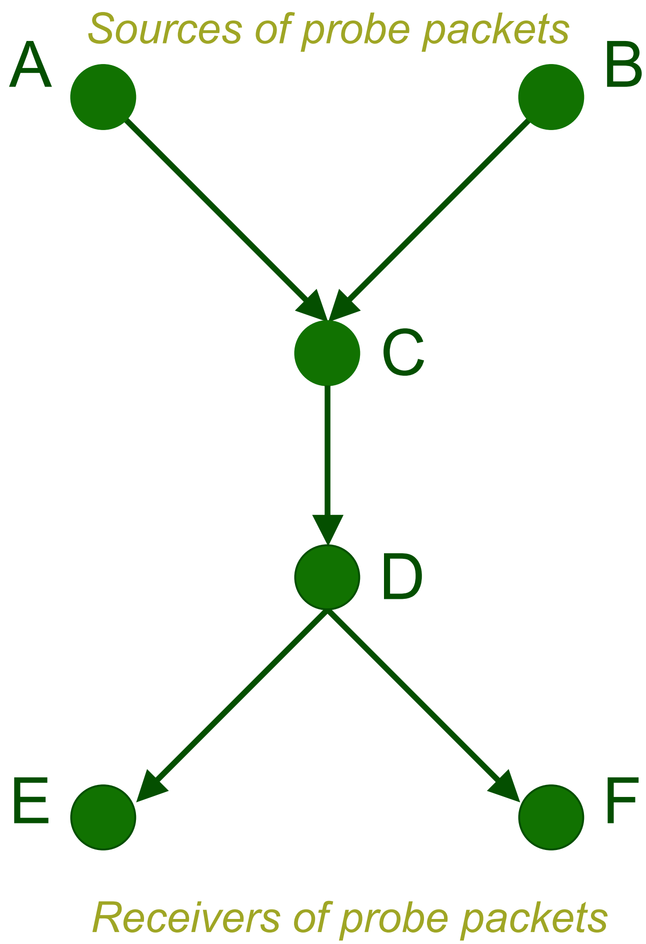

Intermediate node (C):

Within a time window

- if received 2 incoming packets,
  - XOR them and forward
- if received 1 incoming packet
  - just forward

Network Coding Approach

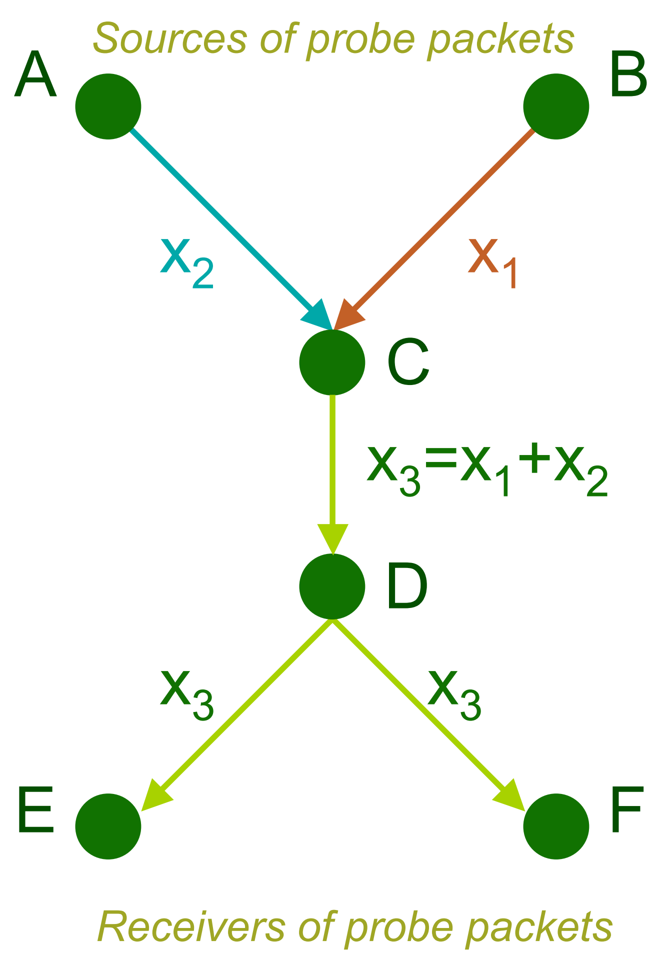

Example:

Nodes A and B send packets

x_1= [1 0], x_2 = [0 1]

Multiple choices for sources and receivers of probe packets

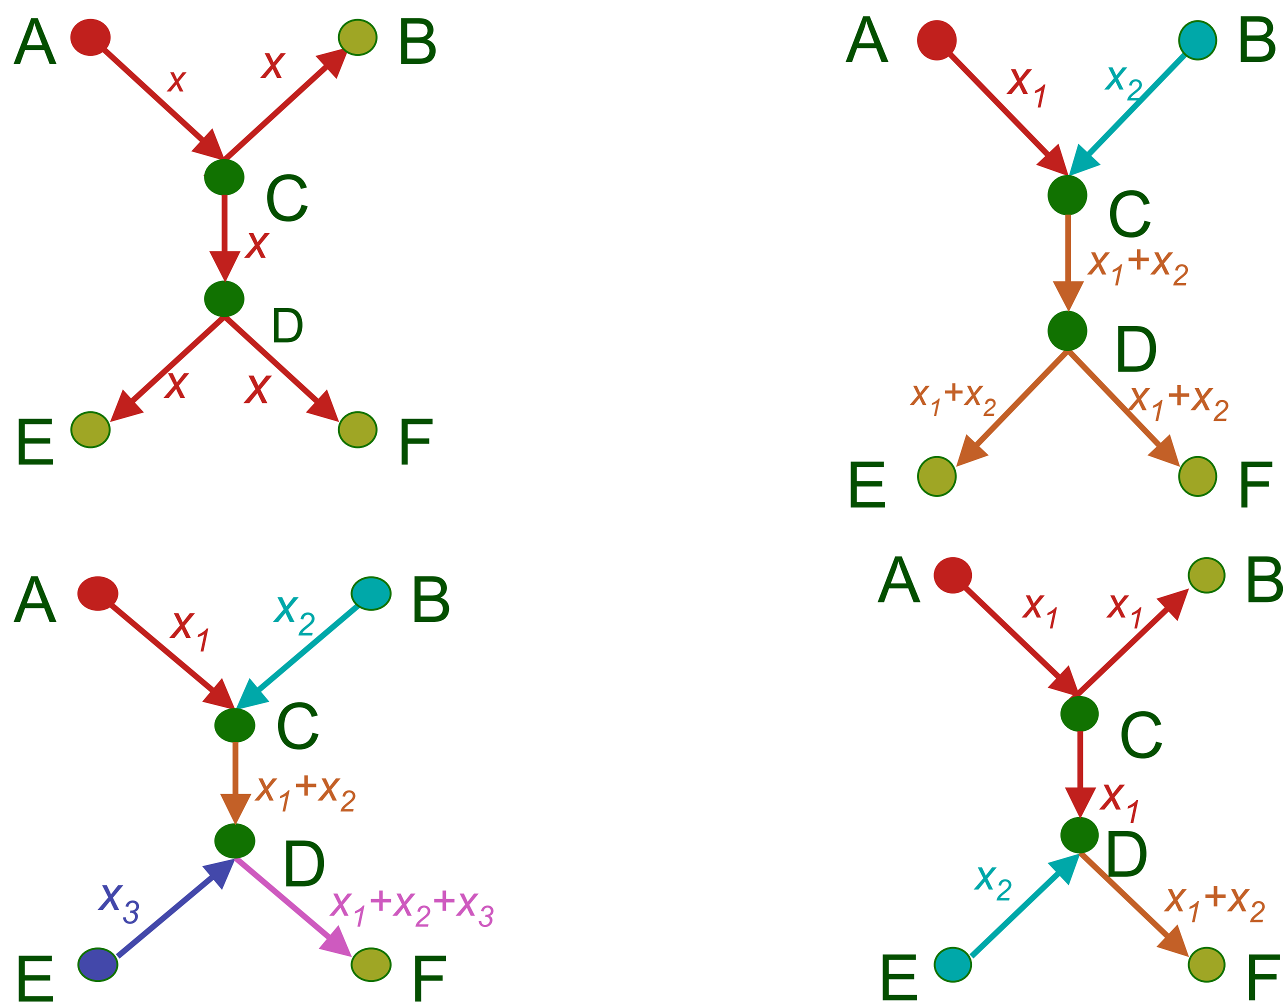

Comparing all four cases
(CR bound - same loss prob. on all links)

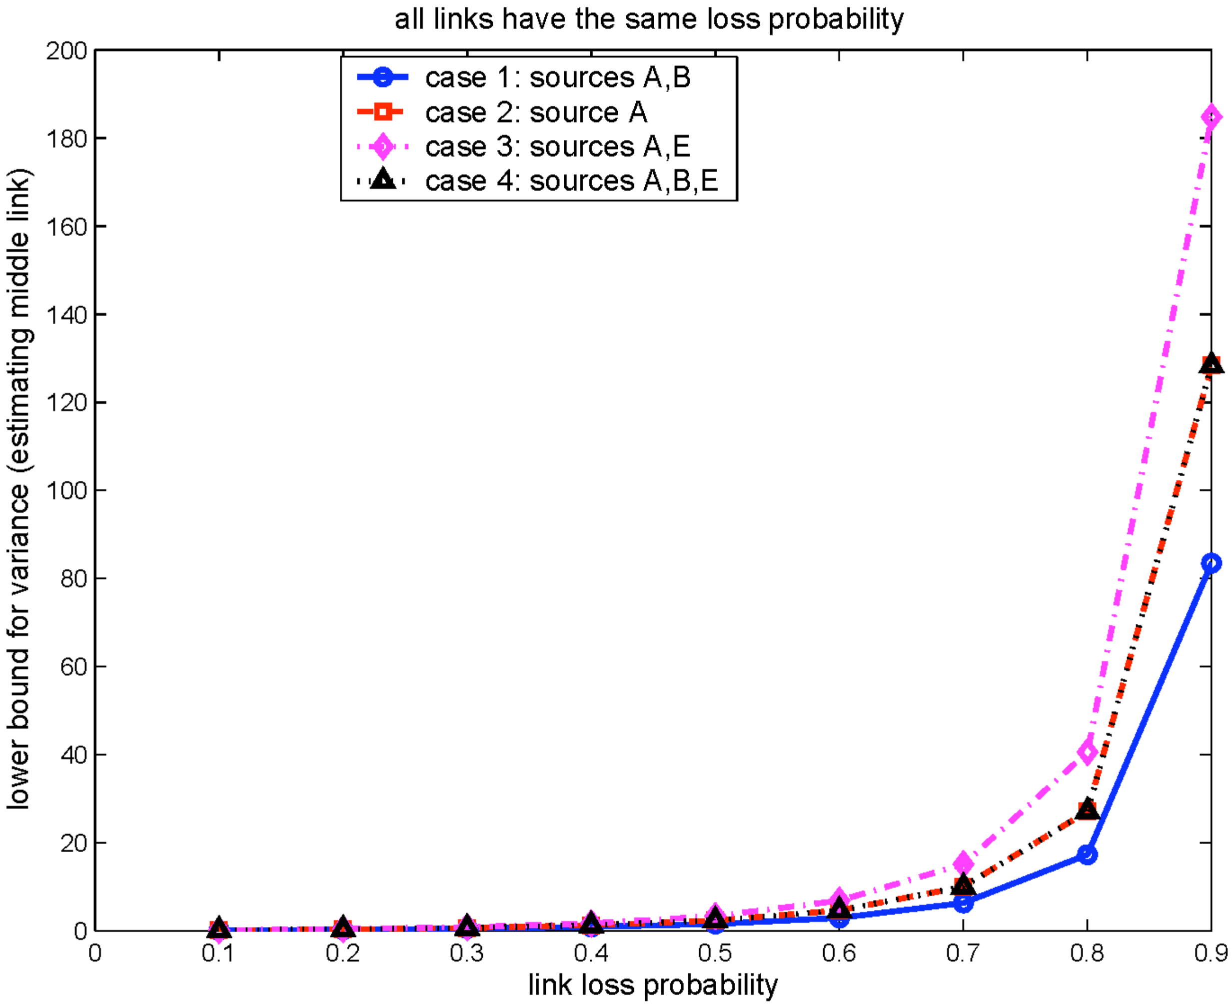

Estimating all links in general topologies.
Questions

- How many sources (and receivers)?
- Where to place them?
- What estimator to use?

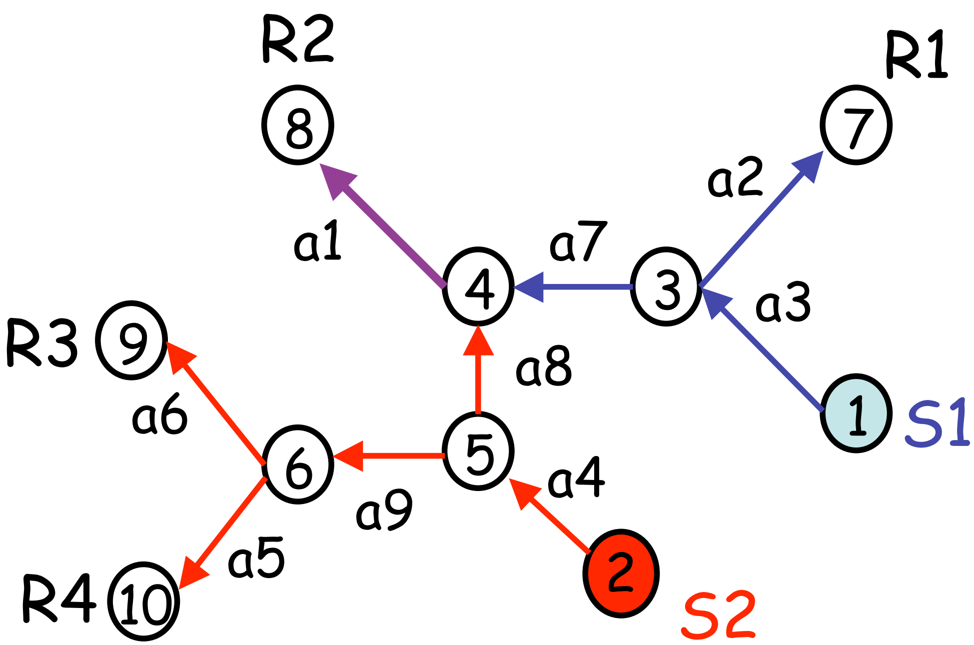

Belief Propagation

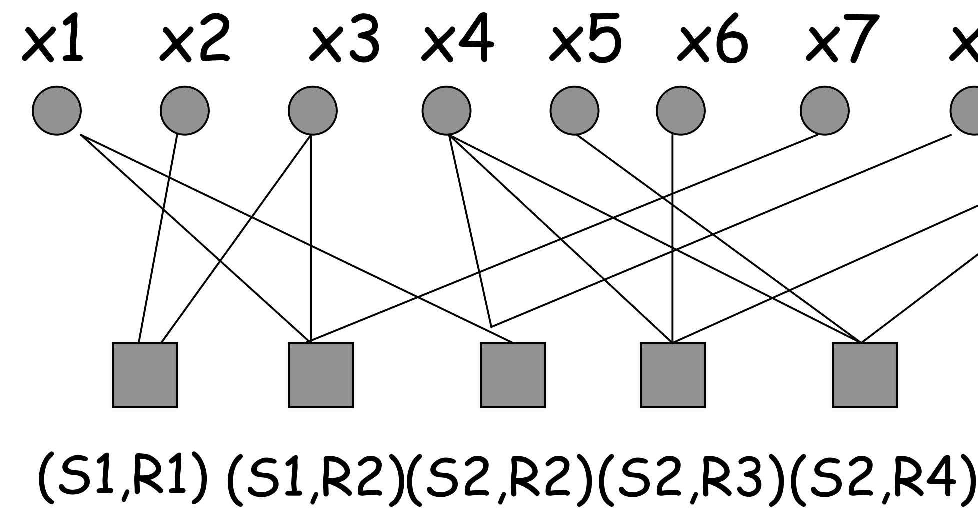

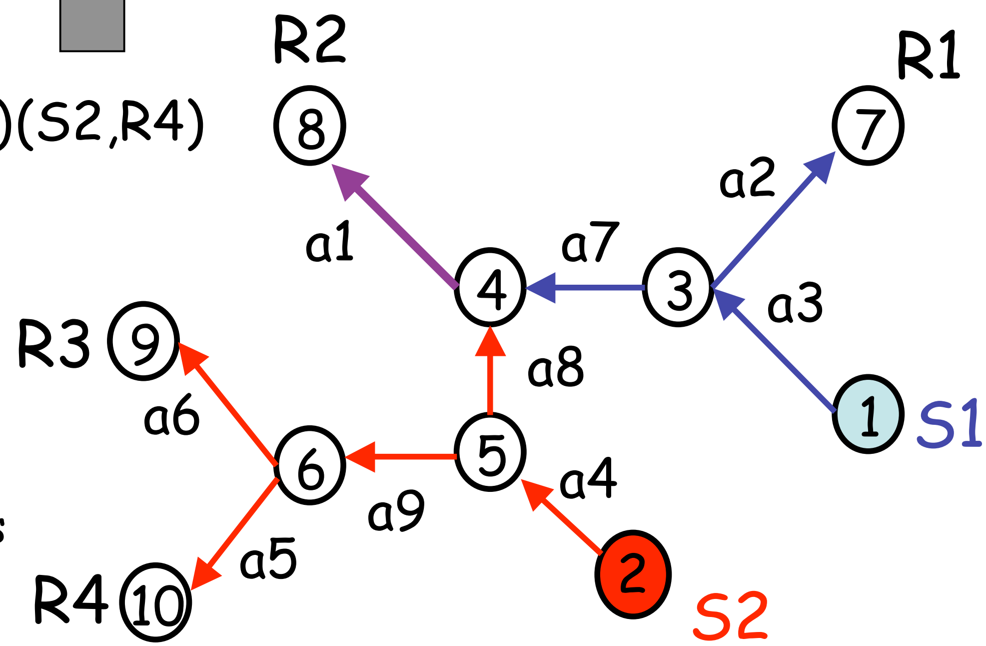

[Mao, Kschischang, Li, Pasupathy, ``A factor graph approach to link loss monitoring in WSN", JSAC 2005]

Estimation Accuracy Metrics
(all links E)

- Entropy Measure:

$$ENT = \sum_{e} \log E\left[\left|\alpha_e - \hat{\alpha}_e\right|^2\right]$$

- where
  - α_e : link loss rate for link e
  - $\hat{\alpha}_e$: estimated link loss rate for link e

Simulations Results

Compare:

- Single source (S1) multicast using MLE
- Two sources (S1,S2) and network coding (at C) using suboptimal estimation

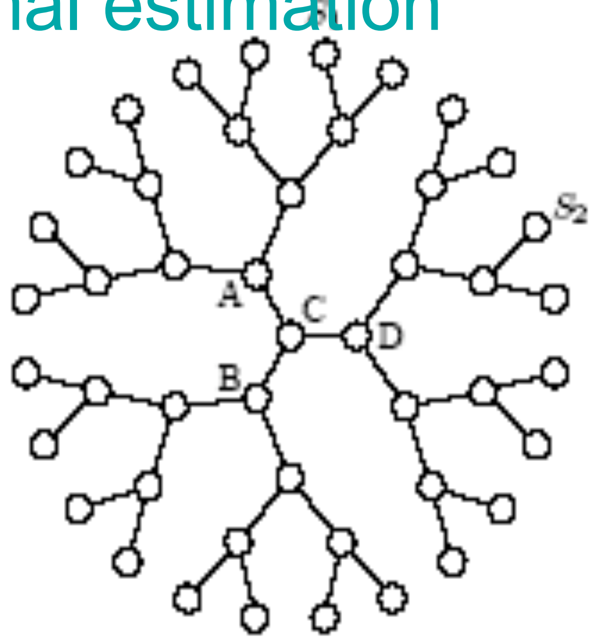

Simulation Results

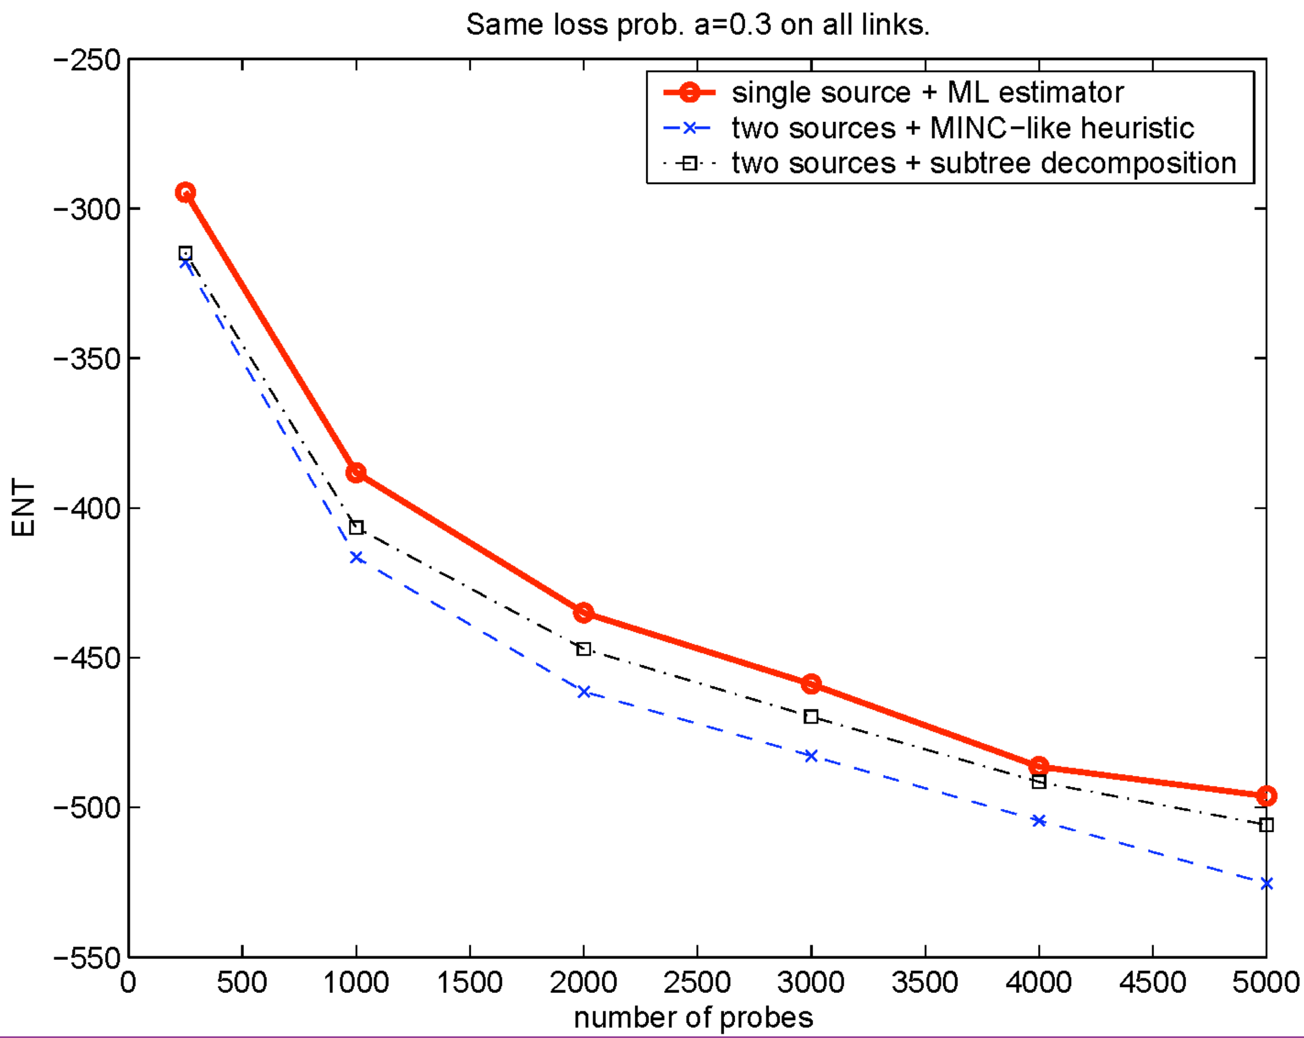

2 sources (with suboptimal estimation) do better than 1 source (with MLE)

Simulation Results

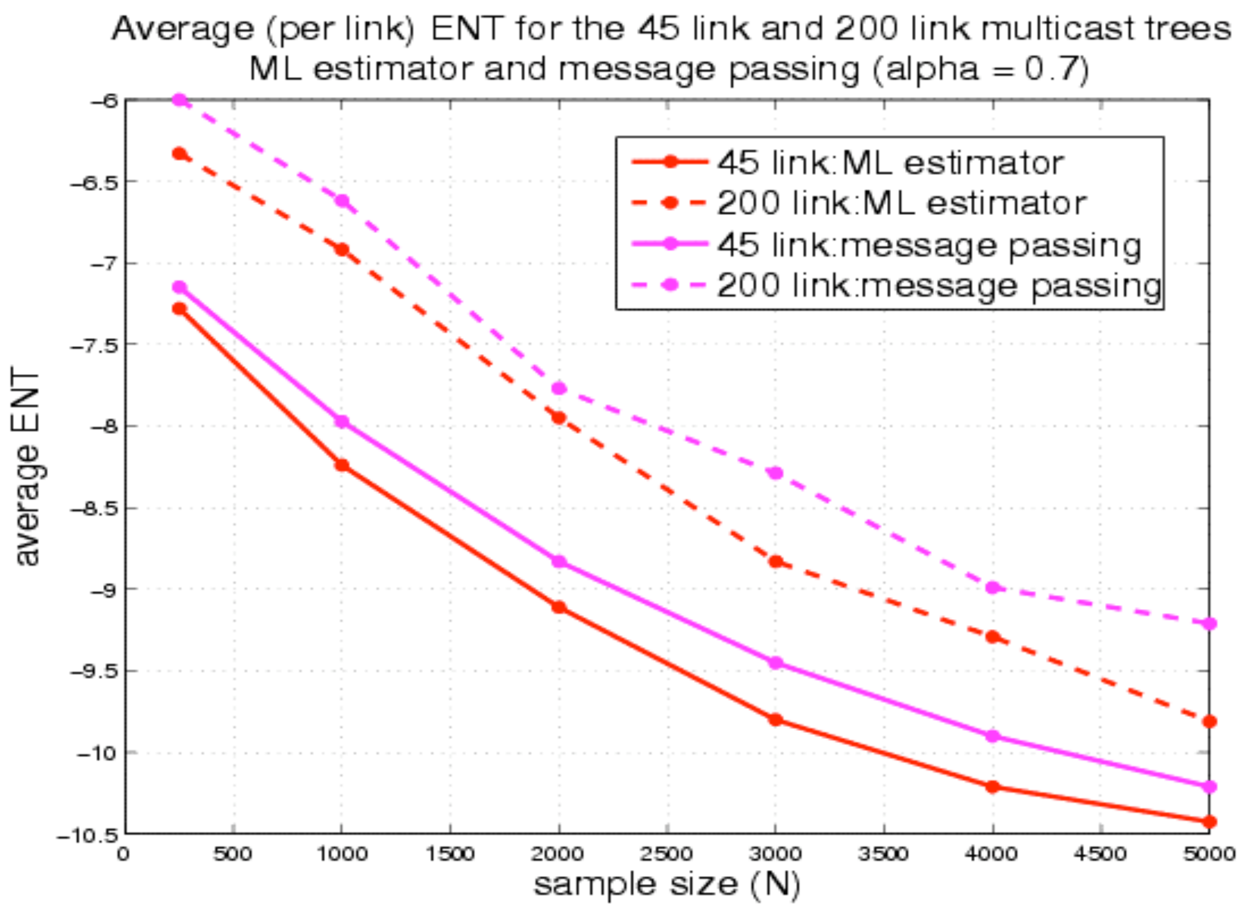

BP approximates the MLE

Applications of Network Coding

- Ad-hoc wireless networks
- Content delivery in P2P networks
- Network tomography
- Sensor networks
- Security
- Chip design
- .....

Data collection in sensor networks

Sensor nodes are static.
Each sensor node observes an independent random variable.

Data collection in sensor networks

Sensor nodes are static.
Each sensor node observes an independent random variable.

Phase 1: Each sensor node broadcasts m times.

Phase 2: A mobile collector queries k sensor nodes uniformly at random.

Data collection in sensor networks

Sensor nodes are static.
Each sensor node observes an independent random variable.

Phase 1: Each sensor node broadcasts m times.

Phase 2: A mobile collector queries k sensor nodes uniformly at random.

What value of k is necessary in order to collect all information

Phase 1:
each node transmits m times

Forwarding:
nodes randomly select and transmit
one of the symbols they have collected.

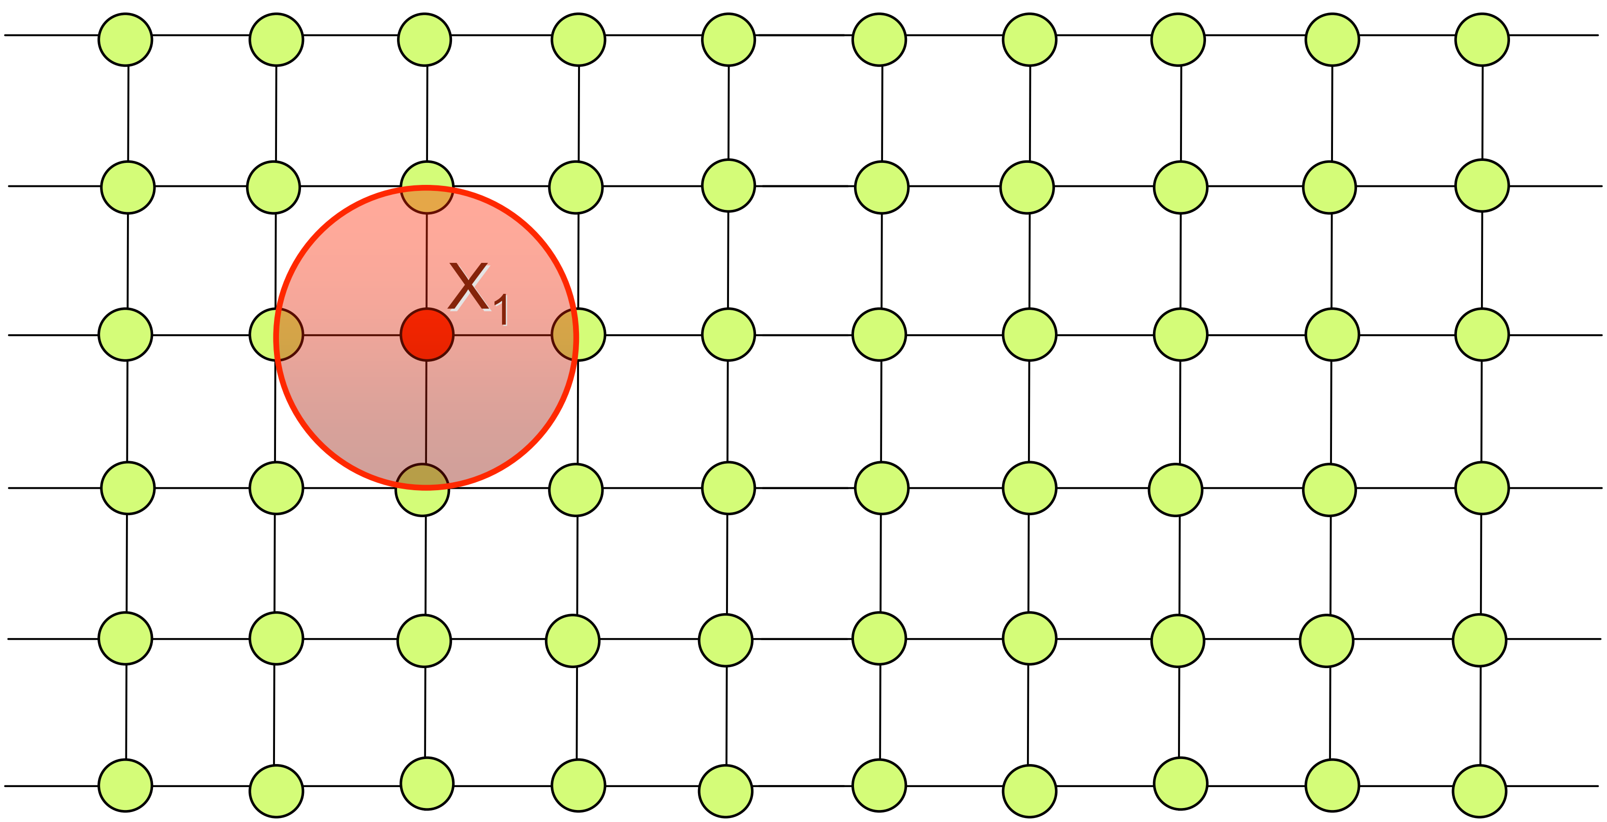

Phase 1: each node transmits m times

Forwarding:
nodes randomly select and transmit
one of the symbols they have collected.

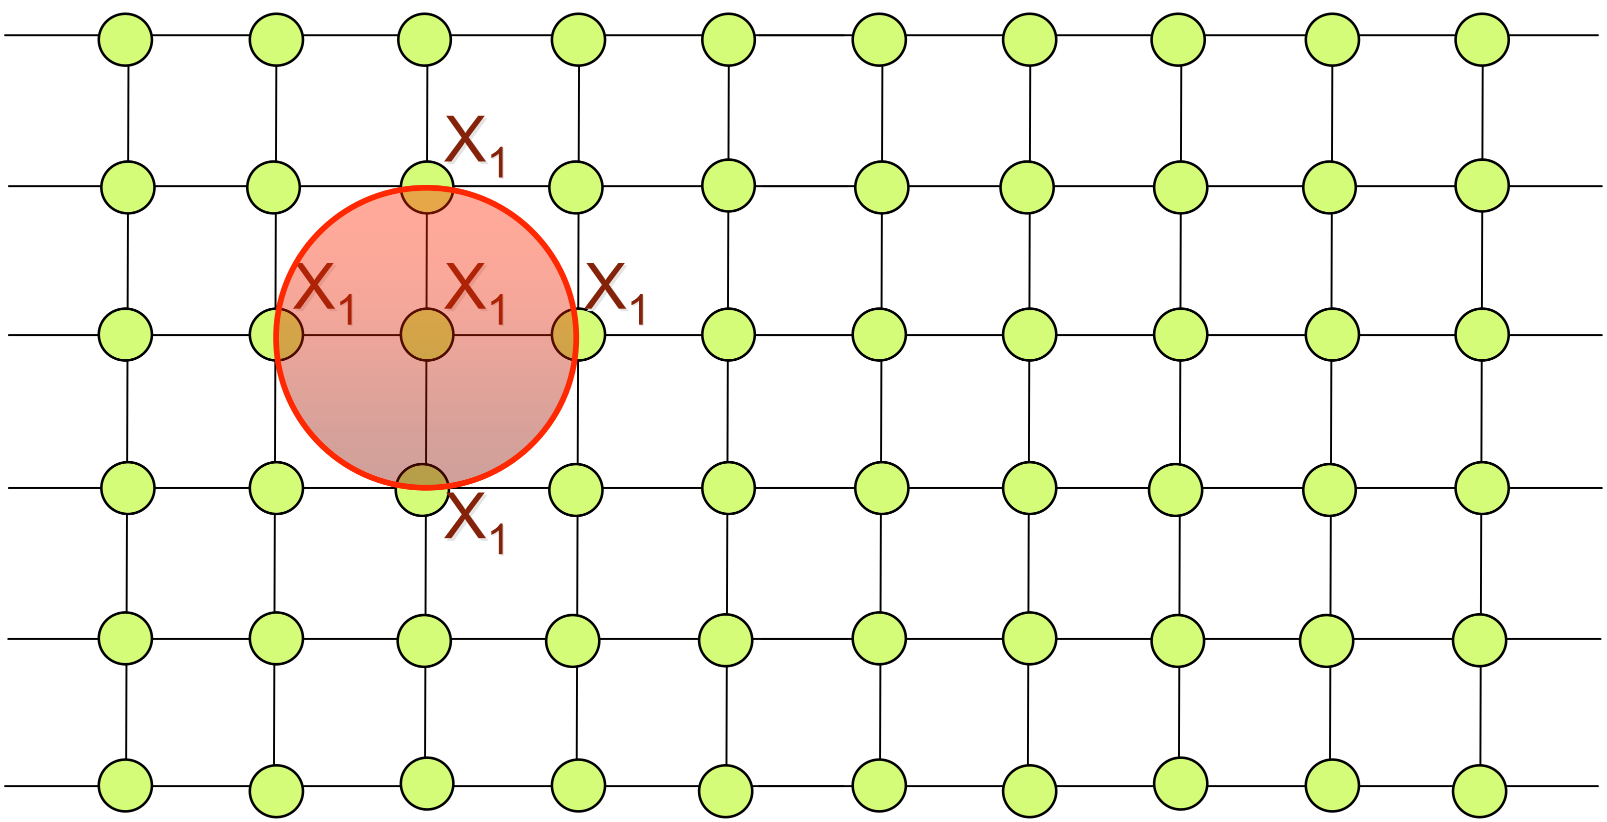

Phase 1: each node transmits m times

Forwarding:
nodes randomly select and transmit
one of the symbols they have collected.

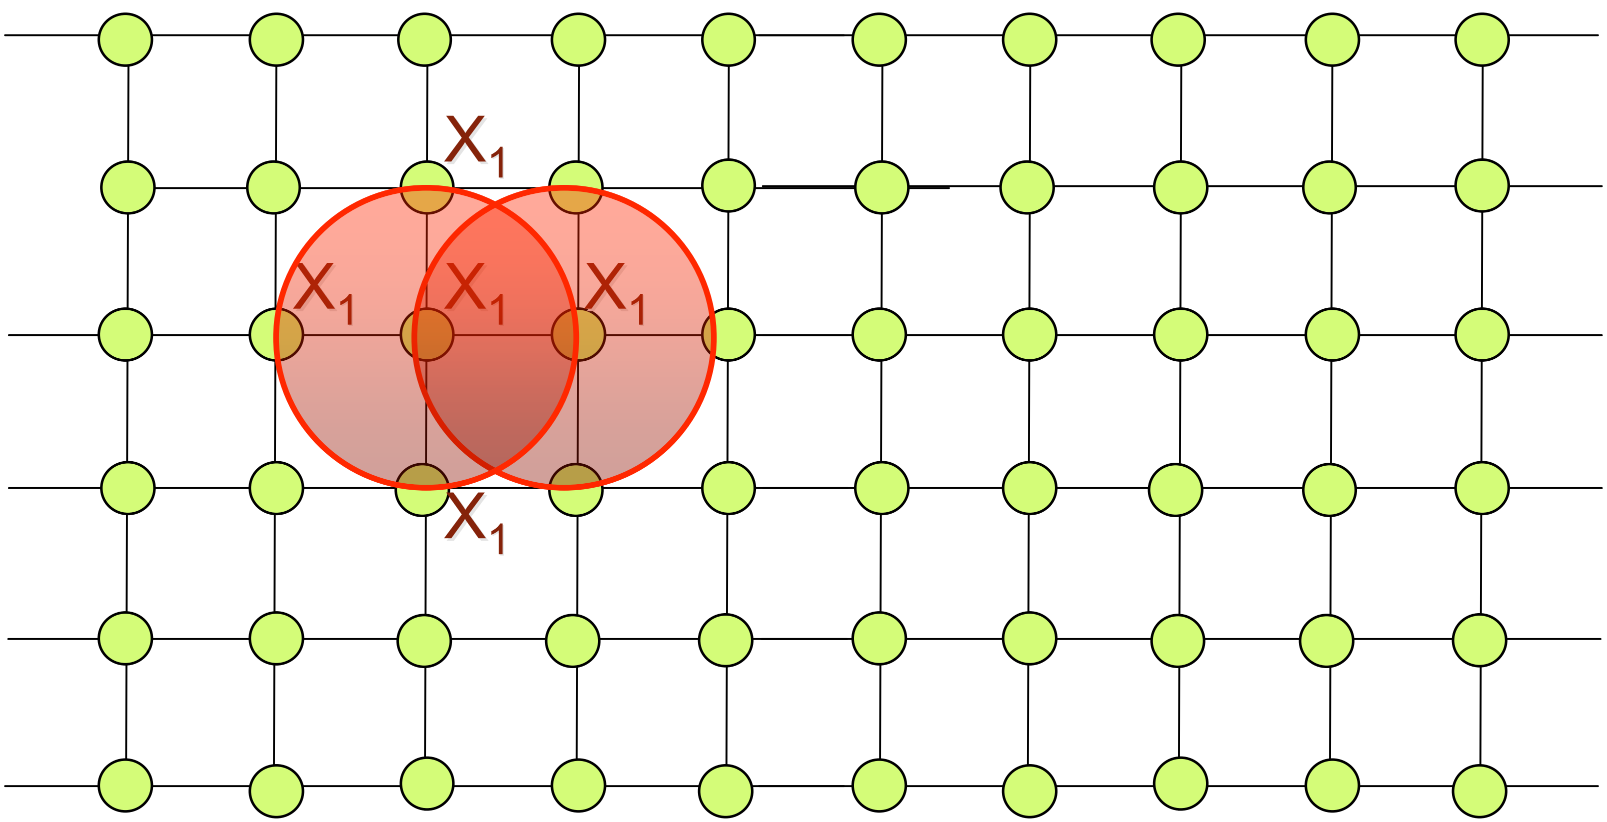

Phase 1: each node transmits m times

Forwarding: nodes randomly select and transmit one of the symbols they have collected.

Each node receives all information from the m nodes within distance $\Theta(\sqrt{m})$

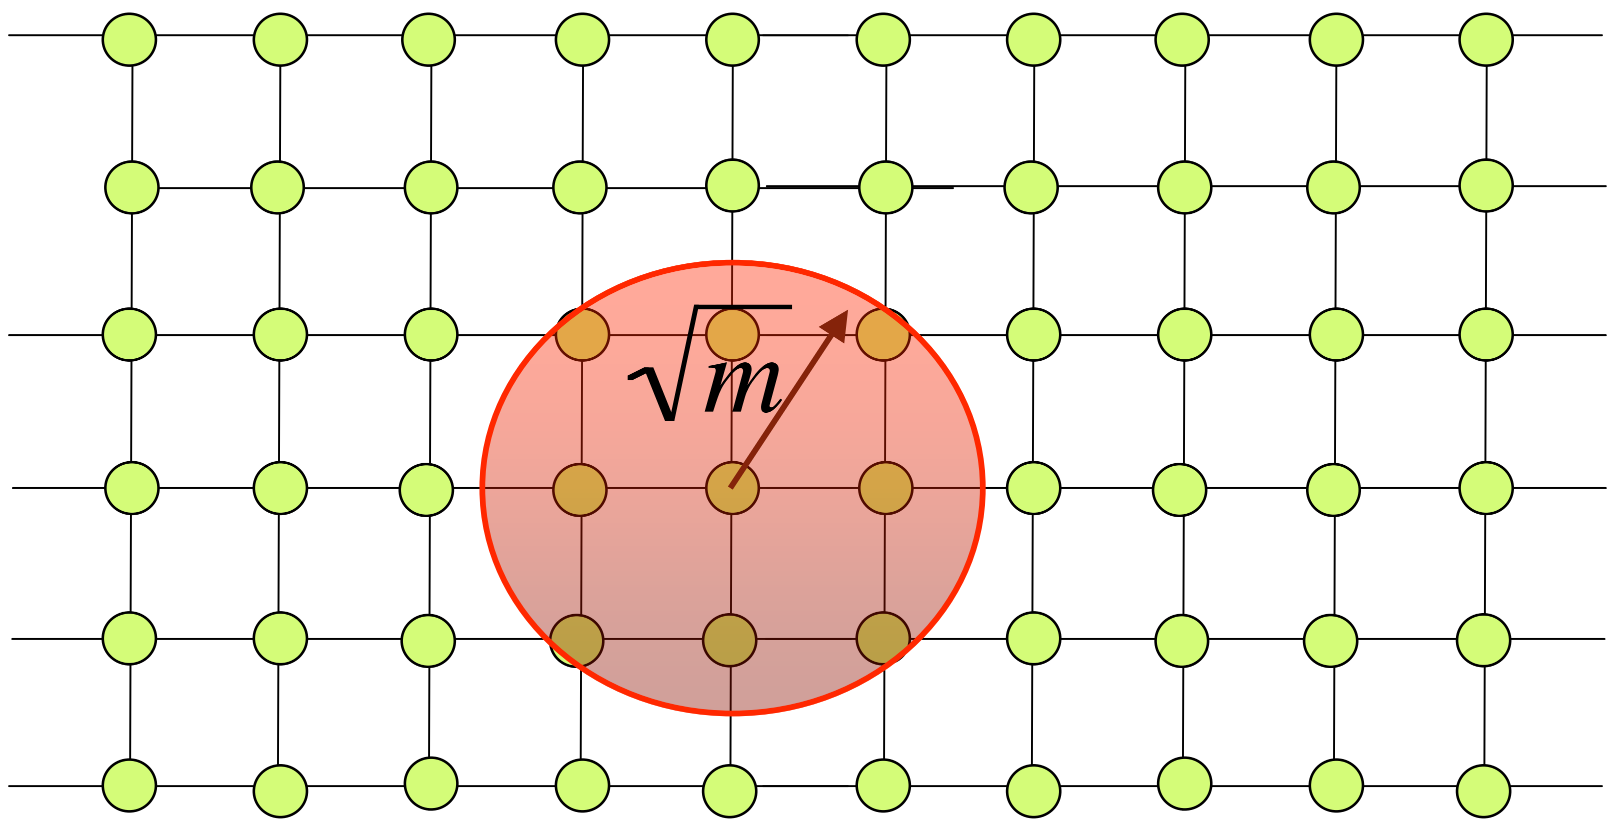

Phase 1:
each node transmits m times

Network coding: nodes transmit a random linear combination of the previously received symbols.

Each node receives all information from the m^2 nodes within distance $\Theta(m)$

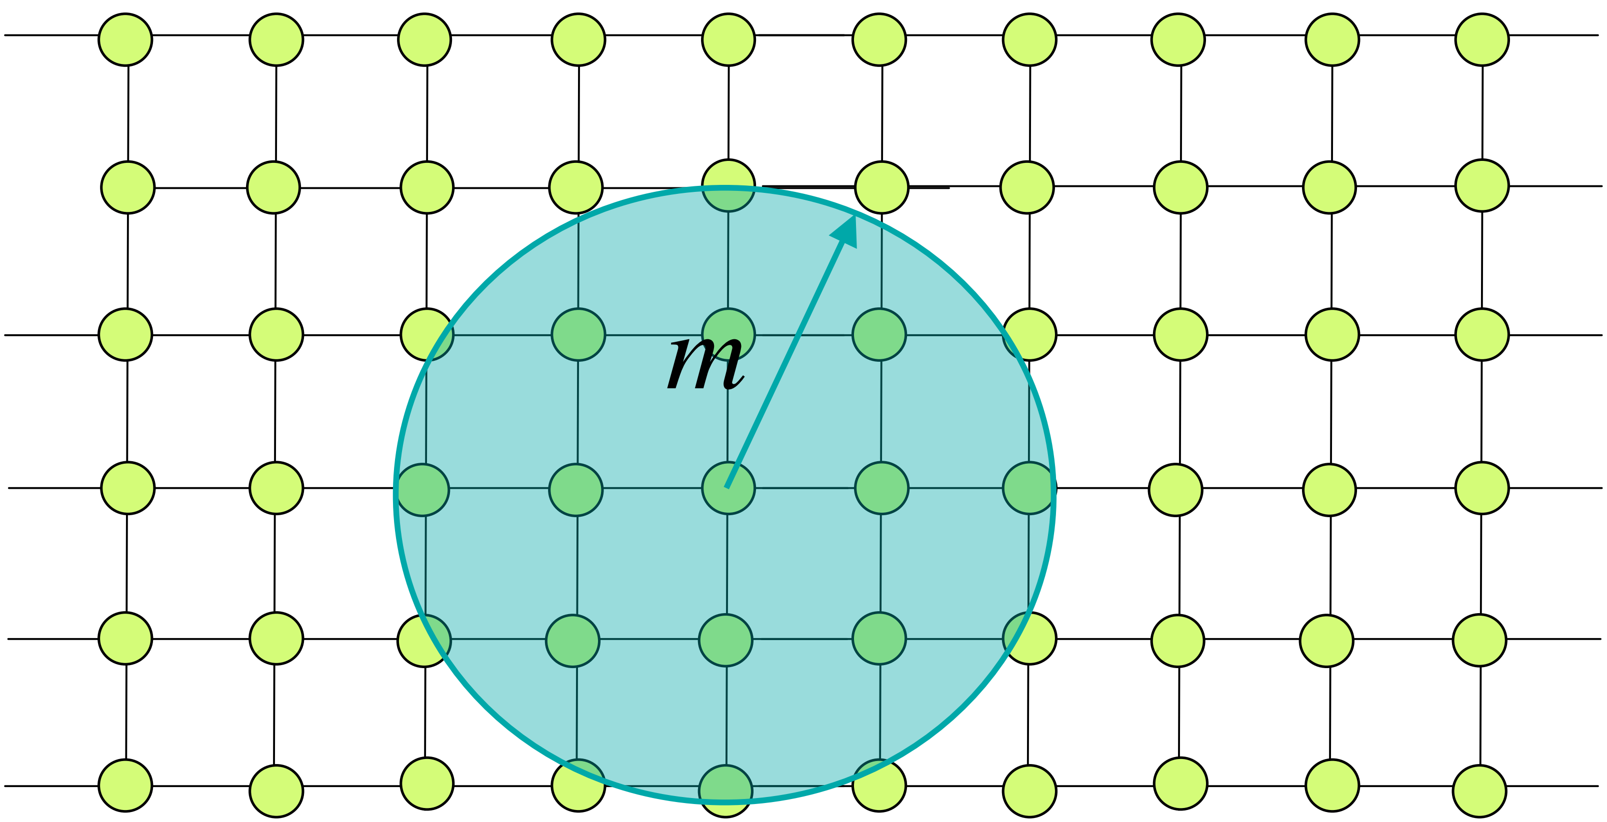

Phase 2: collector randomly queries k nodes

$$k \geq \frac{n}{4m+1}$$

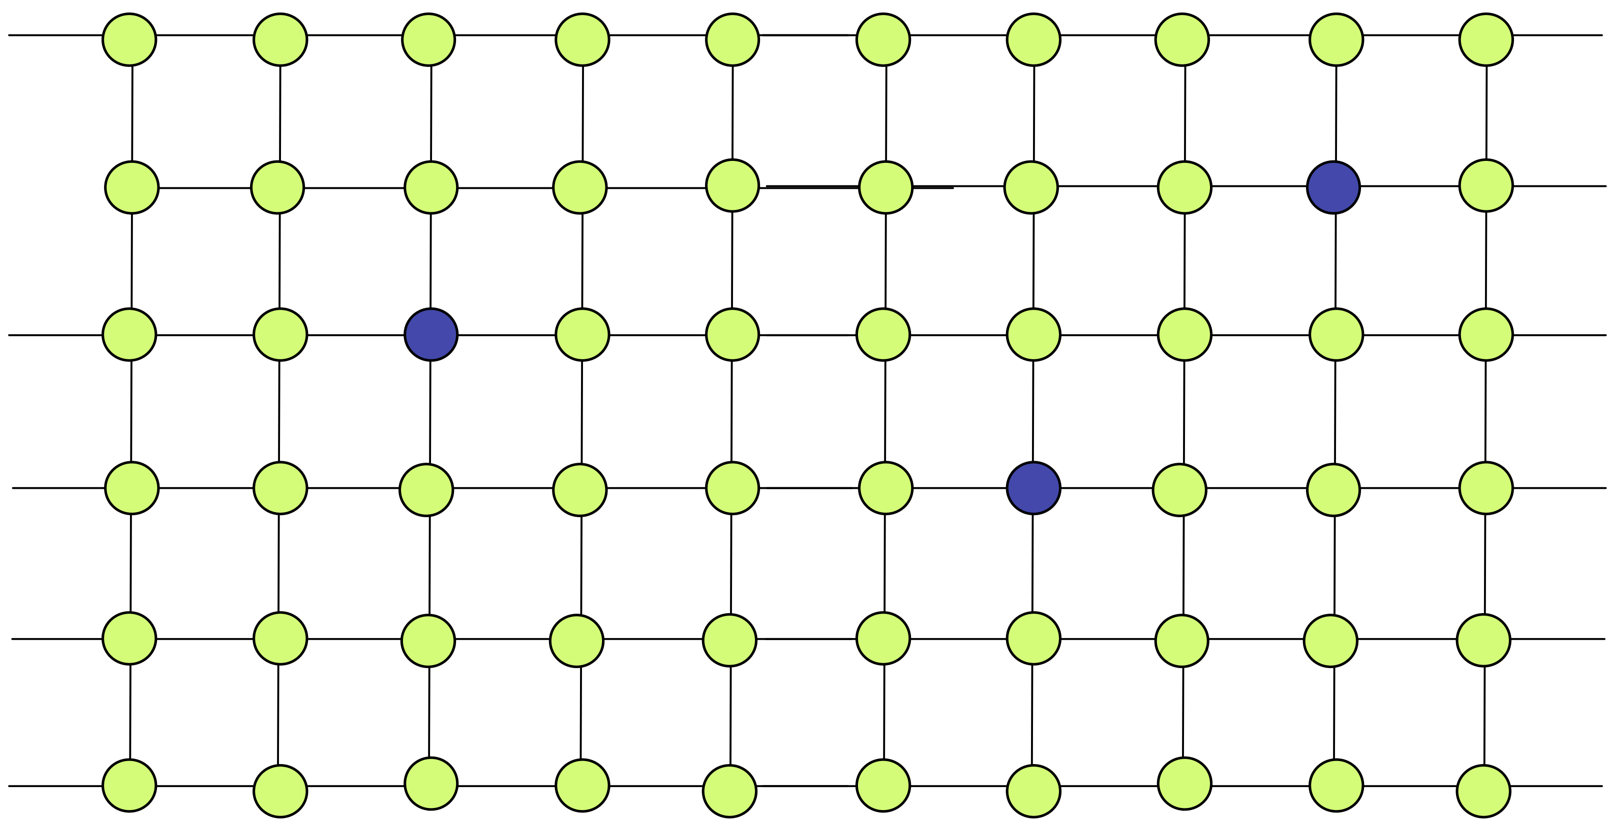

Phase 2: collector randomly queries k nodes

$$k \geq \frac{n}{4m+1}$$

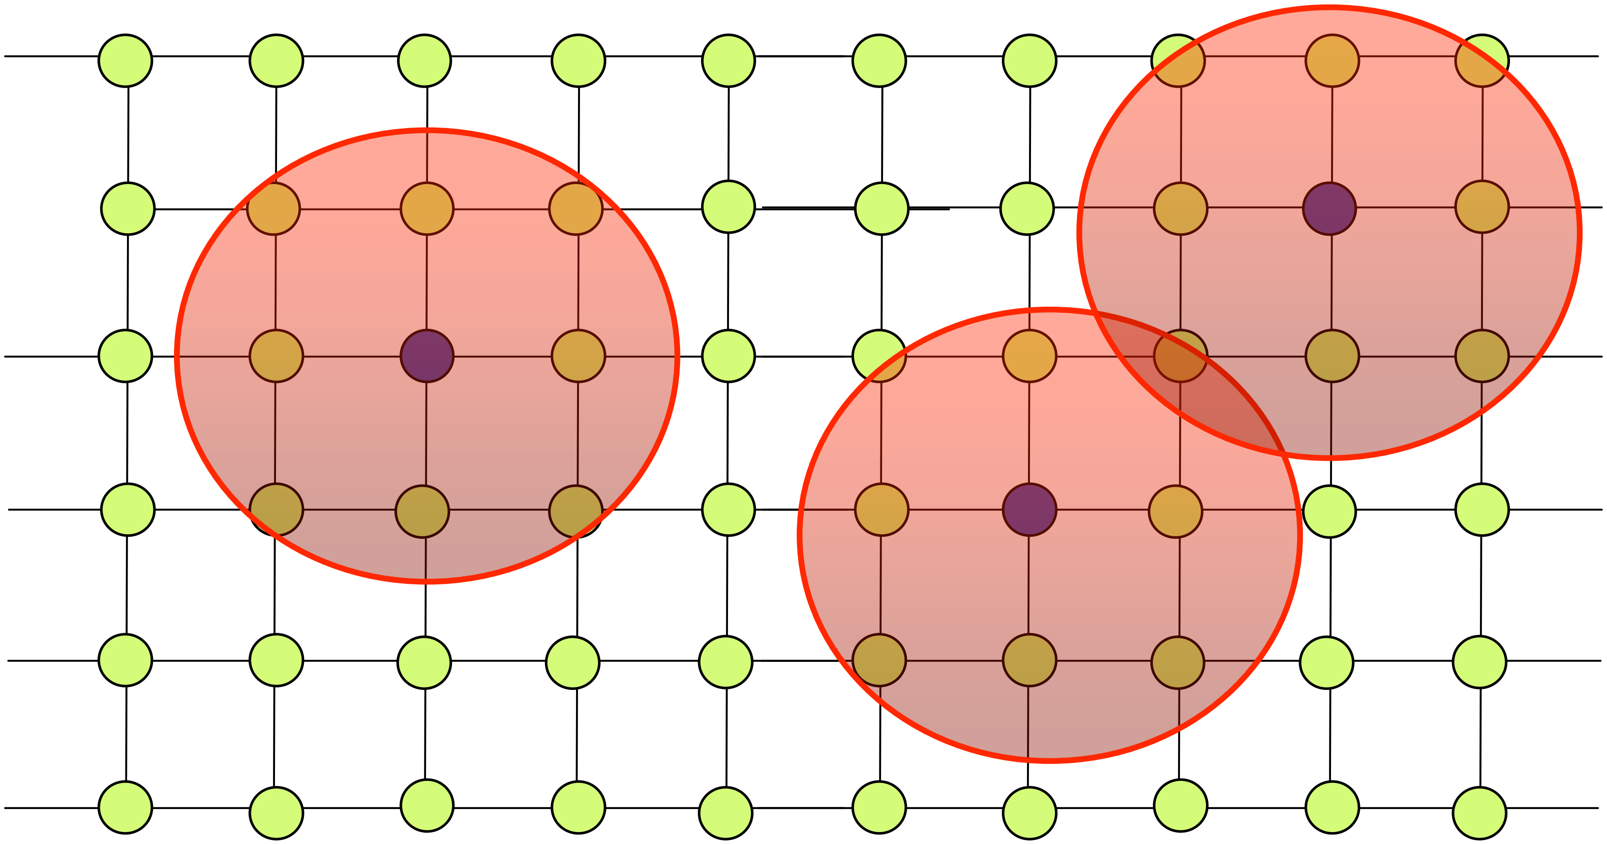

Phase 2: collector randomly queries k nodes

$$k \geq \frac{n}{4m+1}$$

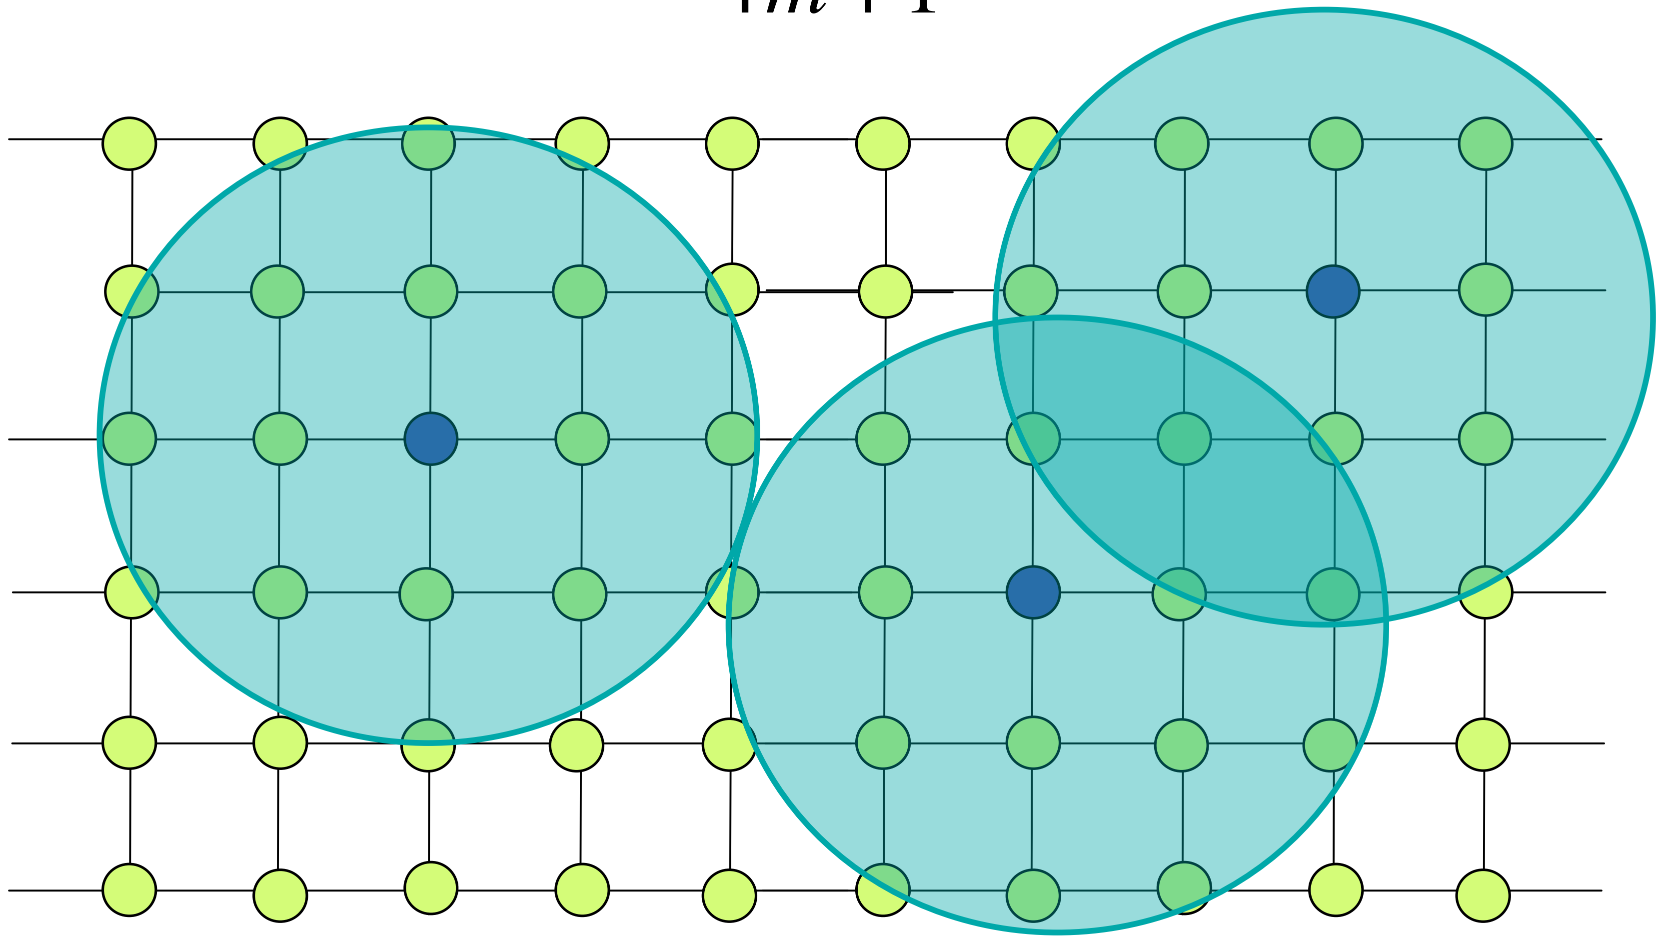

Probability that node is not covered by a disk

$$(1-\frac{r^2}{n})^k \approx e^{-k\frac{r^2}{n}} \approx \frac{1}{n} \Rightarrow k \geq$$

$$k \geq \frac{n}{4m+1}$$

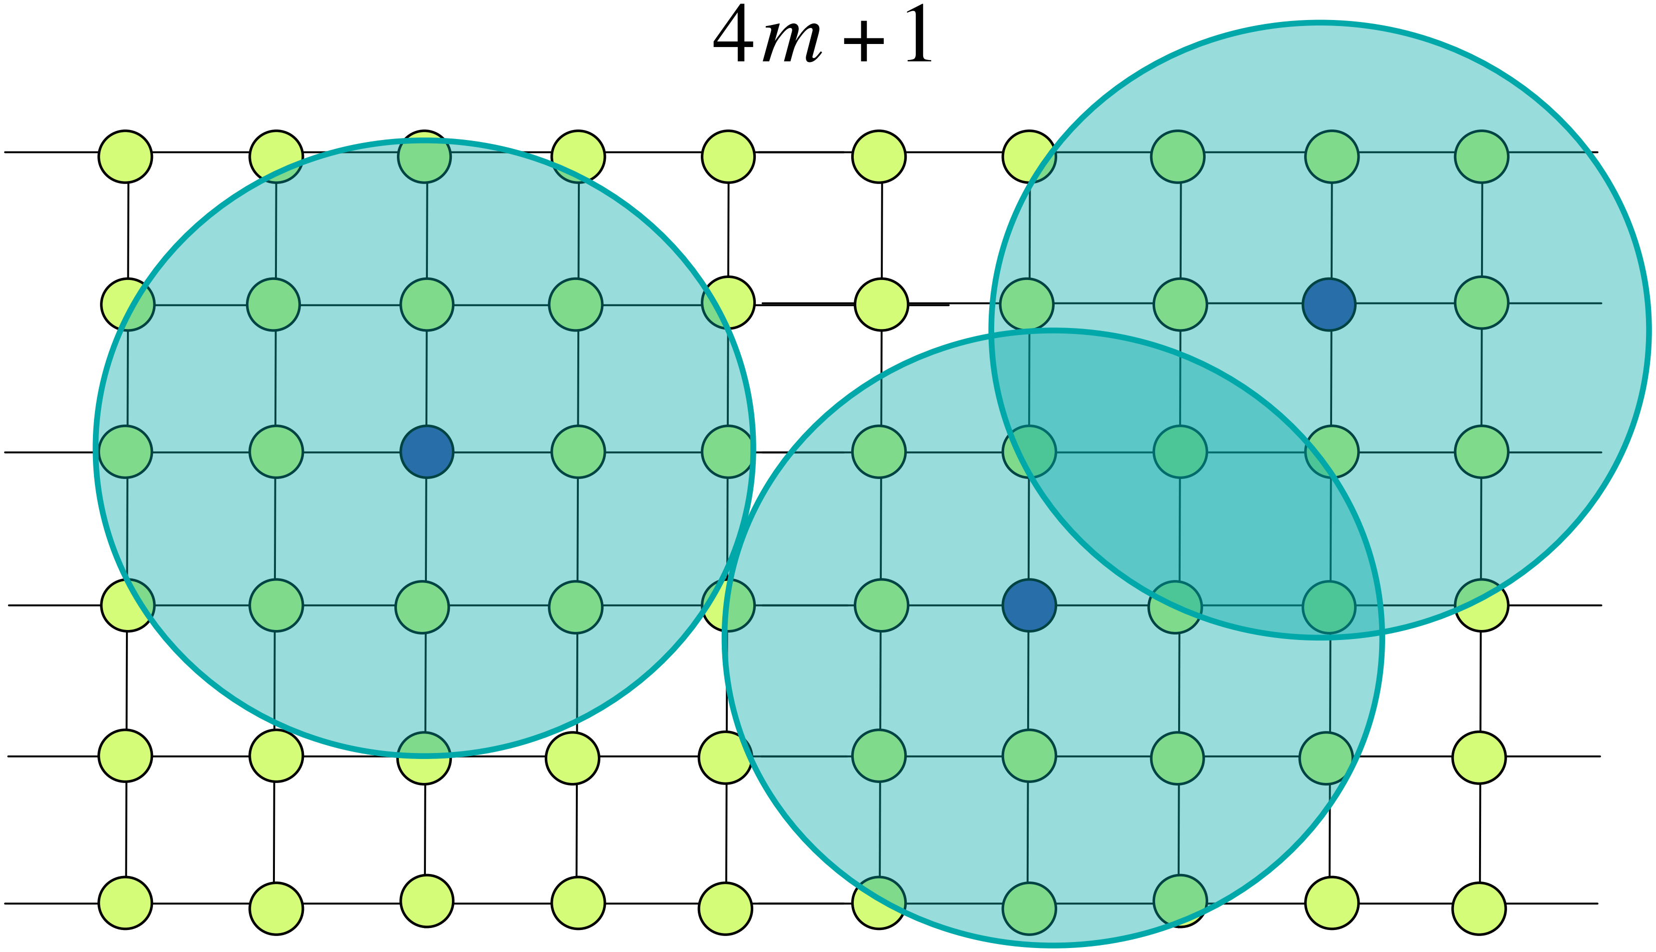

Simulation results: random network

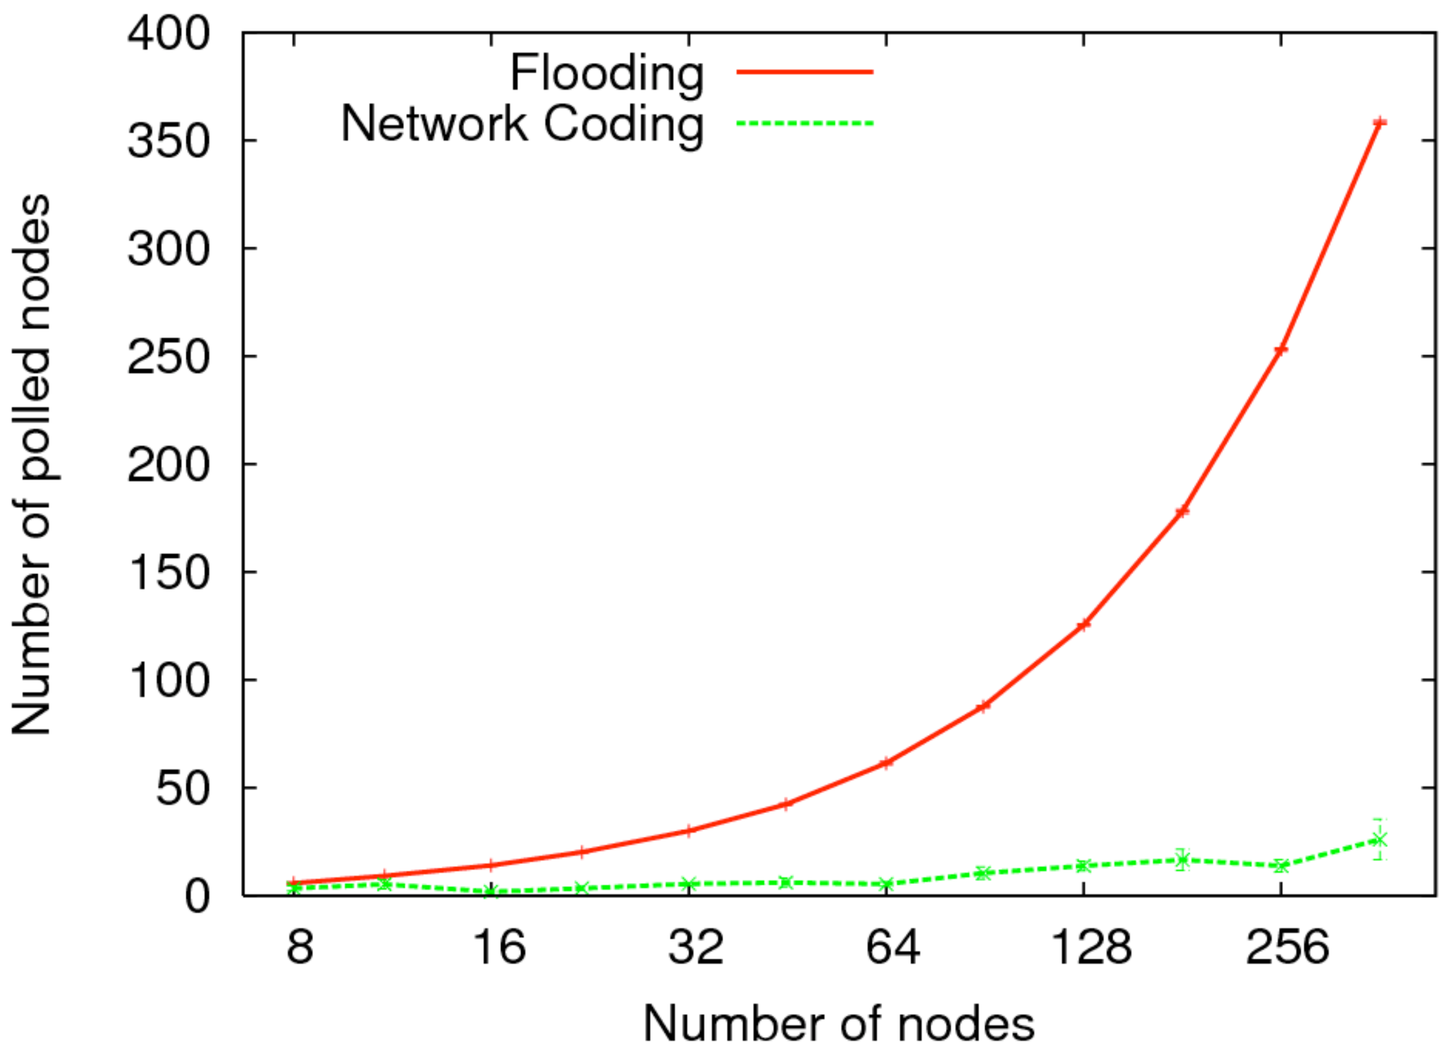

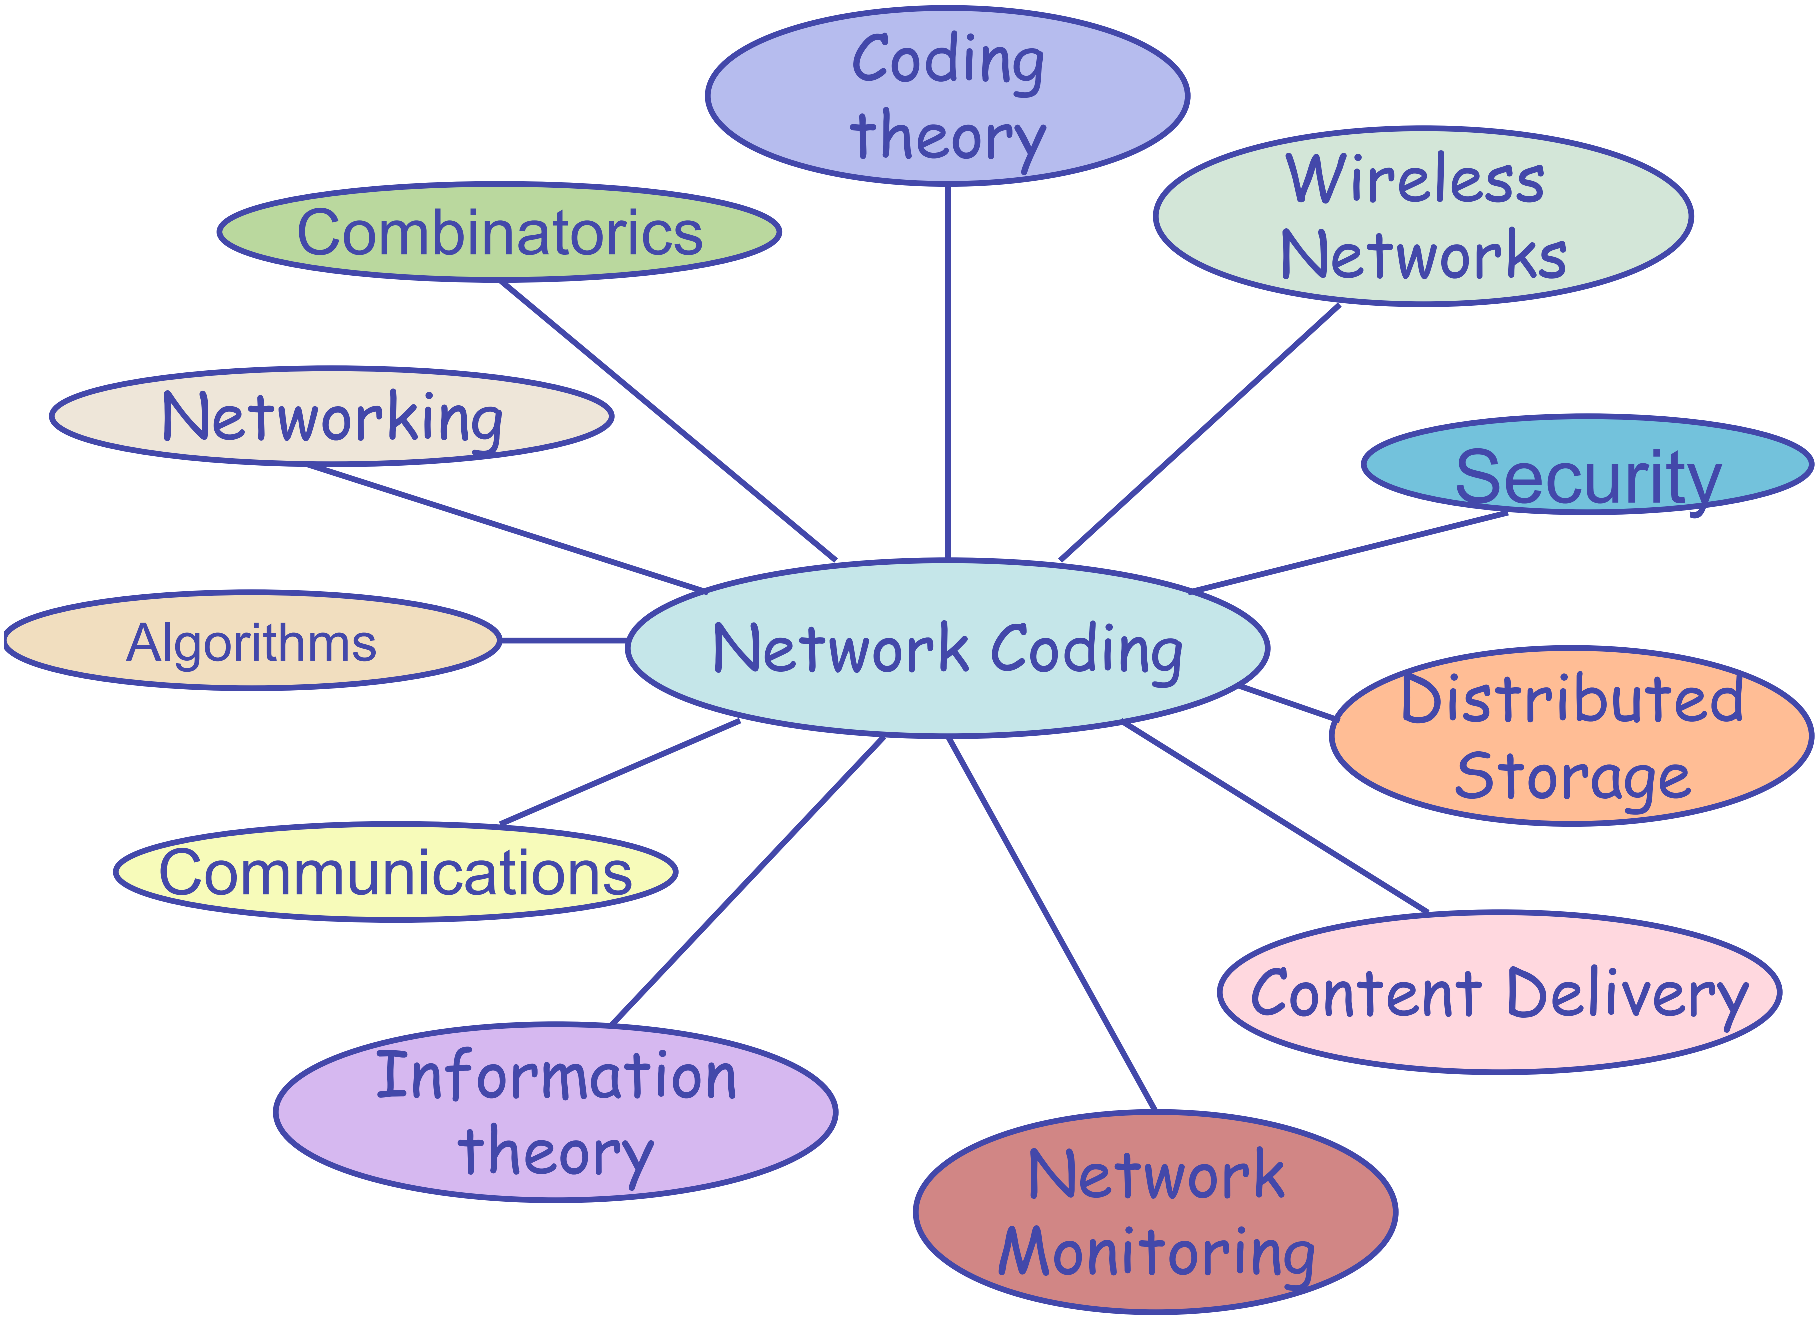
Coding theory
Combinatorics
Wireless Networks
Networking
Security
Algorithms
Network Coding
Distributed Storage
Communications
Content Delivery
Information theory
Network Monitoring