

DIMACS tutorial

Network coding: an Algorithmic Perspective

Tracey Ho - California Institute of
Technology

Alex Sprintson - Texas A&M University

Tutorial outline

Part I

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Tutorial outline

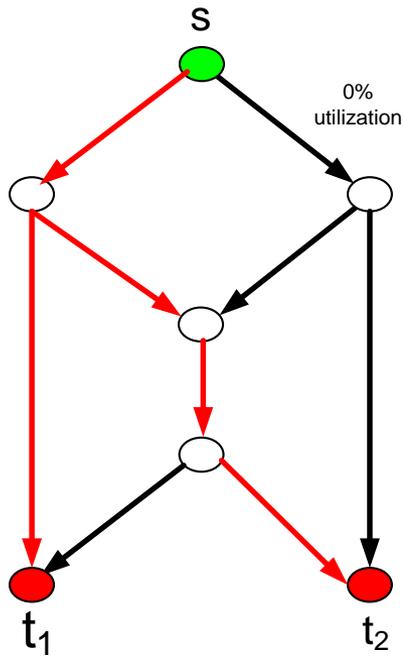
Part I I

- Coding Advantage
- Encoding Complexity
- Network Erasure Correction
- Open problems

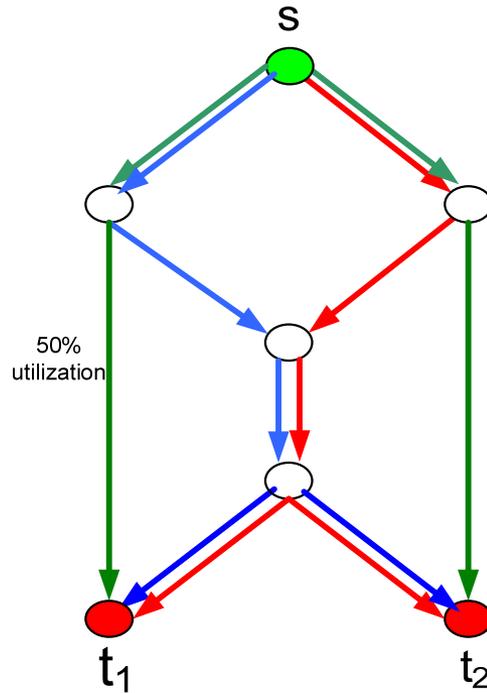
Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

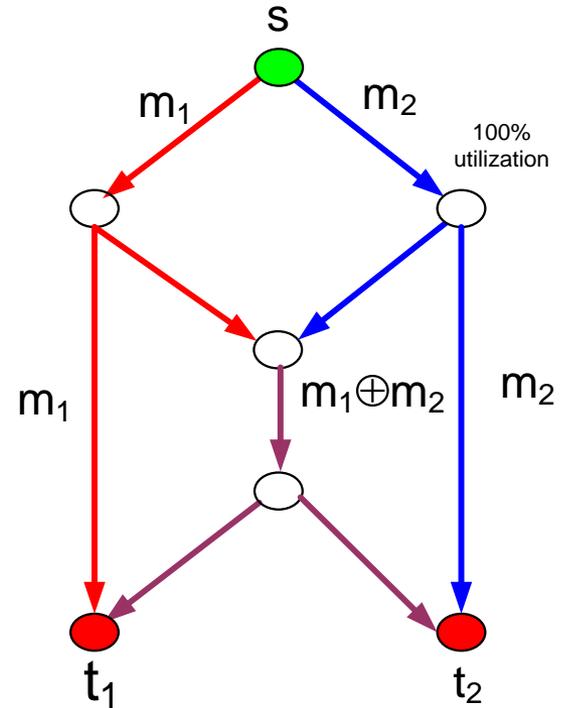
Wired multicast



One message per time unit



1.5 messages per time unit

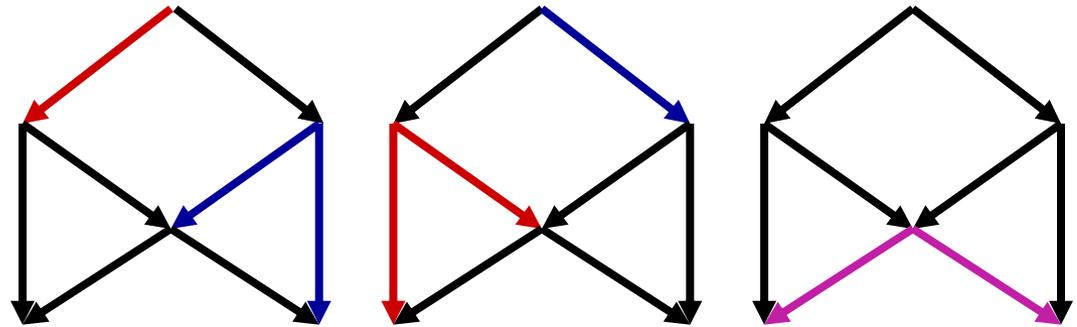
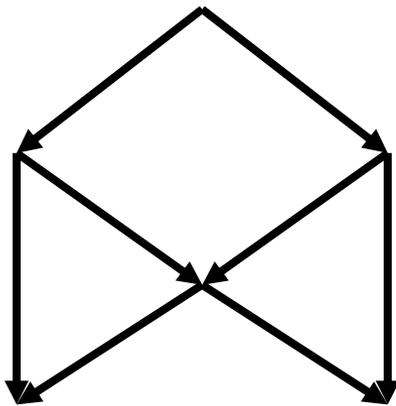


2 messages per time unit

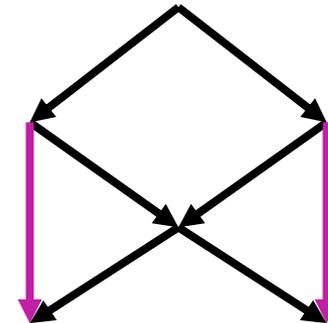
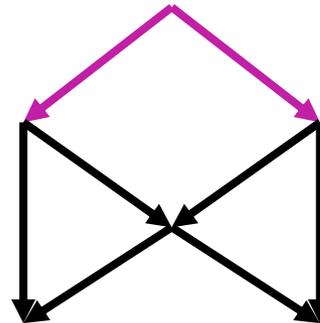
R. Ahlswede et al., "Network information flow," *IEEE Trans. Inform. Theory*, IT-46: 1204-1216, 2000.

Wireless multicast

- Static wireless model with fixed link rates and a half-duplex constraint
- E.g. [SaEp05]



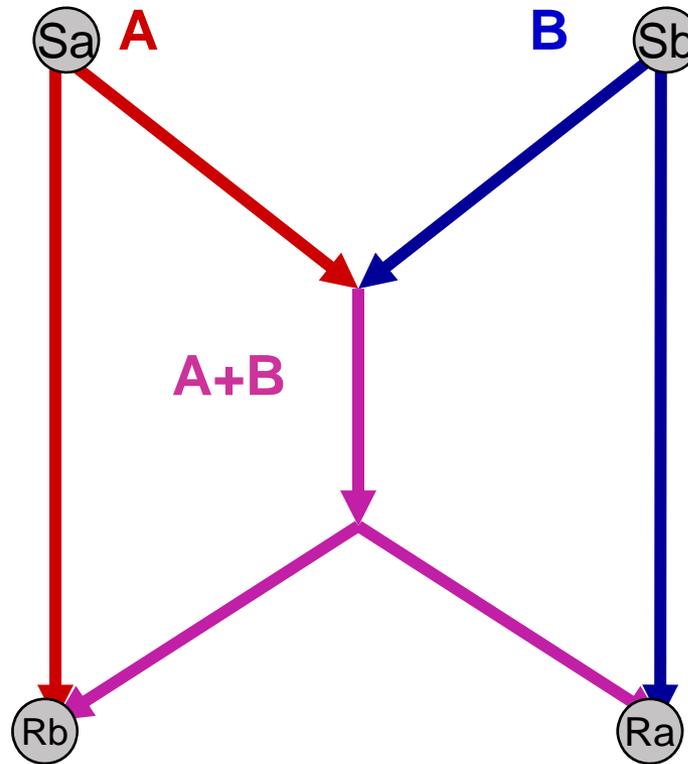
Coding



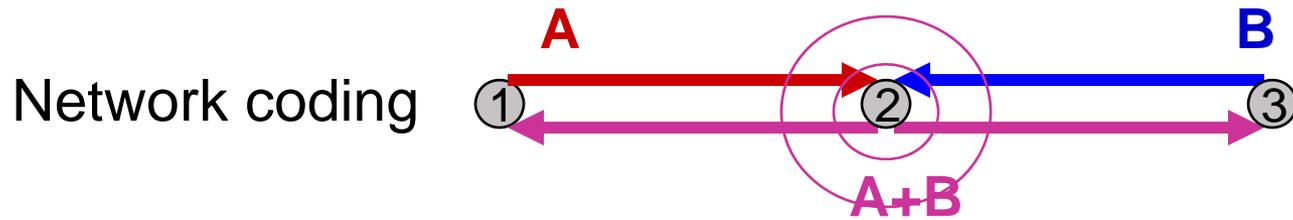
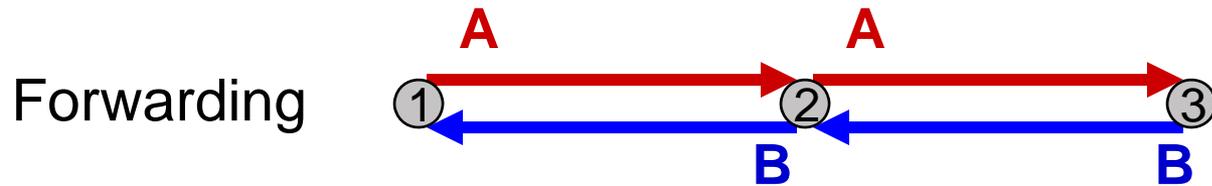
Routing

$$\frac{\text{Coding rate}}{\text{Routing rate}} = \frac{4}{3}$$

Wired unicasts



Wireless unicasts



R. W. Yeung and Z. Zhang, "Distributed source coding for satellite communications," *IEEE Trans. Inform. Theory*, IT-45: 1111-1120, 1999.

Problem description

μ input random processes at v :

$$\mathcal{X}(v) = \{X(v, 1), X(v, 2), \dots, X(v, \mu(v))\}$$

ν Output random processes at u :

$$\mathcal{Z}(u) = \{Z(u, 1), Z(u, 2), \dots, Z(u, \nu(u))\}$$

Random processes on edges: $Y(e)$

A connection:

$$c = (v, u, \mathcal{X}(v, u)), \mathcal{X}(v, u) \subseteq \mathcal{X}(v)$$

A connection is **established** if $\mathcal{Z}(u) \supset \mathcal{X}(v, u)$

Set of connections: \mathcal{C}

The pair $(\mathcal{G}, \mathcal{C})$ defines a **network coding problem**.

Problem statement

Is the problem $(\mathcal{G}, \mathcal{L})$ solvable?

How do we find a solution?

This is fairly idealized (synchronization, protocol, dynamic behaviour, error free operation,...) but gives insights into possible limits and opportunities.

Simplifying assumptions

$C(e) = 1$ (links have the same capacity)

$H(X(v, i)) = 1$ (sources have the same rate)

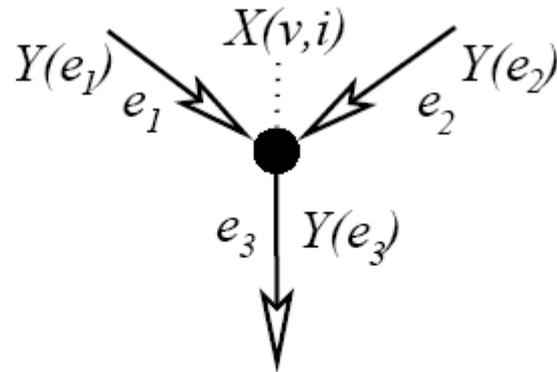
The $X(v, i)$ are mutually independent.

Vector symbols of length m elements in \mathbb{F}_{2^m} .

(\mathbb{F}_{2^m} is the finite field with 2^m elements we can add, subtract, divide and multiply elements in \mathbb{F}_{2^m} without going crazy!)

Scalar linear network coding

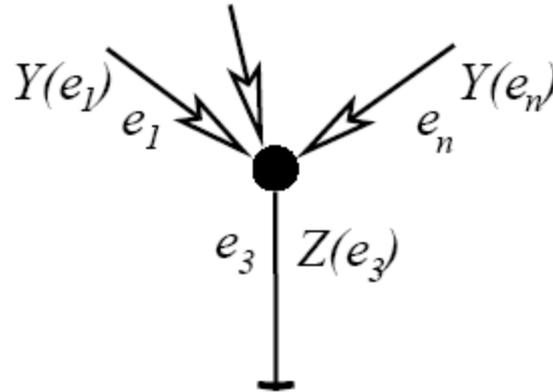
All operations at network nodes are linear!



$$Y(e_3) = \sum_i \alpha_i X(v, i) + \sum_{j=1,2} \beta_j Y(e_j)$$

Scalar linear network coding

At a receiver (terminal) node:

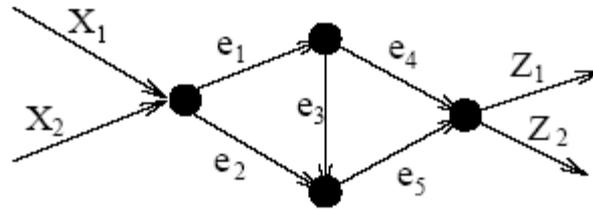


$$Z(v, j) = \sum_{j=1}^n \varepsilon_j Y(e_j).$$

Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

A simple example



$$Y(e_1) = \alpha_{1,e_1}X_1 + \alpha_{2,e_1}X_2$$

$$Y(e_2) = \alpha_{1,e_2}X_1 + \alpha_{2,e_2}X_2$$

$$Y(e_3) = \beta_{e_1,e_3}Y(e_1)$$

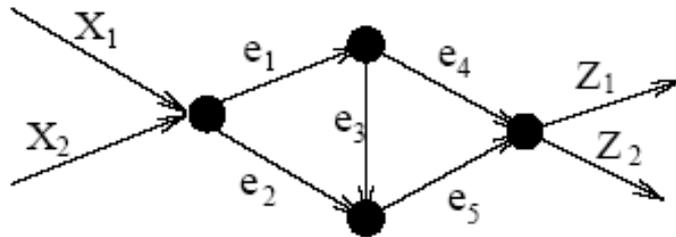
$$Y(e_4) = \beta_{e_1,e_4}Y(e_1)$$

$$Y(e_5) = \beta_{e_2,e_5}Y(e_2) + \beta_{e_3,e_5}Y(e_3)$$

$$Z_1 = \varepsilon_{e_4,1}Y(e_4) + \varepsilon_{e_5,1}Y(e_5)$$

$$Z_2 = \varepsilon_{e_4,2}Y(e_4) + \varepsilon_{e_5,2}Y(e_5)$$

Transfer matrix



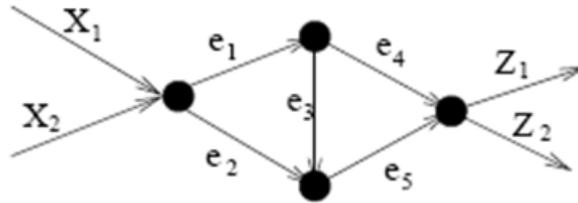
$$F = \begin{pmatrix} 0 & 0 & \beta_{e_1, e_3} & \beta_{e_1, e_4} & 0 \\ 0 & 0 & 0 & 0 & \beta_{e_2, e_5} \\ 0 & 0 & 0 & 0 & \beta_{e_3, e_5} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$A = \begin{pmatrix} \alpha_{1, e_1} & \alpha_{1, e_2} & 0 & 0 & 0 \\ \alpha_{2, e_1} & \alpha_{2, e_2} & 0 & 0 & 0 \end{pmatrix}$$

$$B = \begin{pmatrix} \varepsilon_{1, e_1} & \varepsilon_{1, e_2} & 0 & 0 & 0 \\ \varepsilon_{2, e_1} & \varepsilon_{2, e_2} & 0 & 0 & 0 \end{pmatrix}$$

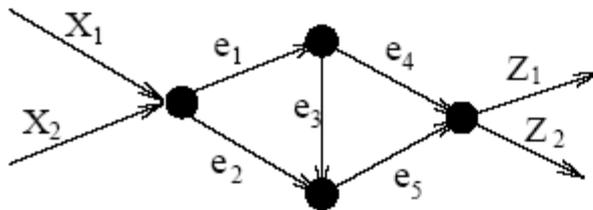
1

Transfer matrix



$$\begin{aligned}
 & \left[X_1 \quad X_2 \right] \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} \left(\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} + \begin{bmatrix} 0 & 0 & \beta_{e_1 e_3} & \beta_{e_1 e_4} & 0 \\ 0 & 0 & 0 & 0 & \beta_{e_2 e_5} \\ 0 & 0 & 0 & 0 & \beta_{e_3 e_5} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} + \right. \\
 & \left. + \begin{bmatrix} 0 & 0 & 0 & 0 & \beta_{e_1 e_3} \beta_{e_3 e_5} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{bmatrix} \right) \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} Z_1 & Z_2 \end{bmatrix}
 \end{aligned}$$

Transfer matrix



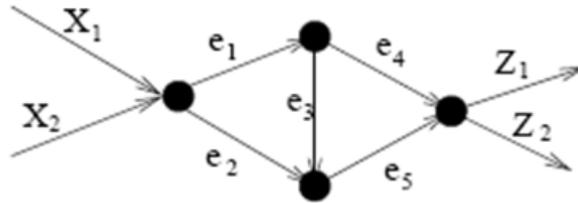
$$F = \begin{pmatrix} 0 & 0 & \beta_{e_1, e_3} & \beta_{e_1, e_4} & 0 \\ 0 & 0 & 0 & 0 & \beta_{e_2, e_5} \\ 0 & 0 & 0 & 0 & \beta_{e_3, e_5} \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

Summing the "path gains":

$$I + F + F^2 + \dots = (I - F)^{-1} = \begin{pmatrix} 1 & 0 & \beta_{e_1, e_3} & \beta_{e_1, e_4} & \beta_{e_1, e_3} \beta_{e_3, e_5} \\ 0 & 1 & 0 & 0 & \beta_{e_2, e_5} \\ 0 & 0 & 1 & 0 & \beta_{e_3, e_5} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{pmatrix}$$

Observe that $G = (I - F)^{-1}$ is polynomial

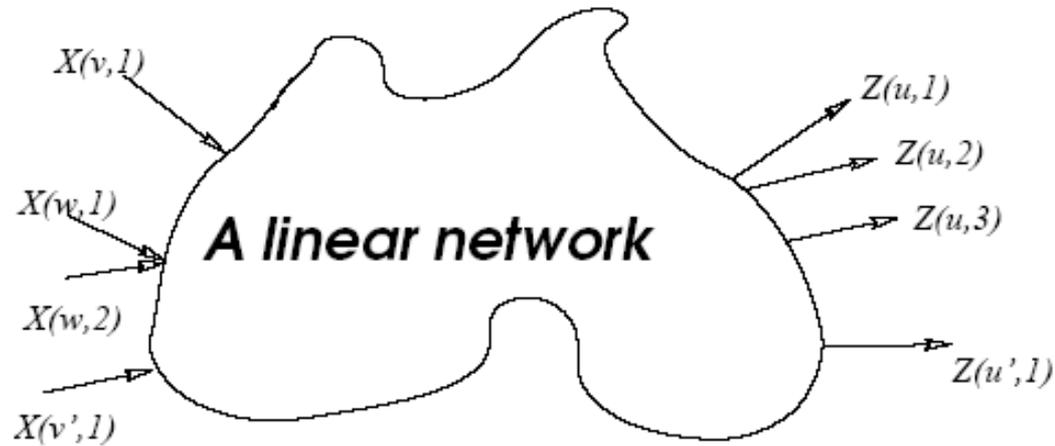
Transfer matrix



$$\begin{bmatrix} X_1 & X_2 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} 1 & 0 & \beta_{e_1 e_3} & \beta_{e_1 e_4} & \beta_{e_1 e_3} \beta_{e_3 e_5} \\ 0 & 1 & 0 & 0 & \beta_{e_2 e_5} \\ 0 & 0 & 1 & 0 & \beta_{e_3 e_5} \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} = \begin{bmatrix} Z_1 & Z_2 \end{bmatrix}$$

$$\begin{bmatrix} X_1 & X_2 \end{bmatrix} \begin{bmatrix} \beta_{e_1 e_4} & \beta_{e_1 e_3} \beta_{e_3 e_5} \\ 0 & \beta_{e_2 e_5} \end{bmatrix} = \begin{bmatrix} Z_1 & Z_2 \end{bmatrix}$$

Linear network system



Input vector: $\underline{x}^T = (X(v, 1), X(v, 2), \dots, X(v', \mu(v')))$

Output vector: $\underline{z}^T = (Z(u, 1), Z(u, 2), \dots, Z(u', \nu(u')))$

Transfer matrix: $M, \underline{z} = M\underline{x}$

$$\underline{z} = M\underline{x} = B^T \cdot \underbrace{(I - F^T)^{-1}}_{G^T} \cdot A^T \underline{x}$$

$$\underline{\xi} = (\xi_1, \xi_2, \dots) = (\dots, \alpha_{e,l}, \dots, \beta_{e',e}, \dots, \varepsilon_{e',j}, \dots)$$

For acyclic networks the elements of G (and hence M) are polynomial functions in **variables** $\underline{\xi} = (\xi_1, \xi_2, \dots)$

\Rightarrow an algebraic characterization of flows....

An algebraic max flow min cut condition

[KM01, 02, 03]

Let network be given with a source v and a sink v' . The following three statements are equivalent:

1. A point-to-point connection $c = (v, v', \mathcal{X}(v, v'))$ is possible.
2. The Min-Cut Max-Flow bound is satisfied for a rate $R(c) = |\mathcal{X}(v, v')|$.
3. The determinant of the $R(c) \times R(c)$ transfer matrix M is nonzero over the ring of polynomials $\mathbb{F}_2[\xi]$

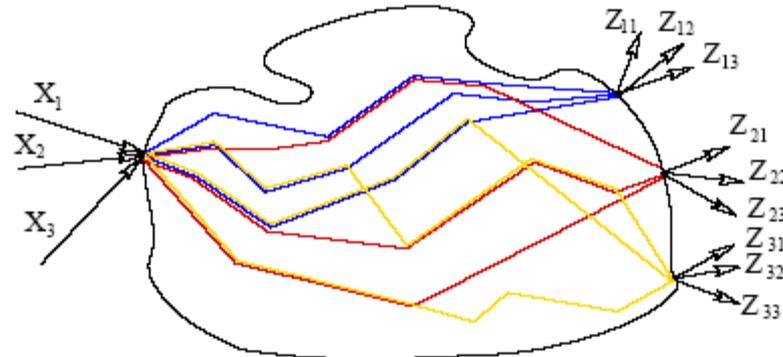
\Rightarrow We have to study the solution sets of polynomial equations.

Solutions

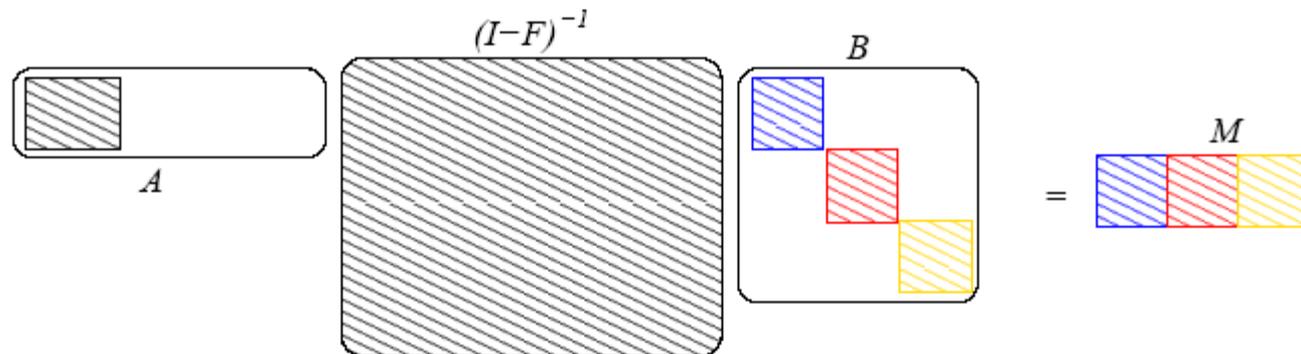
Let $\mathbb{F}[X_1, X_2, \dots, X_n]$ be the ring of polynomials over an infinite field \mathbb{F} in variables X_1, X_2, \dots, X_n . For any non-zero element $f \in \mathbb{F}[X_1, X_2, \dots, X_n]$ there exists an infinite set of n -tuples $(x_1, x_2, \dots, x_n) \in \mathbb{F}^n$ such that $f(x_1, x_2, \dots, x_n) \neq 0$.

Multicast

$$\mathcal{C} = \{(v, u_1, \mathcal{X}(v)), (v, u_2, \mathcal{X}(v)), \dots, (v, u_K, \mathcal{X}(v))\}$$



Multicast network



M is a $|\mathcal{X}(v)| \times K|\mathcal{X}(v)|$ matrix.

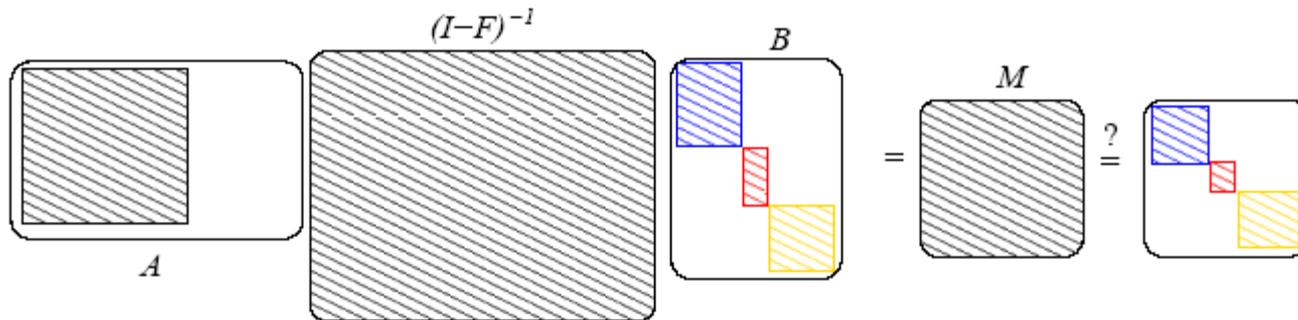
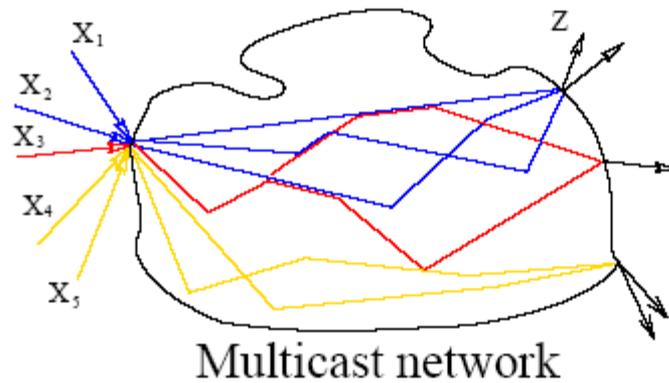
Multicast

[ACLY 00, KM01, 02, 03]

Theorem Let $(\mathcal{G}, \mathcal{C})$ be a multicast network coding problem. There exists a linear network coding solution for $(\mathcal{G}, \mathcal{C})$ over a finite field \mathbb{F}_{2^m} for some large enough m if and only if there exists a flow of sufficient capacity between the source and each sink **individually**.

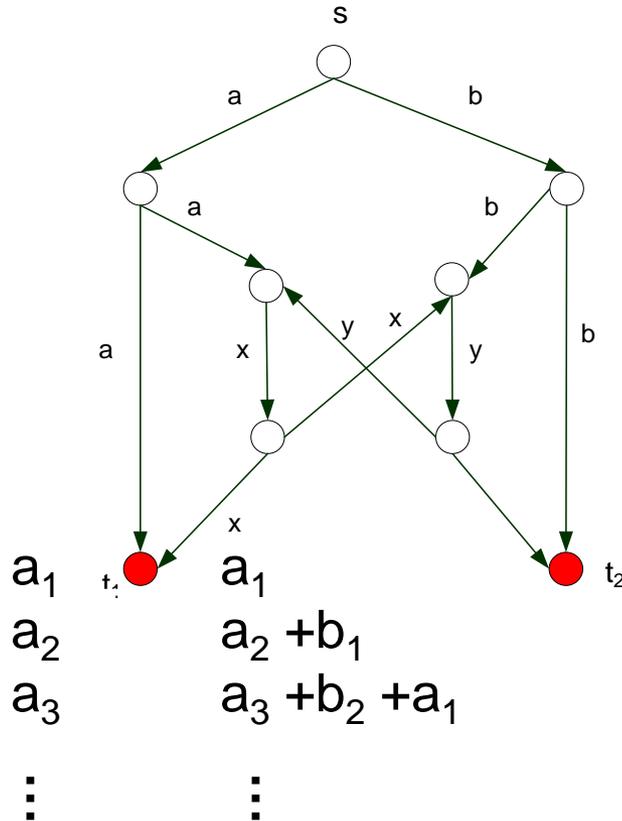
One source, disjoint multicasts

$$\mathcal{C} = \{(v, u_j, \mathcal{X}(v, u_j)) : j = 1, 2, \dots, K\}, \mathcal{X}(v, u_j) \cap \mathcal{X}(v, u_i) = \emptyset$$



Networks with cycles

- Mix messages from different rounds



$$x_i = \begin{cases} a_1 & \text{if } i = 1 \\ a_i \oplus y_{i-1} & \text{otherwise} \end{cases}$$

$$y_i = \begin{cases} b_1 & \text{if } i = 1 \\ b_i \oplus x_{i-1} & \text{otherwise} \end{cases}$$

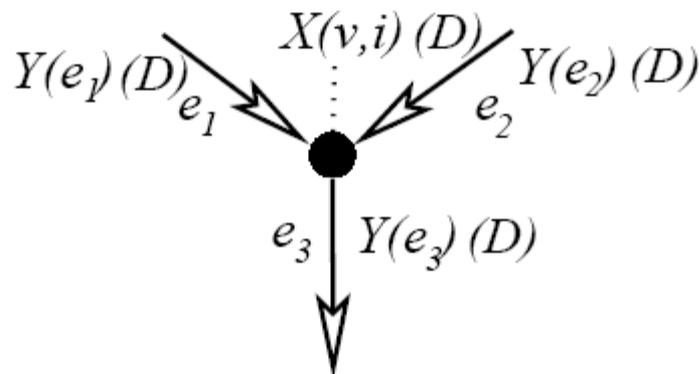
Delays

We transmit random processes in a delay variable D on links, i.e.

$$\begin{aligned}X(v, j)(D) &= \sum_{\ell=0}^{\infty} X_{\ell}(v, j) D^{\ell}, \\Z(v, j)(D) &= \sum_{\ell=0}^{\infty} Z_{\ell}(v, j) D^{\ell}, \\Y(e)(D) &= \sum_{\ell=0}^{\infty} Y_{\ell}(e) D^{\ell}.\end{aligned}$$

Conceptually, we consider an entire sequence in D as one symbol and work over the field of formal power series.

Delays

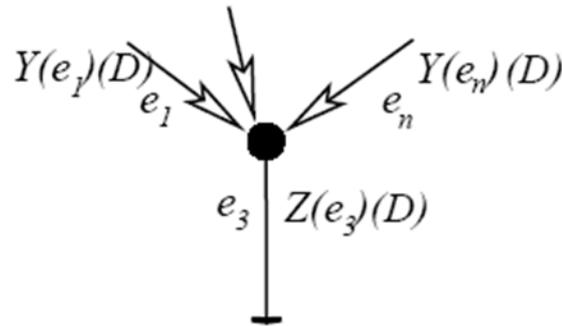


$$Y(e_3)(D) = \sum_i \alpha_i D X(v,i)(D) + \sum_{j=1,2} \beta_j D Y(e_j)(D)$$

(other functions with memory are possible but not necessary)

Delays

At a receiver (terminal) node we have to allow for "rational" functions:



$$Y(e)(D) = \sum_{\ell=0}^{\infty} Y_{\ell}(e)D^{\ell}, \quad Z(v, j)(D) = \sum_{\ell=0}^{\infty} Z_{\ell}(v, j)D^{\ell}$$

$$Z_{\ell}(v, j) = \sum_{j=1}^n \sum_{k=0}^{\mu} \varepsilon_{j,k} Y_{\ell-k}(e_j) + \sum_{k=1}^{\mu} \lambda_k Z_{\ell-k}(v, j)$$

or

$$Z(v, j)(D) = \sum_{j=1}^n \frac{\varepsilon_{j,k}(D)}{\lambda(D)} Y(e_j)(D)$$

Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Centralized multicast code construction

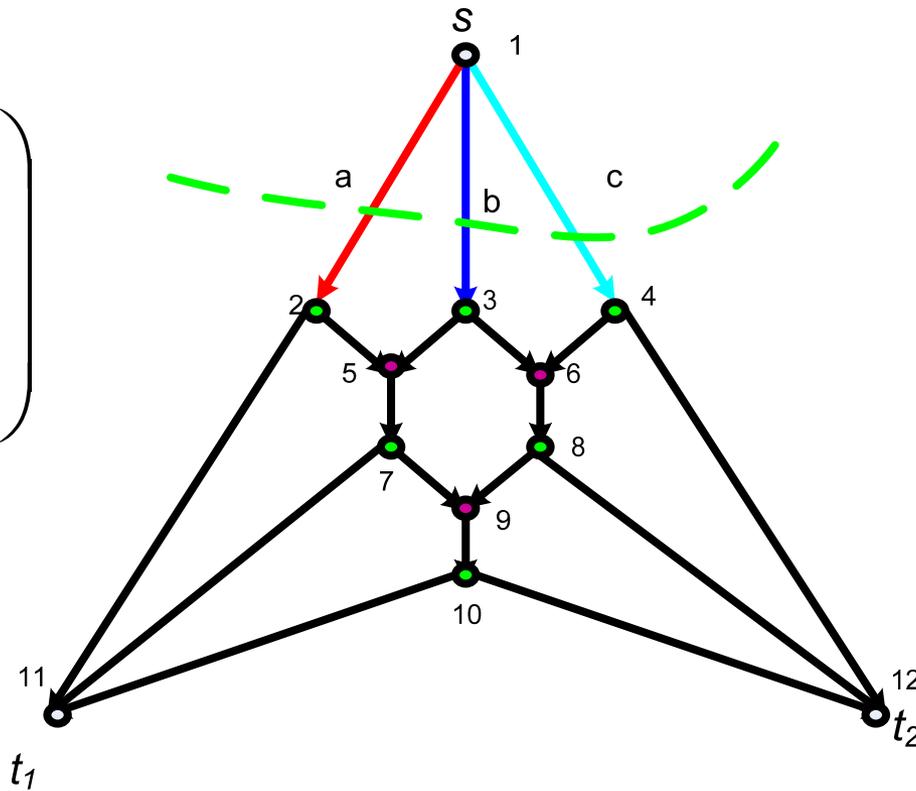
- Centralized polynomial-time construction for acyclic graphs
 - Choose a flow solution for each sink individually
 - Consider the links in the union of these flow solutions
 - Set the code coefficients of these links in ancestral order starting from the source, ensuring that at each step the "frontier set" for each sink has linearly independent coefficient vectors

S. Jaggi, P. Sanders, P. A. Chou, M. Effros, S. Egnér, K. Jain, and L. M. Tolhuizen, "Polynomial time algorithms for multicast network code construction," *IEEE Trans. Inform. Theory*, 51(6):1973–1982, 2005

Example

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

B_{t_1}



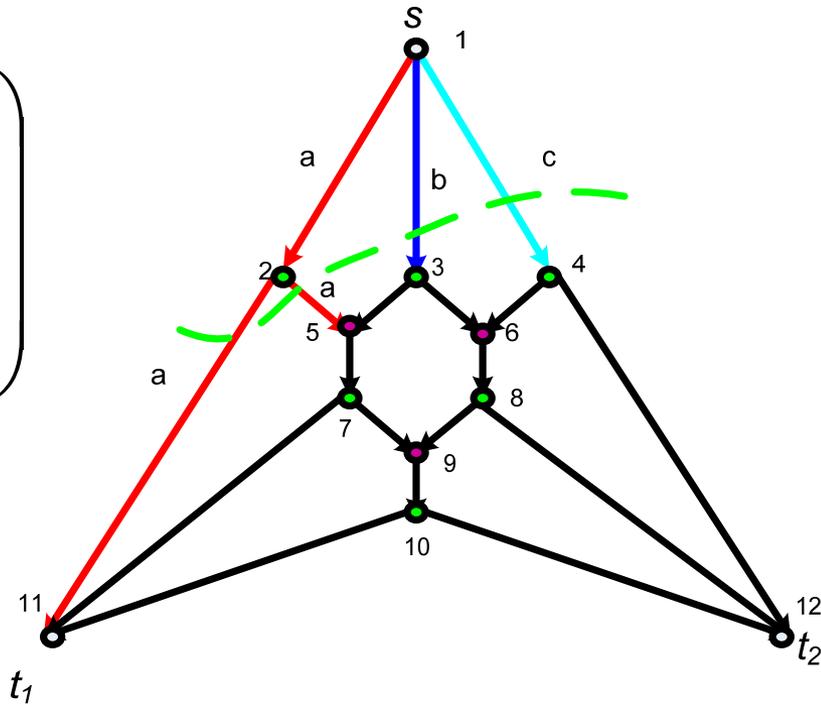
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

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

B_{t_1}



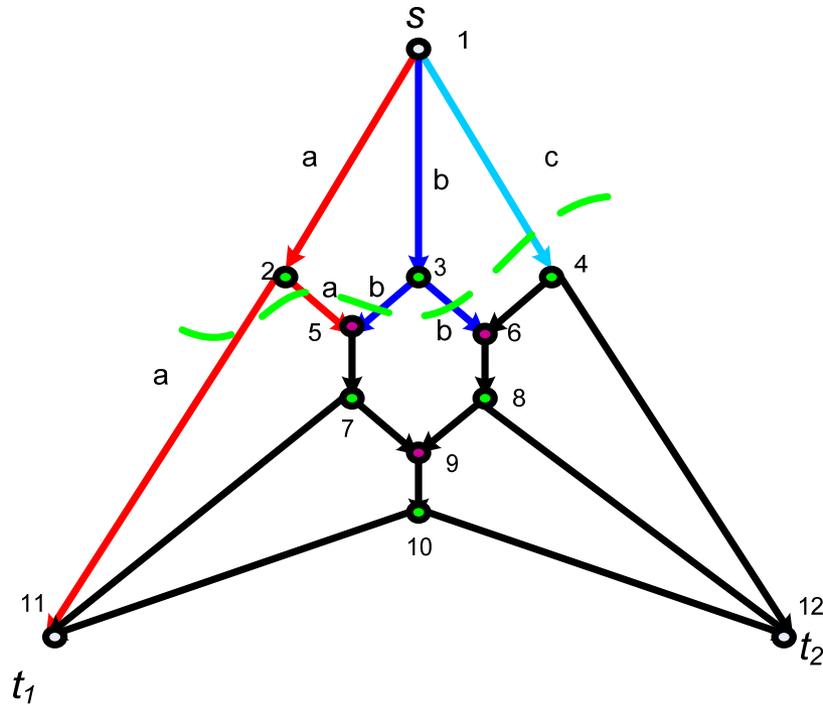
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

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

B_{t_1}



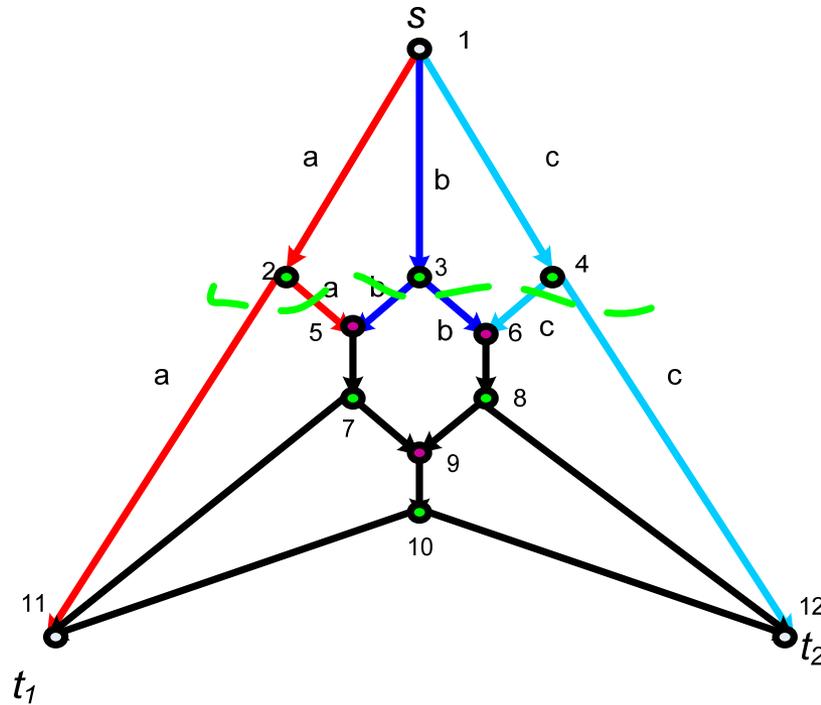
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

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

B_{t_1}



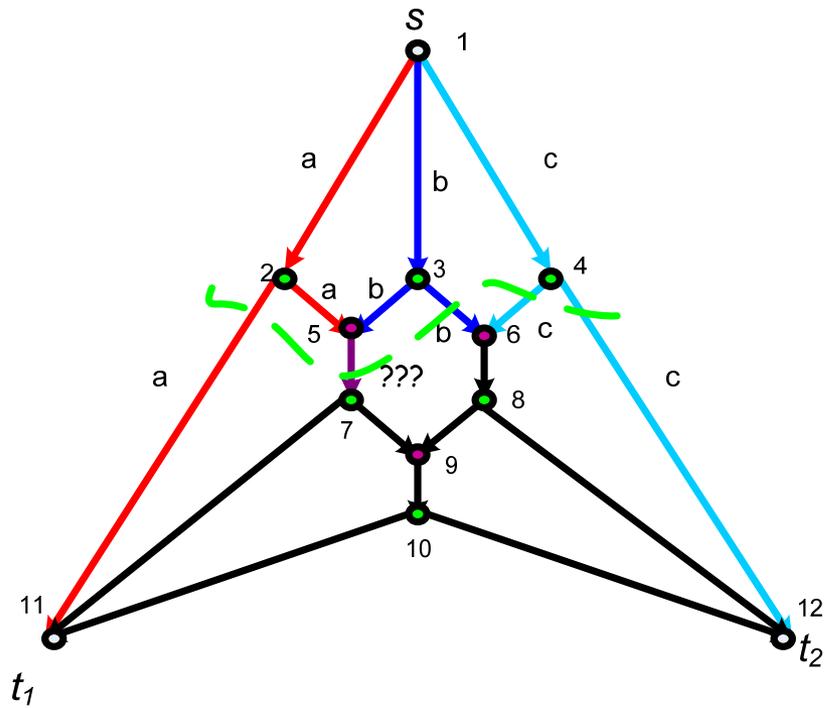
$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

$$\begin{pmatrix} 1 & ? & 0 \\ 0 & ? & 0 \\ 0 & ? & 1 \end{pmatrix}$$

B_{t_1}



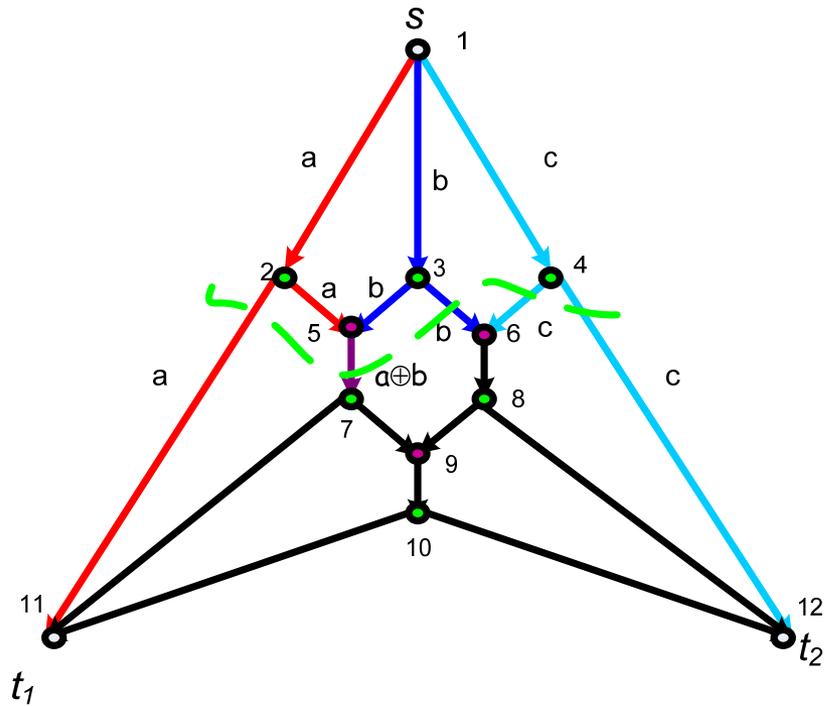
$$\begin{pmatrix} ? & 0 & 0 \\ ? & 1 & 0 \\ ? & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

$$\begin{pmatrix} 1 & \boxed{1} & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_1}



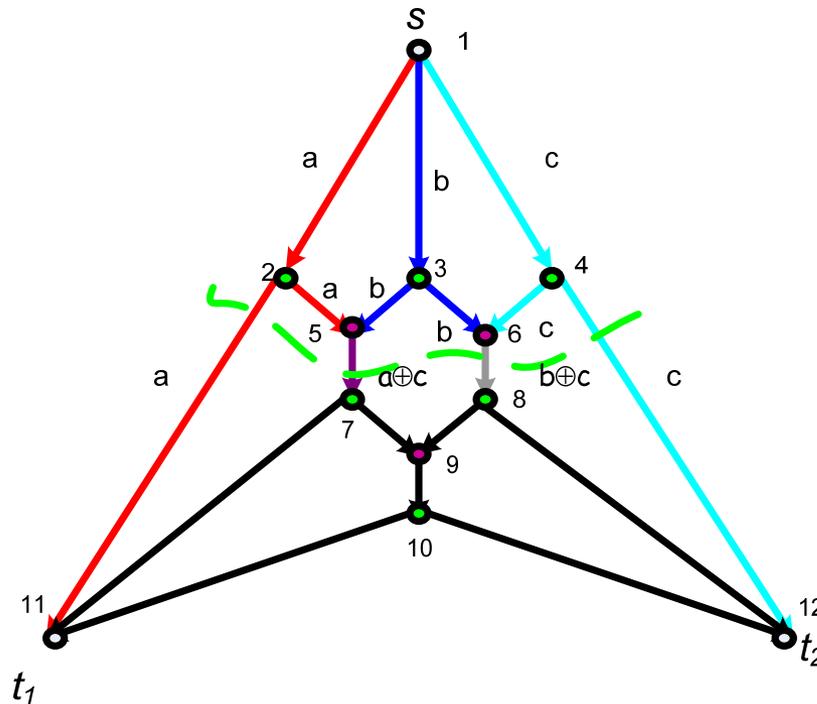
$$\begin{pmatrix} \boxed{1} & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

B_{t_2}

Example

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

B_{t_1}

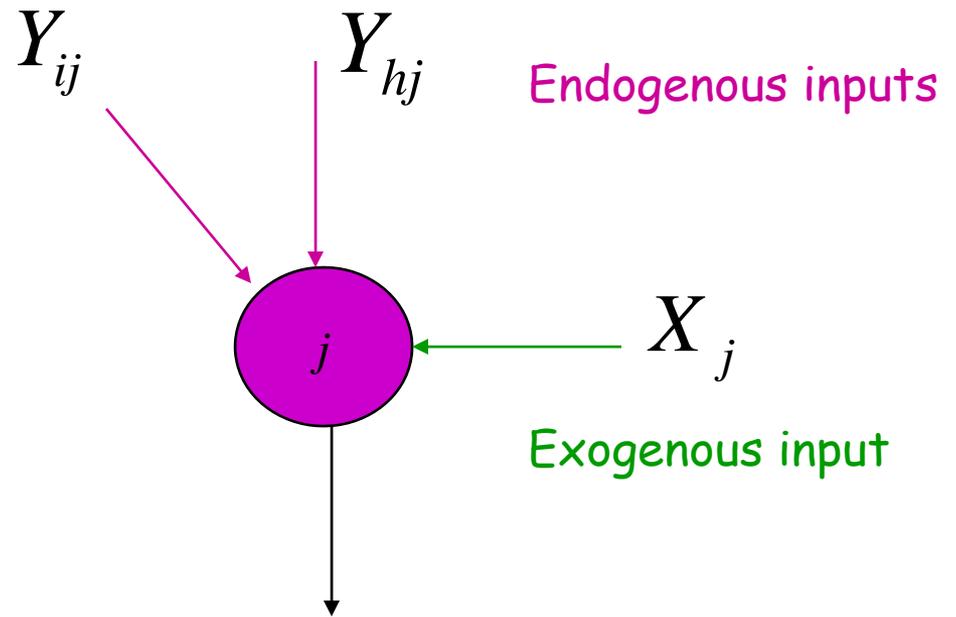


$$\begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix}$$

t2 matrix

Random multicast code construction

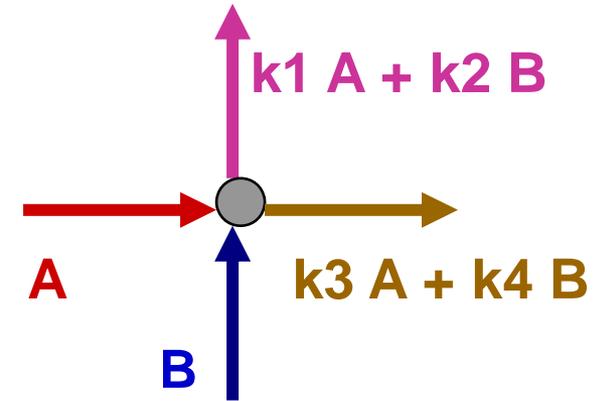
- The effect of the network code is that of a transfer matrix from source inputs to receiver outputs
- To recover source symbols, receivers need sufficient degrees of freedom - an invertible transfer matrix
- The realization of the determinant of the matrix will be non-zero with high probability if the code coefficients are chosen independently and randomly from a sufficiently large field



$$Y_{jm} = \alpha_{jm}^{ij} Y_{ij} + \alpha_{jm}^{hj} Y_{hj} + \alpha_{jm}^j X_j$$

Distributed random network coding

- Random linear coding among packets of a single multicast or unicast session
- Nodes independently choose random linear mappings from inputs to outputs in some field
- Header scheme to communicate transfer matrix to receivers: vector of code coefficients in packet headers, to which same linear mappings are applied



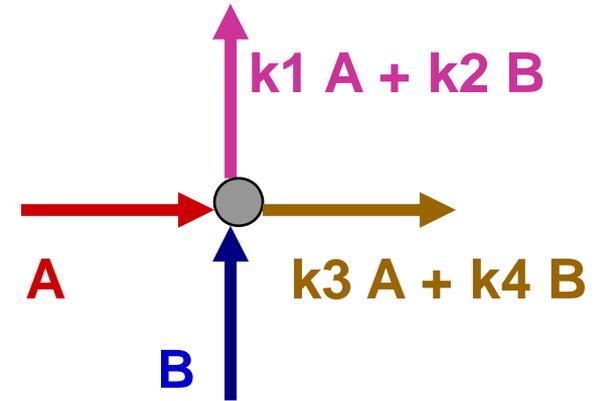
1	0	A
0	1	B

k1	k2	k1 A + k2 B
k3	k4	k3 A + k4 B

T. Ho, R. Koetter, M. Médard, D. R. Karger and M. Effros, "The Benefits of Coding over Routing in a Randomized Setting", International Symposium on Information Theory 2003.

Distributed random network coding

- For any multicast subgraph which satisfies min-cut max-flow bound for each receiver, probability of failure over field F is roughly inversely Polynomials in field size
- Random network coding can be used to distribute a group of packets from any number of sources
- Advantages: decentralized, optimal throughput, robust to link failures / packet losses



1	0	A
0	1	B

k_1	k_2	$k_1 A + k_2 B$
k_3	k_4	$k_3 A + k_4 B$

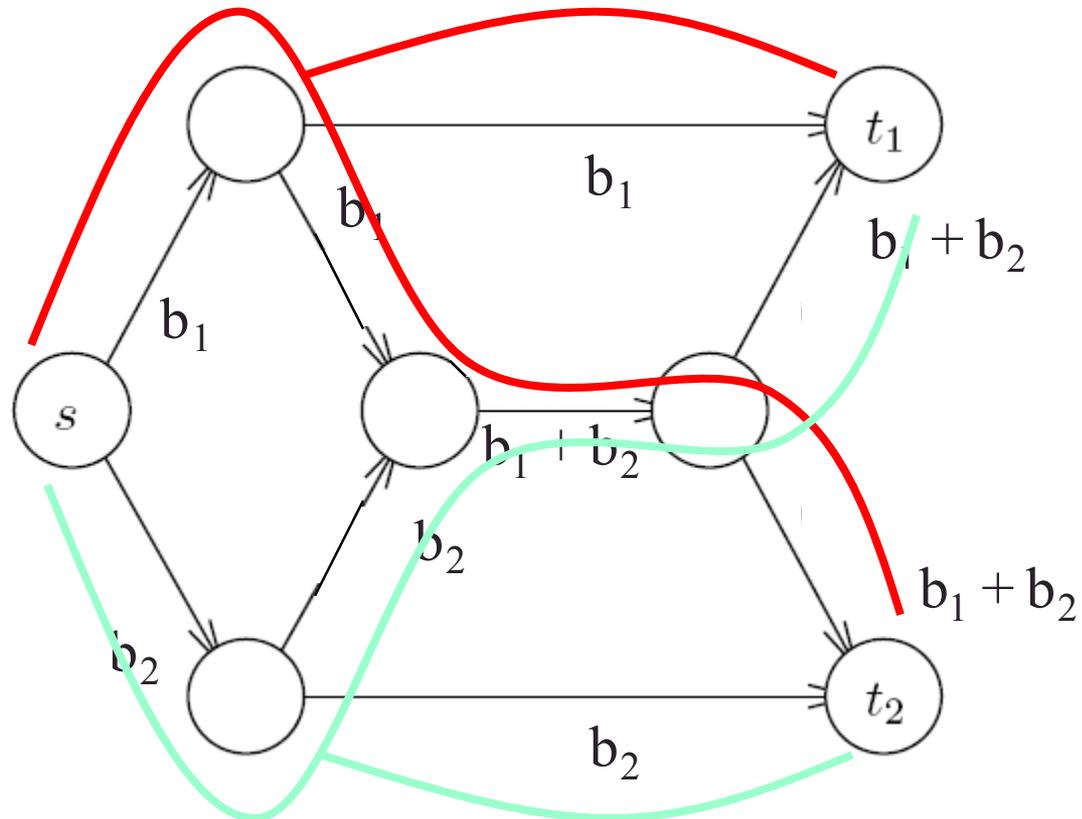
T. Ho, R. Koetter, M. Médard, D. R. Karger and M. Effros, "The Benefits of Coding over Routing in a Randomized Setting", International Symposium on Information Theory 2003.

Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Minimum cost multicast optimization

- Without network coding, minimum cost multicast optimization problem is NP-complete
- E.g., in the illustration, integrality constraint arises in time-sharing between the blue and red trees



Minimum cost multicast optimization with network coding

- With network coding, minimum cost multicast optimization problem becomes a polynomial-complexity linear optimization

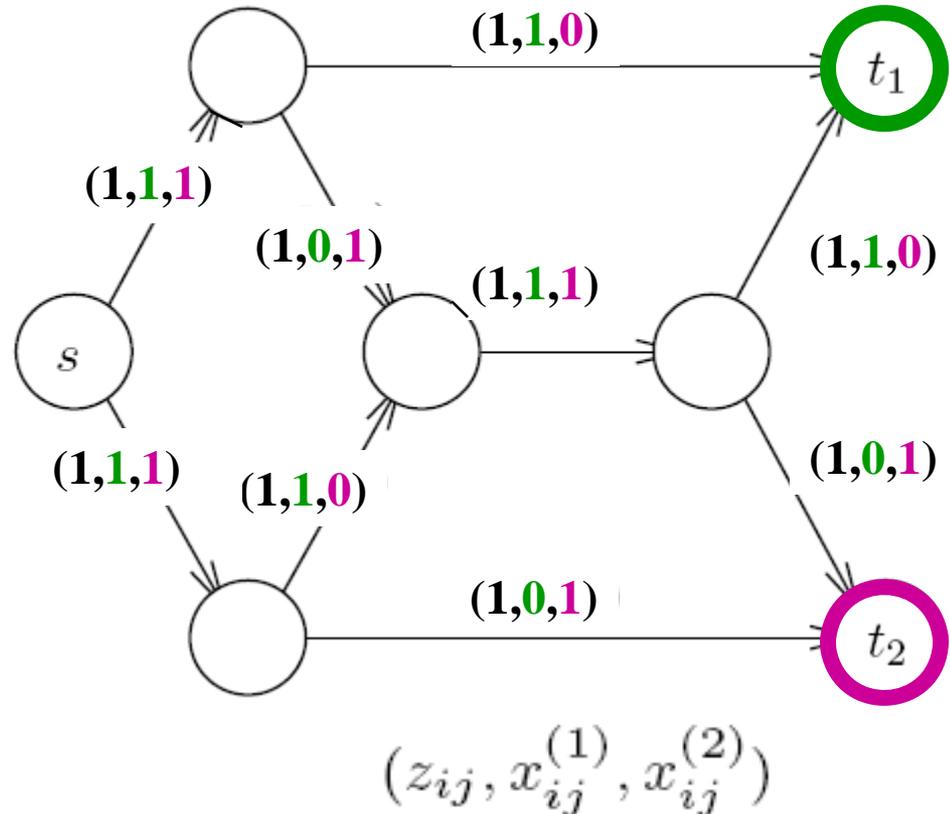
minimize $f(z)$

subject to $z \in Z$

$z_{ij} \geq x_{ij}^{(t)} \geq 0,$

$$\sum_{\{j|(i,j) \in \mathcal{A}\}} x_{ij}^{(t)} - \sum_{\{j|(j,i) \in \mathcal{A}\}} x_{ji}^{(t)}$$

$$= \begin{cases} R, & \text{if } i = s \\ -R, & \text{if } i = t \\ 0 & \text{otherwise.} \end{cases}$$



D. S. Lun, N. Ratnakar, M. Médard, R. Koetter, D. R. Karger, T. Ho, E. Ahmed, and F. Zhao. Minimum-cost multicast over coded packet networks. *IEEE Trans. Inform. Theory*, 52(6):2608-2623, June 2006.

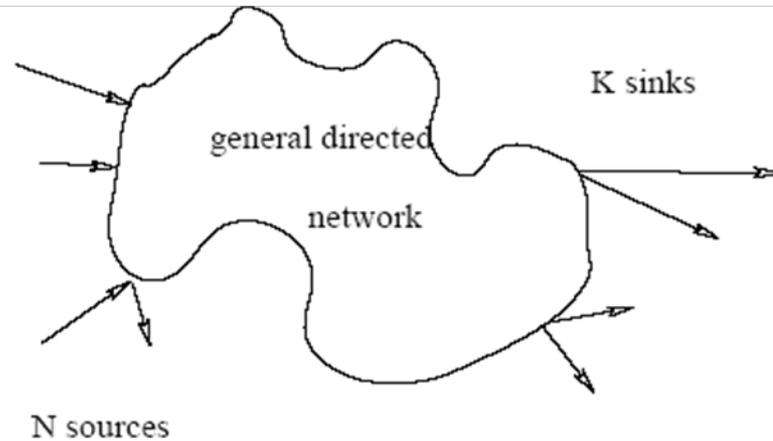
Multicast optimization with network coding

- The minimum cost multicast optimization problem can be solved in polynomial-time in a distributed way by a subgradient algorithm [Lun et al 05]
- For multiple multicast sessions with intra-session network coding, a distributed back pressure approach can be used [Ho & Viswanathan 05]
 - Generalization of back pressure approach for multi-commodity routing of Tassiulas & Ephremides 92, Awerbuch & Leighton 93
 - Network coding greatly reduces complexity for multicast

Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Network coding for non-multicast



$$\mathcal{C} = \{(v_i, u_j, \mathcal{X}(v_i, u_j))\}$$

$$M = \begin{pmatrix} M_{1,1} & M_{1,2} & \dots & M_{1,K} \\ M_{2,1} & M_{2,2} & & M_{2,K} \\ \vdots & & & \vdots \\ M_{N,1} & M_{N,2} & \dots & M_{N,K} \end{pmatrix}$$

$M_{i,j}$ corresponds to $c_{i,j} = (v_i, u_j, \mathcal{X}(v_i, u_j))$.

Theorem [Generalized Min-Cut Max-Flow Condition] Let an acyclic, delay-free scalar linear network problem $(\mathcal{G}, \mathcal{C})$ be given and let $M = \{M_{i,j}\}$ be the corresponding transfer matrix relating the set of input nodes to the set of output nodes. The network problem is solvable if and only if there exists an assignment of numbers to $\underline{\xi}$ such that

1. $M_{i,j} = 0$ for all pairs (v_i, v_j) of vertices such that $(v_i, v_j, \mathcal{X}(v_i, v_j)) \notin \mathcal{C}$.

2. If \mathcal{C} contains the connections

$(v_{i_1}, v_j, \mathcal{X}(v_{i_1}, v_j)), (v_{i_2}, v_j, \mathcal{X}(v_{i_2}, v_j)), \dots, (v_{i_\ell}, v_j, \mathcal{X}(v_{i_\ell}, v_j))$
the determinant of $[M_{i_1,j}^T, M_{i_2,j}^T, \dots, M_{i_\ell,j}^T]$ is nonzero.

For the general case we need to find **solutions** to some system of polynomial equations!

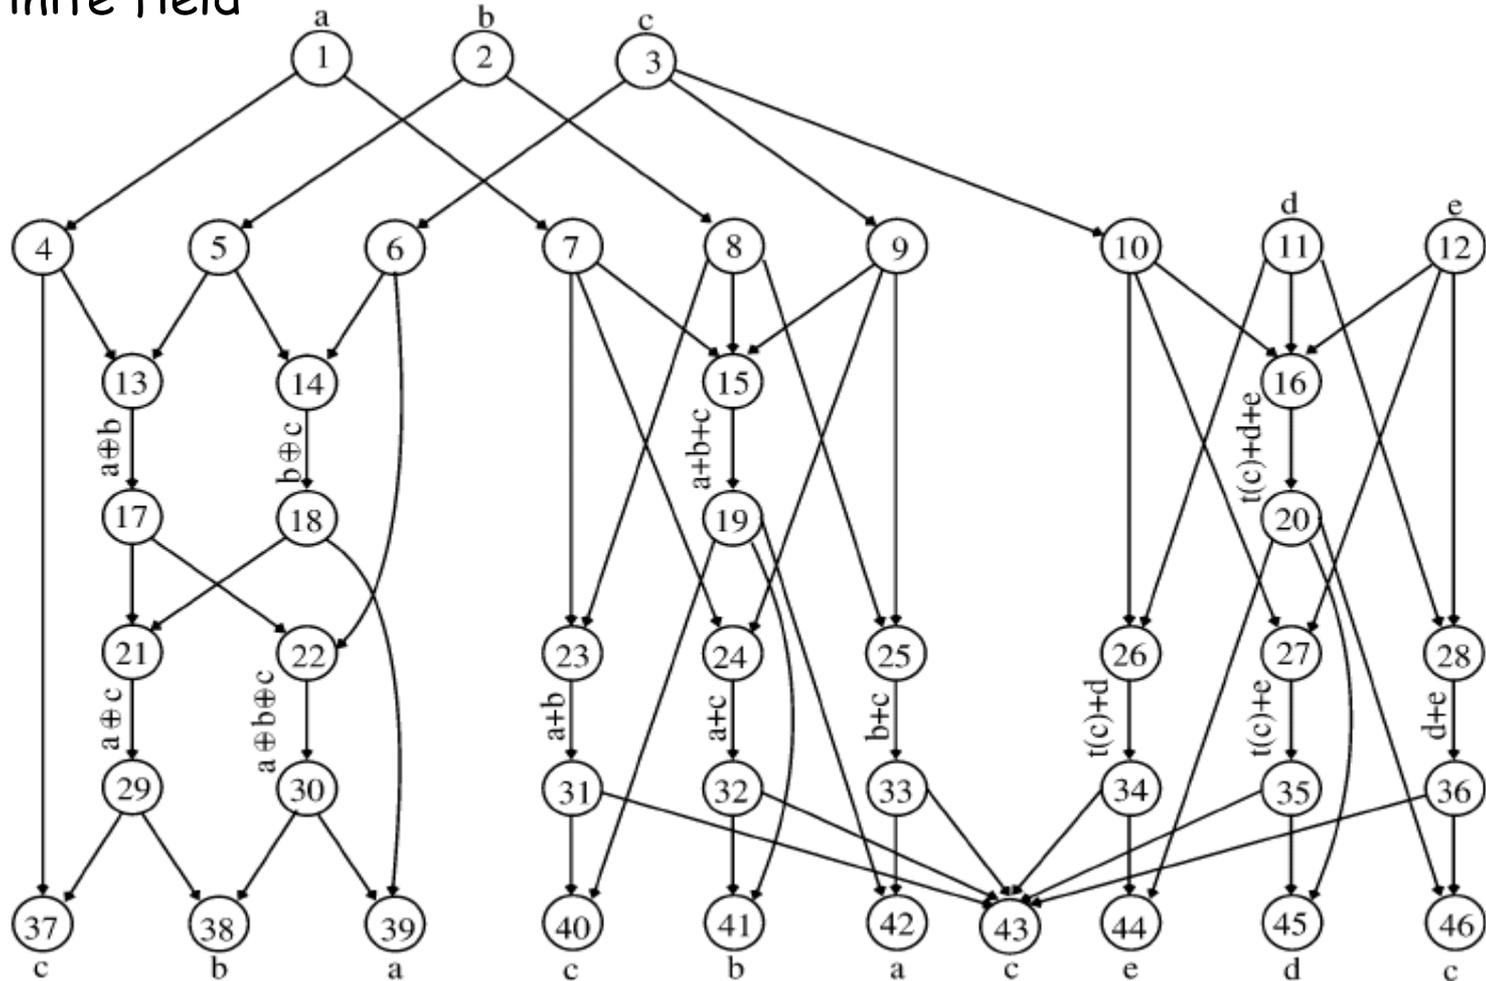
For the multicast case we need to find **non solutions** to some system of polynomial equations!

Another way to phrase this is: In a multicast setup everybody wants everything so the issue of interference is moot!

For the general case we may have carefully balanced solutions where some unwanted information cancels out in clever ways.....

Linear coding may not suffice

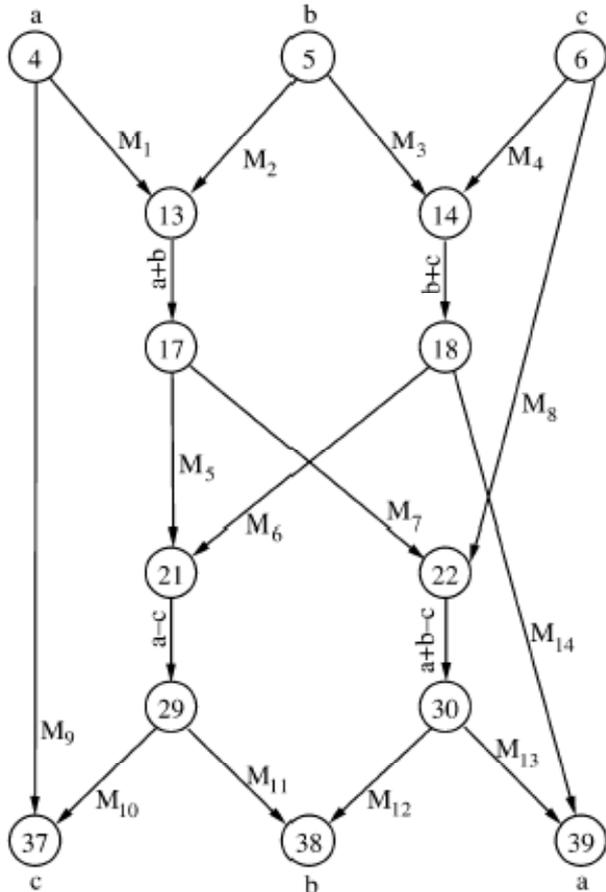
Network with no linear solution for any vector dimension over any finite field



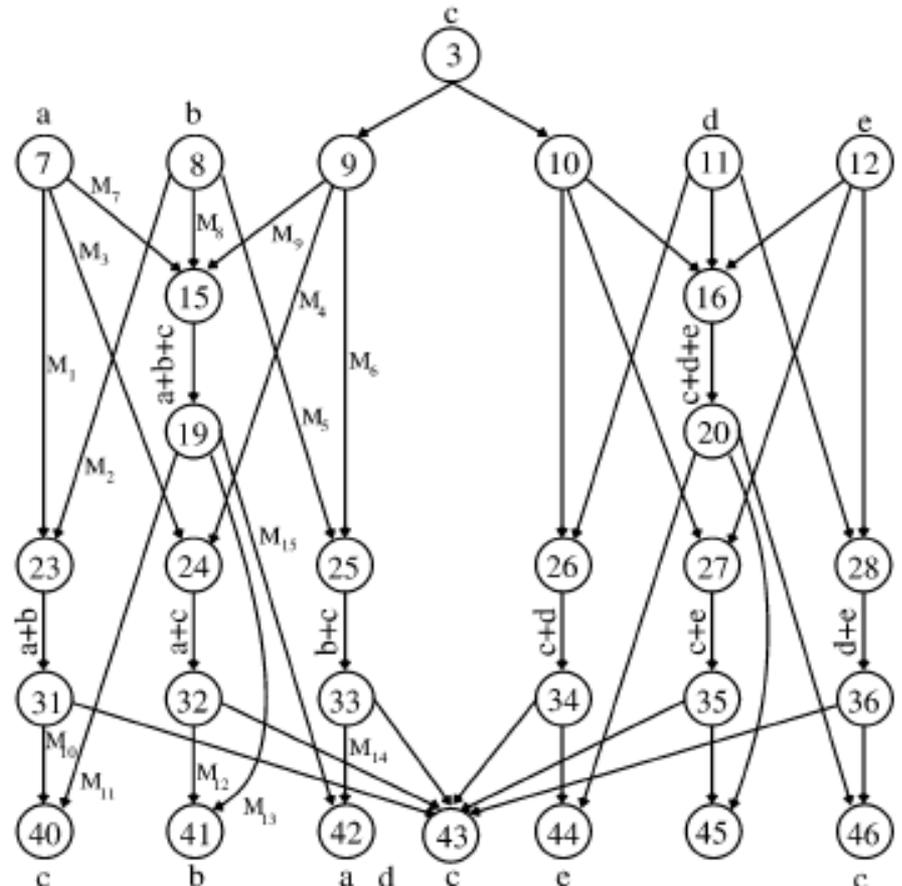
Dougherty, R. Freiling, C. Zeger, K., "Insufficiency of linear coding in network information flow", IEEE Transactions on Information Theory, Aug. 2005

Linear coding may not suffice

No linear solution for any vector dimension over a finite field with odd characteristic



No linear solution for any vector dimension over a finite field with characteristic 2

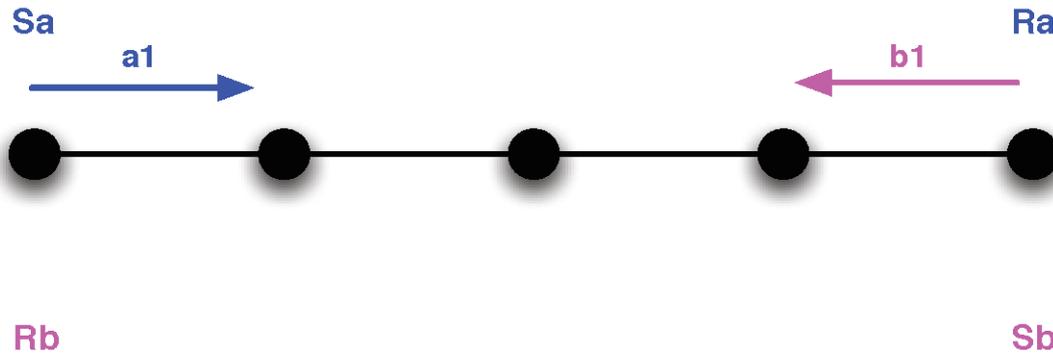


Questions

- When is coding advantageous in terms of throughput or cost and by how much?
- What types of codes are needed?
- How do we construct such codes?

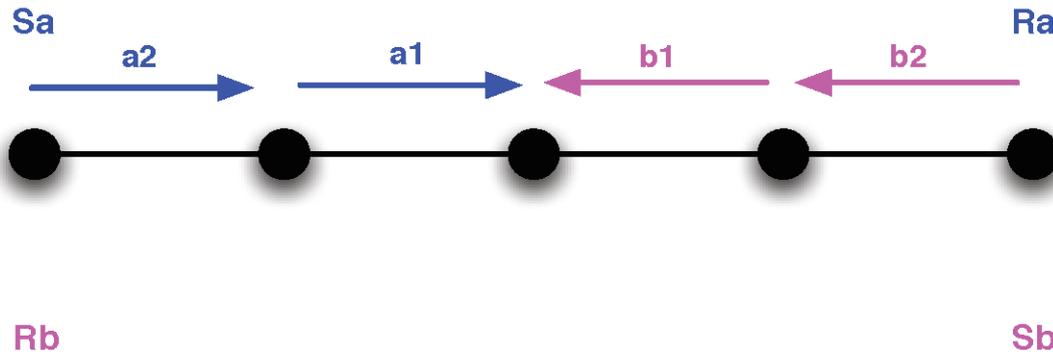
Coding advantages

- information exchange between two nodes



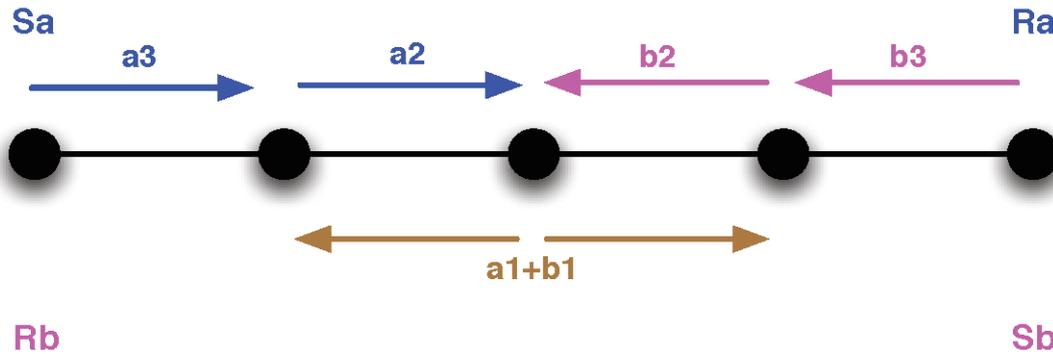
Coding advantages

- information exchange between two nodes



Coding advantages

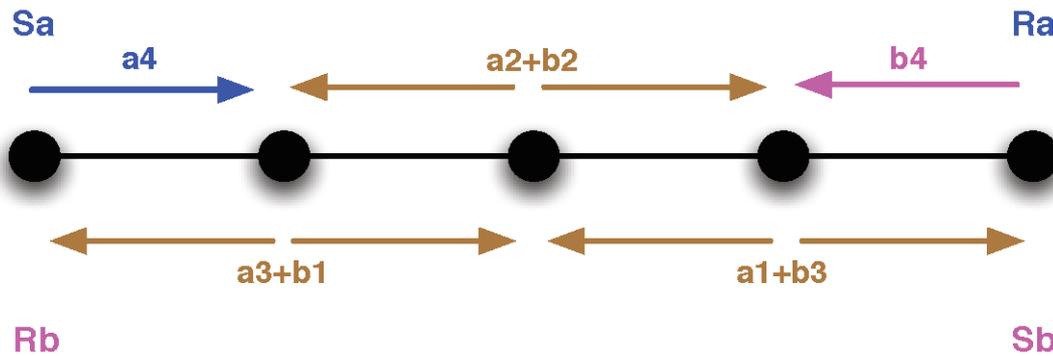
- information exchange between two nodes



Y. Wu, P. A. Chou, S.-Y. Kung, "Information exchange in wireless networks with network coding and physical-layer broadcast," Microsoft Technical Report, MSR-TR-2004-78, Aug. 2004

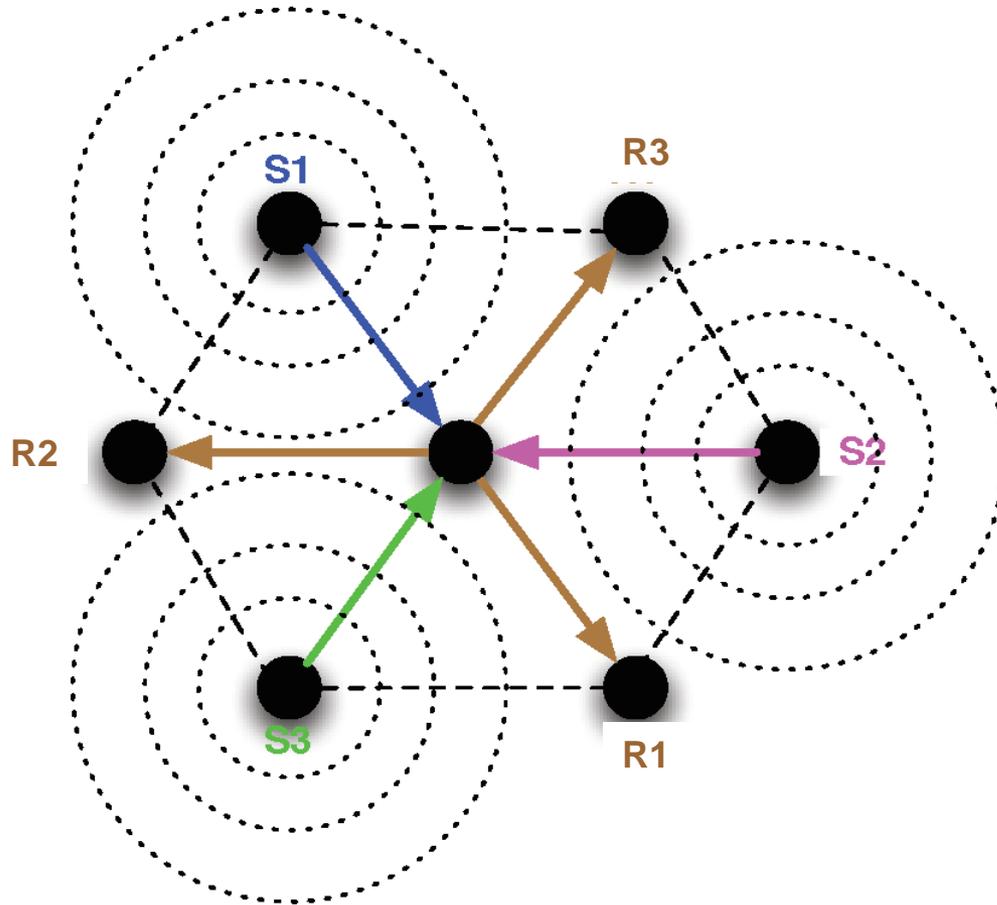
Coding advantages

- information exchange between two nodes



Coding advantages

- Intersection of paths



Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Opportunism (1)

Opportunistic Listening:

- Every node listens to all packets
- It stores all heard packets for a limited time
- Node sends **Reception Reports** to tell its neighbors what packets it heard
 - Reports are annotations to packets
 - If no packets to send, periodically send reports

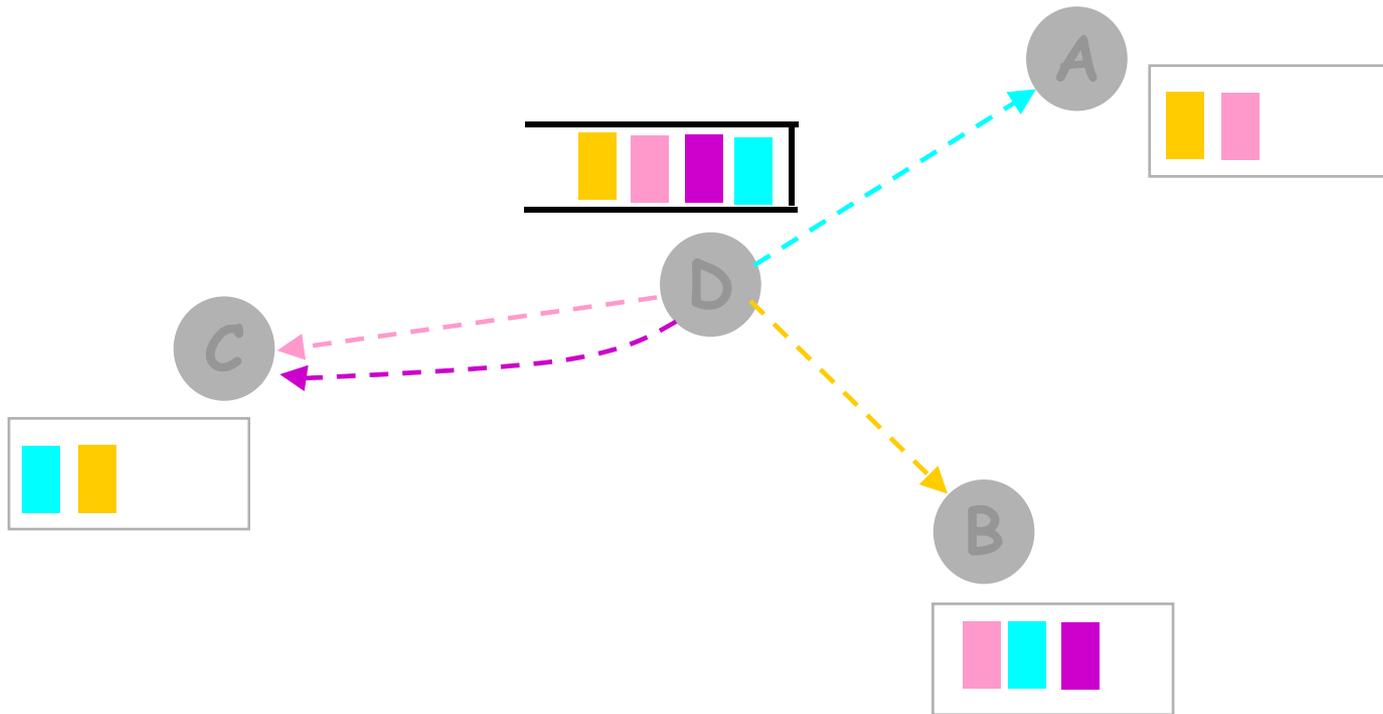
Sachin Katti, Hariharan Rahul, Wenjun Hu, Dina Katabi, Muriel Medard and Jon Crowcroft
“XORs in the Air: Practical Wireless Network Coding”, ACM SIGCOMM 2006.

Opportunism (2)

Opportunistic Coding:

- Works with any routing protocol
- To send packet p to neighbor A , can XOR p with packets already known to A
 - Thus, A can decode
- We would like to benefit multiple neighbors from a single transmission

Efficient coding

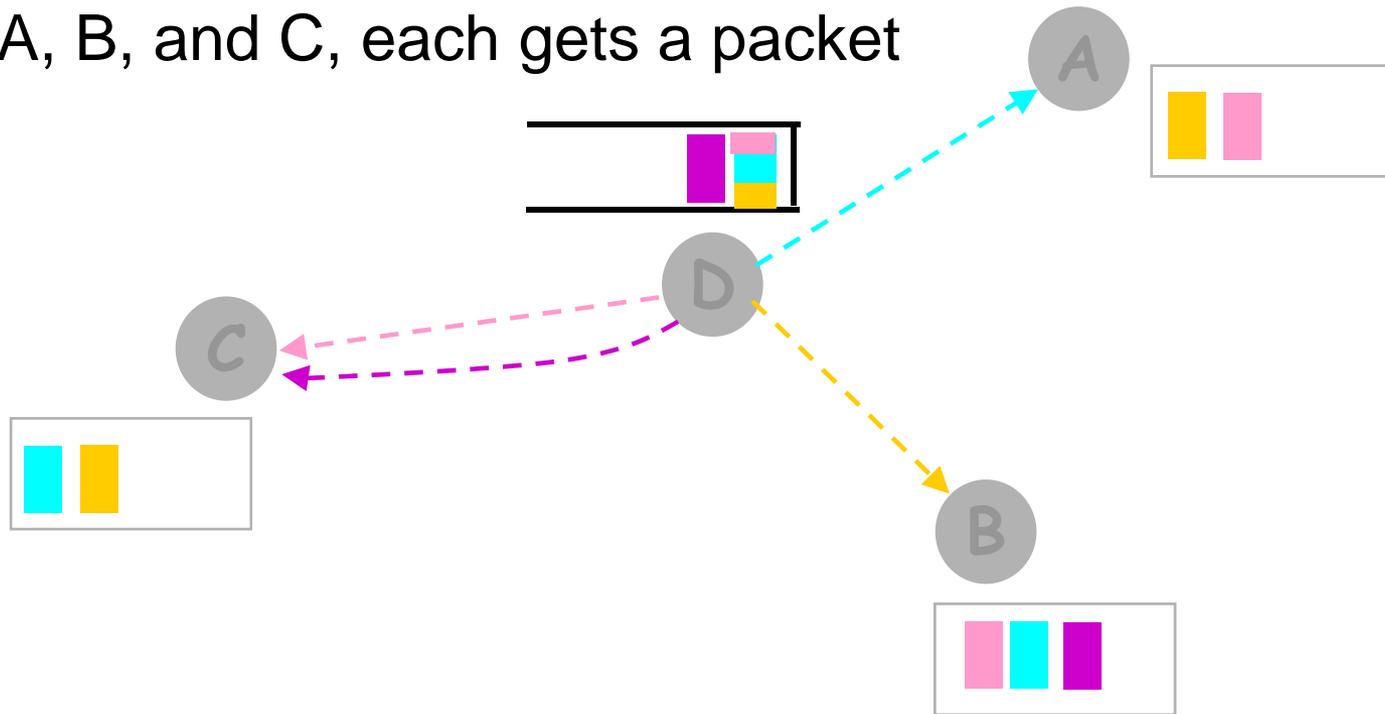


Arrows show next-hop

Efficient coding

Best Coding

A, B, and C, each gets a packet



To XOR n packets, each next-hop should have the $n-1$ packets encoded with the packet it wants

But how does a node know what packets a neighbor has?

- Reception Reports
- But reception reports may get lost or arrive too late
- Use *Guessing*
 - If I receive a packet I assume all nodes closer to sender have received it

Experiment

- Piggyback on 802.11 unicast which has collision detection and backoff
 - Each XOR-ed packet is sent to the *MAC* address of one of the intended receivers
 - Put all cards in promiscuous mode
- 40 nodes
- 400m x 400m
- Senders and receivers are chosen randomly
- Flows are duplex (e.g., ping)
- Metric:
 - Total Throughput of the Network

Current 802.11

Net. Throughput (KB/s)

2500

2000

1500

1000

500

0

1

2

4

6

8

10

12

14

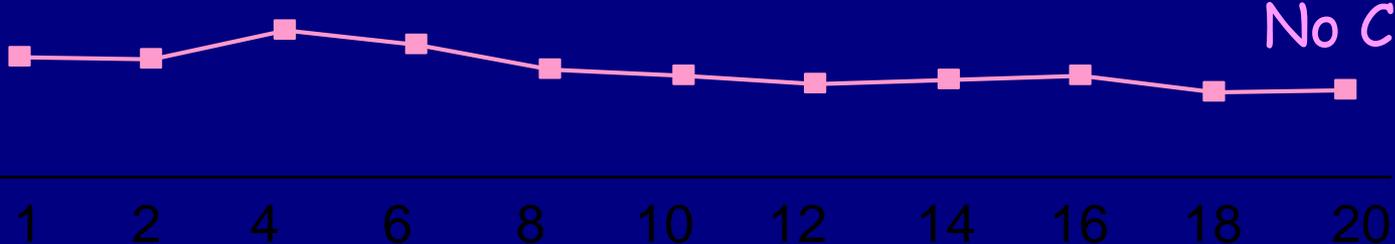
16

18

20

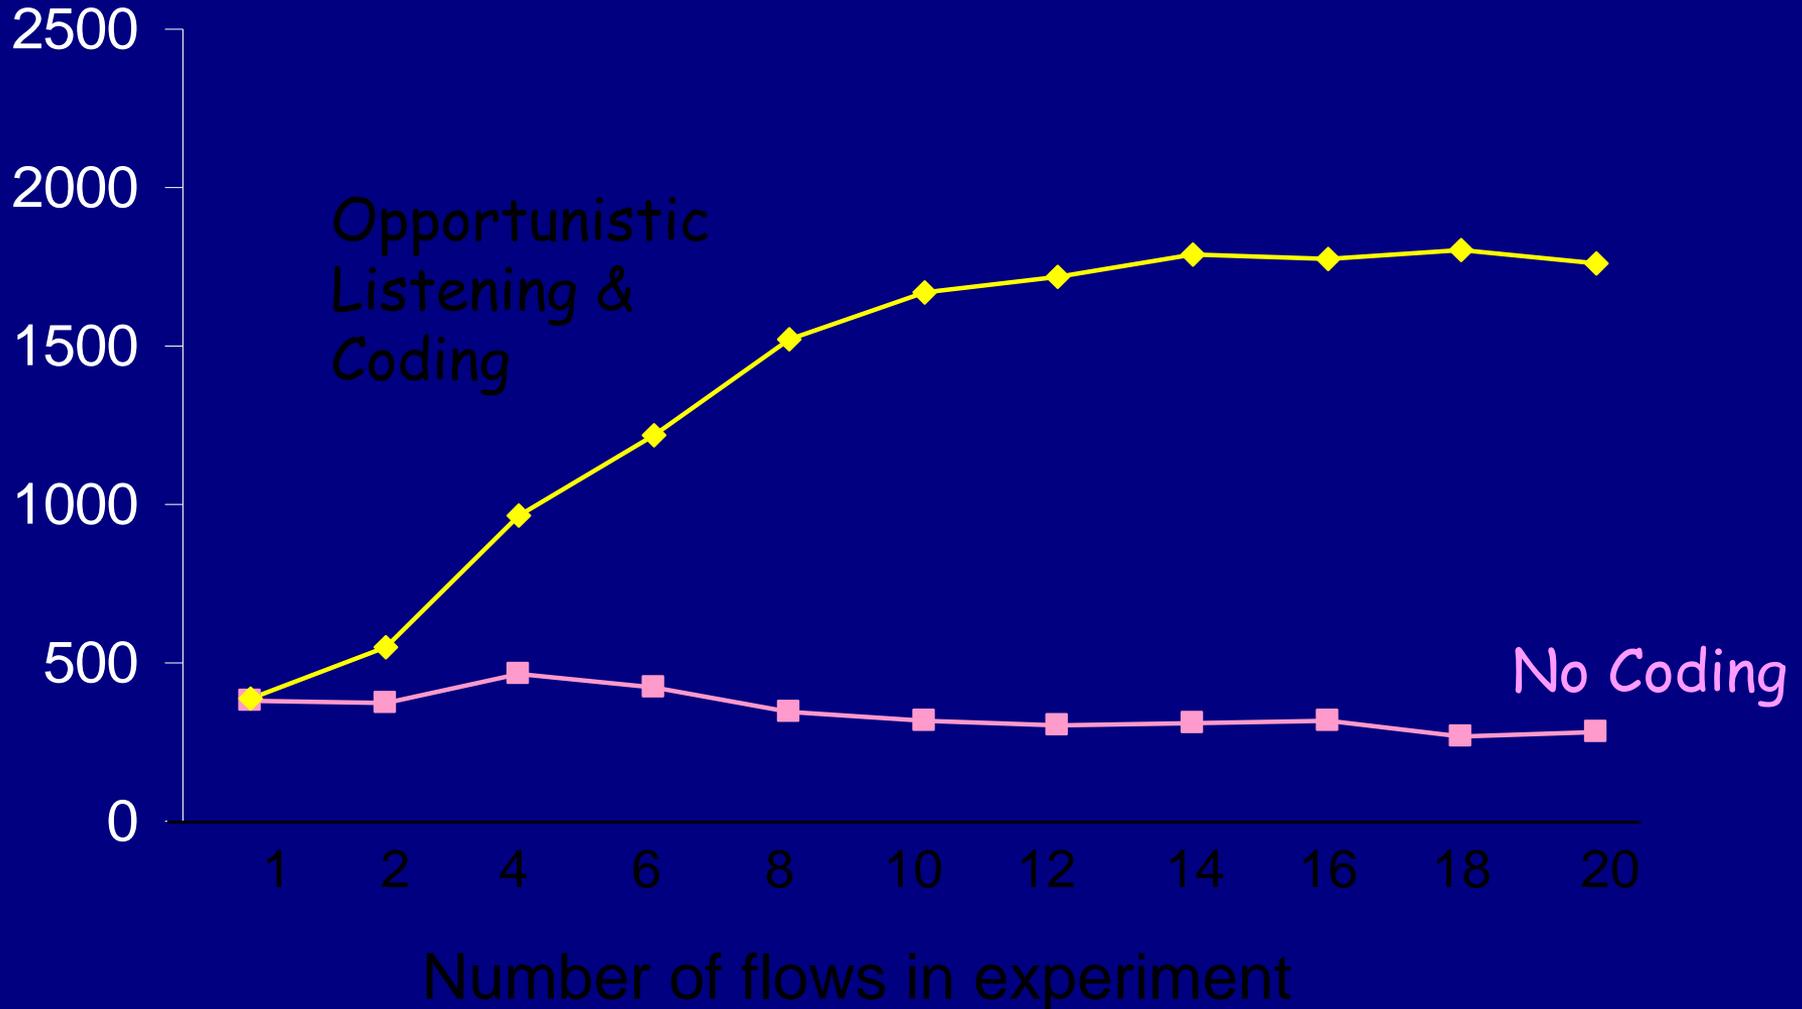
Number of flows in experiment

No Coding



Opportunistic Listening & Coding

Net. Throughput (KB/s)

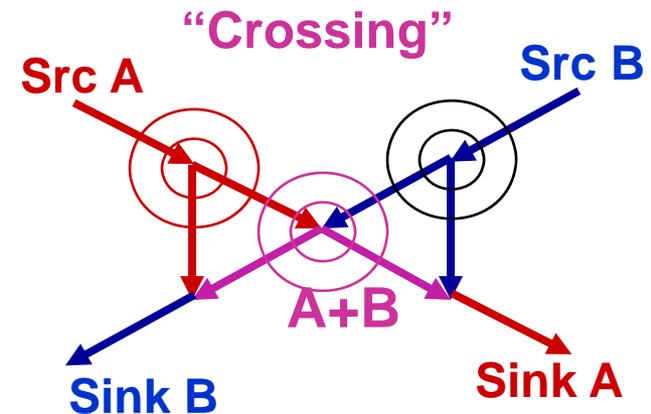
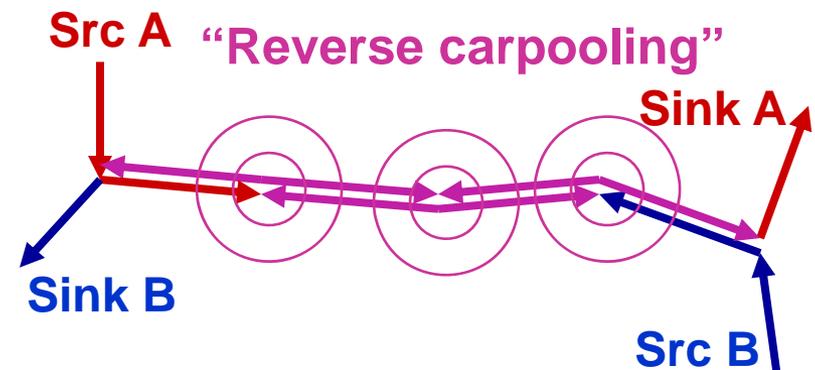


Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

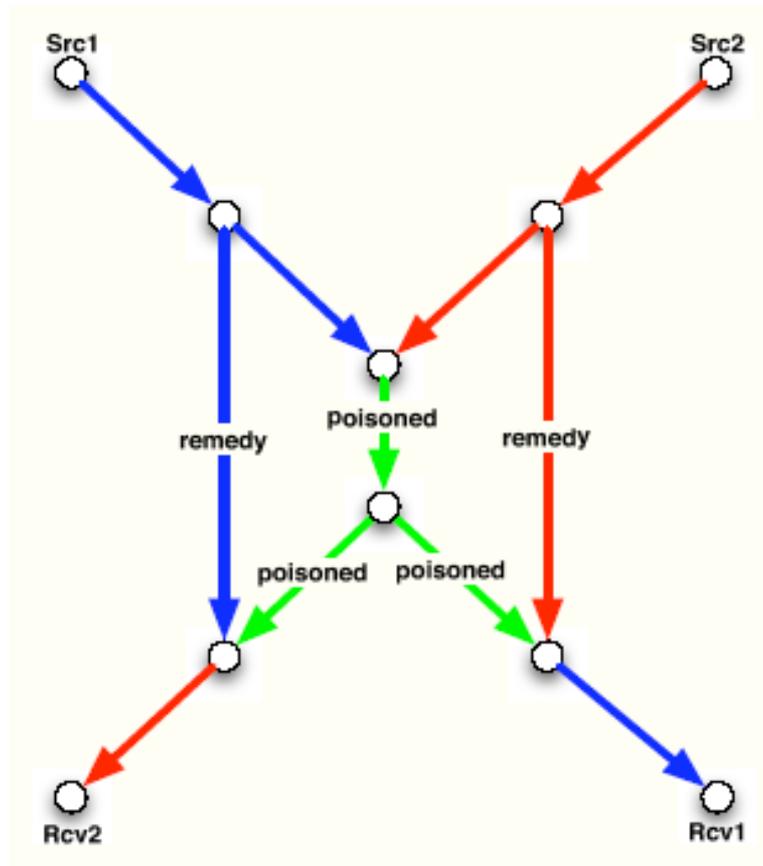
Systematic coding across unicast sessions

- Systematic constructions for network coding across multiple unicasts [RKH05, TRLKM06, EHK06, H06]
 - Aim: to choose routes and network codes so as to take maximum possible advantage of network configurations where network coding is known to improve throughput or reduce transmissions



XOR coding across pairs of unicasts

- Canonical module [RKH05]:



- Reverse carpooling is a special case where remedies do not travel

XOR coding across pairs of unicasts

- If we limit coding to pairs of uncoded or decoded flows, the problem becomes one of optimally fitting together canonical modules
- Can form a linear optimization problem whose constraints are:
 - Conservation of uncoded, poison and remedy flows
 - Conversion rules
- Combinatorial approximation algorithm for reduced complexity [H06]

Part I outline

- Introduction and problem description
- Multicast
 - Algebraic model
 - Constructing multicast network codes
 - Network optimization
- Non-multicast
 - Insufficiency of linear codes
 - Opportunistic wireless network coding
 - Multiple unicast network coding
- Adversarial errors

Adversarial errors

- Network coding needed for optimal rate in multicast and in networks with packet losses and failures
 - Promising near-term applications in peer-to-peer and ad hoc networks; possibility of compromised participating nodes
- Information theoretic techniques for detecting and correcting errors introduced by an adversary who observes and controls unknown subsets of links/transmissions
 - Network coding facilitates use of a subgraph containing multiple paths to each sink, which can help security
 - However, coding at intermediate nodes causes error propagation → traditional approaches not suitable

Model

- Adversary knows the entire message and the coding scheme, but possibly not some of its random choices, e.g. random coding coefficients/random hash functions
- We consider a batch of exogenous source packets transmitted by distributed random network coding to a sink node which may be part of a unicast or multicast
- Adversary injects packets that may contain arbitrary errors; sink receives packets that are random linear combinations of the source and adversarial packets
- We will consider a few variants of this model which differ in terms of the adversary's knowledge and transmission capacity

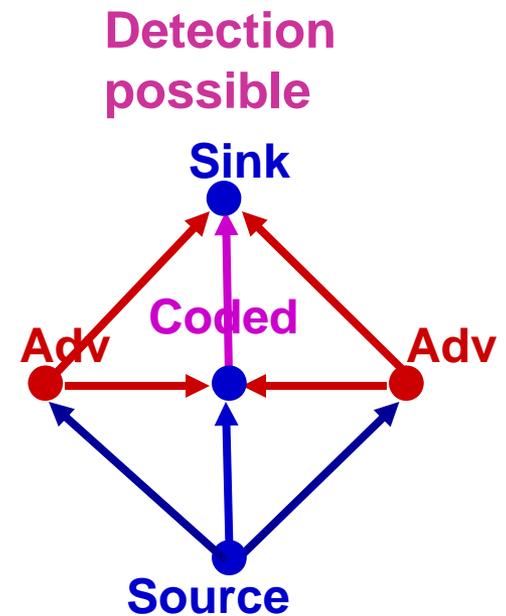
Detection and correction of adversarial errors

- Error detection/error correction capability added to random network coding scheme by adding appropriately designed redundancy; the only changes are at source and sink
- For error correction, overhead is lower bounded in terms of the number of adversarial transmissions as a proportion of the source-sink minimum cut
- For error detection, overhead can be traded off flexibly against detection probability and coding field size
- Error detection scheme can be used for low overhead monitoring when an adversary is not known to be present, in conjunction with a higher overhead error correction scheme activated upon detection of an adversary

Detection of adversarial errors

- Augment each source packet with a flexible number of hash symbols
- As long as not all adversarial packets have been designed with knowledge of the random coding combinations present in all packets received at the sink, adversarial errors result in decoded packets having non-matching data and hash values w.h.p.
- No limit on adversary's transmission capacity, require only that adversary has imperfect knowledge of random code

1	0	A	h_A
0	1	B	h_B



T. Ho, B. Leong, R. Koetter, M. Médard, M. Effros and D. R. Karger, "Byzantine Modification Detection in Multicast Networks with Random Network Coding", submitted to IEEE Transactions on Information Theory, 2006.

Error detection scheme

- Let each source packet contain n header/payload symbols x_1, \dots, x_n and $k < n$ hash symbols h_1, \dots, h_k , where n and k are design parameters which determine overhead

- $h_1 = \varphi(x_1, \dots, x_t) = x_1^2 + \dots + x_t^{t+1}$

...

$$h_k = \varphi(x_{\lfloor (k-1)t \rfloor + 1}, \dots, x_n)$$

where $t = \lceil n/k \rceil$

- Sink observation $Y = TX + UZ$ is the sum of a random linear transform T of source data X and a random linear transform U of adversarial errors Z
- Decoded packets given by $X + T^{-1}UZ$

Error detection performance

- For symbol length $\log q$ bits, if the sink receives s linearly independent combinations of source packets (which may be coded together with any number of adversarial packets), and at least one packet is erroneous, then

a) for at least s decoded packets, the adversary cannot determine which of a set of at least $q-1$ possible values will be obtained

(values can be partitioned into sets of the form $\{v + \lambda w \mid \lambda \in F_q\}$)

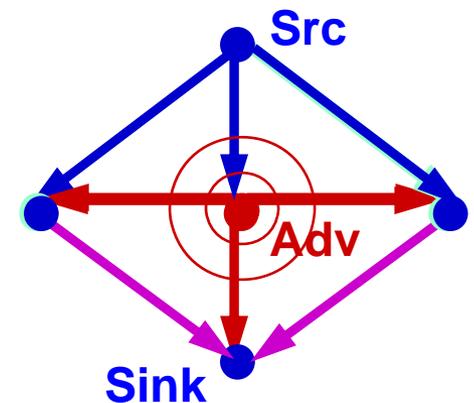
b) the detection probability is at least $1 - ((t+1)/q)^s$

- Example:
 - With 2% overhead ($t=50$), symbol length=7 bits, $s=5$, the detection probability is 98.9%
 - With 1% overhead ($t=100$), symbol length=8 bits, $s=5$, the detection probability is 99.0%

Correction of adversarial errors

- C = capacity from source to sink
- z = capacity from adversary to sink
- n = length of each packet
- Sink receives $Y = TX + UZ = T'X + E$ where
 - coefficient matrix $T' = T + UL$ is $C \times b$
 - source matrix X is $b \times n$
 - error matrix $E = U(Z - LX)$ is $C \times n$, $\text{rank} \leq z$
- Note that if $b \leq C - z$, the column spaces of T' and E are linearly independent w.h.p.

Correction possible



Case 1: Shared secret algorithm

- Source and sink share a low-rate secure secret channel, adversarial capacity $z < C$
- Source uses secret channel to send C random symbols r_1, \dots, r_C and corresponding hash vectors $h(r_i, X) = X [r_i \ r_i^2 \ \dots \ r_i^n]^T$
- Sink calculates syndrome matrix S whose i^{th} column is the difference between $T'h(r_i, X)$ and the corresponding hash of the received data $h(r_i, Y)$, which is in the column space of E
- Since the adversary does not know r_1, \dots, r_C w.h.p. S has the same column space as E
- Since column spaces of T' and E are linearly independent, can solve $Y = T'X + E$ for X
- For $b = C - z$, asymptotically achieves optimal rate

X :

1 0...0	Pkt 1 data
⋮	⋮
0 0...1	Pkt b data

secret:

$r_1 \ \dots \ r_C$
$h(r_1, X) \ \dots \ h(r_C, X)$

Case 2: Omniscient adversary algorithm

- Adversary knows everything, has transmission capacity $z < C/2$
- Source adds $(z+\epsilon)n$ redundant symbols to header/data symbols s.t. resulting value X satisfies $(z+\epsilon)n$ randomly chosen linear constraints and forms $C-z$ packets of n symbols each
- W.h.p. over random constraints and random code, for all q^{zn} possible values of the set of adversarial packets, the sink can construct and solve a system of linear equations to obtain source data
- Optimal rate of $C-2z$ is achieved asymptotically with n

	n		
X:	1 0...0	Pkt 1 data	Redun- dant symbols
	⋮	⋮	
	0 0...1	Pkt C-z data	

Case 3: Limited adversary algorithm

- Adversary observes y transmissions and controls z , where $2z+y < C$
- A small fraction of each packet consists of redundant information generated as follows:
 - Use shared secret algorithm to generate secret hash information
 - Use omniscient adversary algorithm to generate additional redundancy protecting a mix of secret hash information with extra random symbols (for secrecy)
- Sink first decodes secret hash information, then decodes message using shared secret algorithm
- Optimal rate of $C-z$ is asymptotically achieved

Correction of adversarial errors

Common intuition behind algorithms:

- A sink observes the sum of a random linear transform T of data X transmitted by source and a random linear transform U of Z transmitted by adversary
- Design redundancy in source transmissions to satisfy constraints that adversarial data cannot (or is unlikely to) satisfy
- Algebraic decoding algorithms using the observations that:
 - U has rank $\leq z$ (#adversarial transmissions)
 - If $b \leq C-z$, the column spaces of T and U are linearly independent w.h.p.

Correction possible

