On Benefits of Network Coding in Bidirected Networks and Hyper-networks

Zongpeng Li
University of Calgary / INC, CUHK

December 14 2011, at INC, CUHK

Joint work with:

Xunrui Yin, Xin Wang, Jin Zhao, Xiangyang Xue

School of Computer Science
Fudan Univeristy, Shanghai, China

To appear in: IEEE INFOCOM 2012

Benefits of Network Coding: Higher Multicast Rate

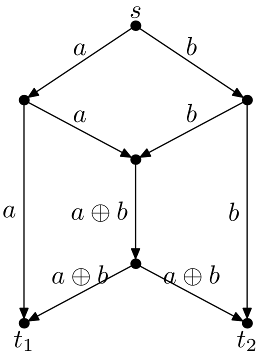

With Network Coding

2 bits / 1 sec

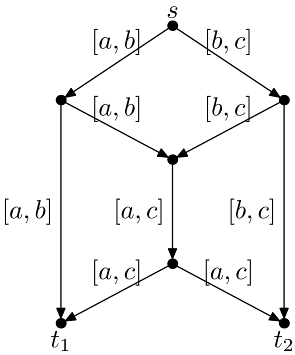

Without Network Coding

3 bits / 2 secs

Motivation

- Coding advantage:
  - ratio of max throughput: with network coding over without network coding
  - always ≥ 1, by definition
- Classic models:
  - Directed Network
    - Arbitrarily large coding advantage in contrived topologies
  - Undirected Network
    - coding advantage upper bounded by 2

Motivation

- We study two parameterized, more general network models
  - Bidirected Networks (with max link imbalance α)
    - Directed network: special case, $\alpha = \infty$.
    - Internet core has small α close to 1.
    - Hope to reveal the connection between the coding advantage and the level of link asymmetry
  - Hyper-Networks (with max link size β)
    - Undirected network: special case, $\beta = 2$.
    - Models Ethernet buses and broadcast links.
    - Hope to generalize the upper-bound of 2 proven for undirected networks
    - Technique: transform a hyper-network into a bidirected network for analysis.

General Network Models

- A network is represented as a (multi-)graph $G(V, E)$ where
  - each link has unit capacity
  - parallel links allowed
- A multicast session (s,T):
  - $s \in V$: the multicast source
  - $T \subset V$: the set of multicast receivers
- A (symmetrical) multicast throughput $\mathcal{R}$ is achieved if *each receiver* receives information at rate $\mathcal{R}$.

Max. Throughput with Network Coding

Theorem 1 (Ahlswede, Cai, Li, Yeung, IT2000)

In a directed network,

$$\mathcal{R}_{nc} = \min_{t \in T}\{\lambda_G(s,t)\}$$

- $\mathcal{R}_{nc}$: max multicast throughput with network coding
- $\lambda_G(s,t)$: edge connectivity from s to t.
  - *i.e.*, the number of edge disjoint paths from s to t.

Example

Recall in the butterfly network, $\mathcal{R}_{nc} = 2$:

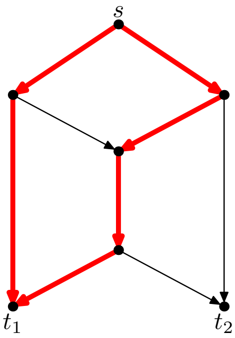

edge disjoint paths to t_1

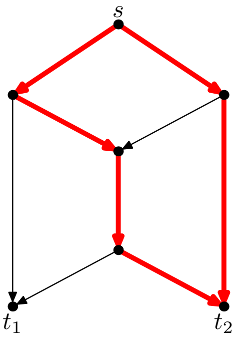

edge disjoint paths to t_2

Max Throughput with Routing

- Without network coding, symbols can still be replicated.
- The trace of each symbol forms a multicast tree.
- Packing multicast trees: deciding transmission rates for possible multicast trees, under link capacity constraints.

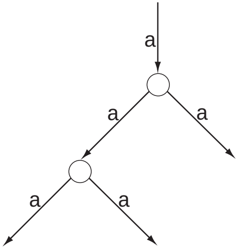

Proposition 1

Without network coding, the max multicast rate $\mathcal{R}_{tree}$ is achieved by an optimal packing of multicast trees.

Example

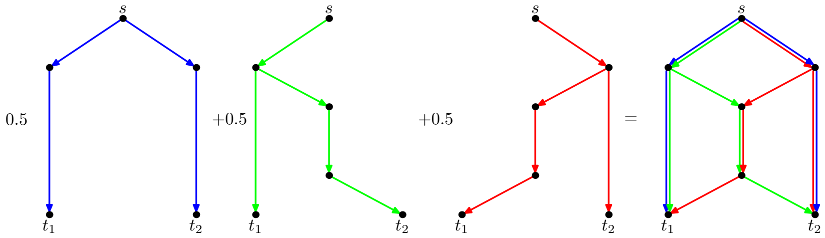

In the butterfly network, $\mathcal{R}_{tree} = 0.5 + 0.5 + 0.5 = 1.5$.

Coding Advantage

Definition

Given topology $G(V, E)$ and multicast session (s, T),
Coding Advantage is defined as $\theta = \mathcal{R}_{nc}/\mathcal{R}_{tree}$

- In the butterfly network, $\theta = 2/1.5 \doteq 1.33$.

Coding Advantage

Question

$$\max_{G,s,T} \theta = ?$$

- In terms of throughput improvement
  - How good can Network Coding be?
  - In which scenario, Network Coding outperforms Routing the most?

Previous Results

Scenario	Coding Advantage	Note
In general	$\theta \to \infty$	Illustrated later.
Unicast or Broadcast	$\theta = 1$	$\|T\| = 1$ or $T = V \backslash \{s\}$
Most P2P Overlay Networks	$\theta = 1$	sufficient down link capacity
Undirected Networks	$\theta \leq 2$	$f(u,v) + f(v,u) \leq c(\{u,v\})$

Previous Results

Scenario	Coding Advantage	Note
In general	$\theta \to \infty$	Illustrated later.
Unicast or Broadcast	$\theta = 1$	$\lvert T \rvert = 1$ or $T = V \backslash \{s\}$
Most P2P Overlay Networks	$\theta = 1$	sufficient down link capacity
Undirected Networks	$\theta \leq 2$	$f(u,v) + f(v,u) \leq c(\{u,v\})$

Previous Results

Scenario	Coding Advantage	Note
In general	$\theta \to \infty$	Illustrated later.
Unicast or Broadcast	$\theta = 1$	$\|T\| = 1$ or $T = V \backslash \{s\}$
Most P2P Overlay Networks	$\theta = 1$	sufficient down link capacity
Undirected Networks	$\theta \leq 2$	$f(u,v) + f(v,u) \leq c(\{u,v\})$

Previous Results

Scenario	Coding Advantage	Note
In general	$\theta \to \infty$	Illustrated later.
Unicast or Broadcast	$\theta = 1$	$\|T\| = 1$ or $T = V \backslash \{s\}$
Most P2P Overlay Networks	$\theta = 1$	sufficient down link capacity
Undirected Networks	$\theta \leq 2$	$f(u,v) + f(v,u) \leq c(\{u,v\})$

Example with Large Coding Advantage

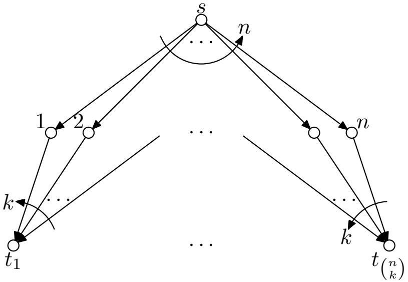

- $\mathcal{R}_{nc} = k$, $\mathcal{R}_{tree} \leq \frac{n}{n-k+1}$.
- $\theta \geq \frac{k(n-k+1)}{n} \to \infty$, as $n = 2k \to \infty$.

Observation

In practice, links are often bidirected.

Bidirected Networks in Practice

- Rocketfuel project: 6 ISP topologies, all bidirected, with symmetrical link capacity in opposite directions
- Fraleight *et al.*: The closer to the Internet backbone, the more symmetric the traffic is.
  - OC-48 links examined: traffic ratio between 1:1 and 5:1.
- John and Tafvelin: backbone OC-192 link, no significant difference in opposite link traffic volume
- Often, a single number is associated with an Internet link for representing its capacity — implicitly assuming the capacities in opposite directions are equal.
- Fastest Internet Link: 2011 experiment
  - Victoria, BC $\rightarrow$ Seattle, WA: 98 GB/sec
  - Seattle, WA $\rightarrow$ Victoria, BC: 88 GB/sec

Case Study

What happens when links are bidirected?

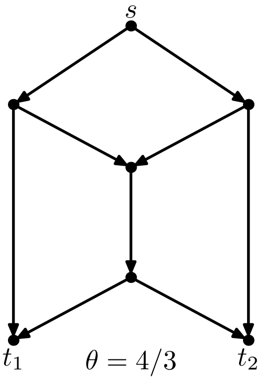

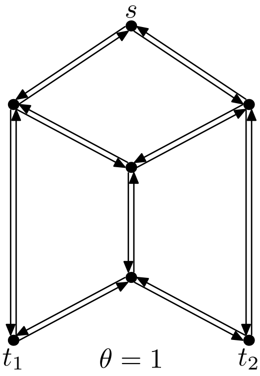

Completely Balanced Networks

It is not a coincidence!

Theorem 2

In completely link balanced networks, $\theta = 1$.

- Let $c(u, v)$ denote the actual link capacity from u to v, *i.e.*, the number of parallel unit capacity links from u to v.
- A bidirected network is completely link balanced, if $c(u, v) = c(v, u), \forall u, v \in V$.

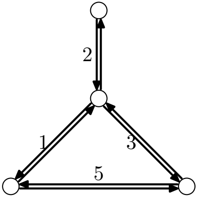

link balanced

Proof of theorem 2

Proof sketch:

1. Convert the completely link balanced network B into a broadcast network D such that
   - $\mathcal{R}_{nc}(B) = \mathcal{R}_{nc}(D)$.
   - $\mathcal{R}_{tree}(B) \geq \mathcal{R}_{tree}(D)$.
2. Apply the fact network coding can not increase multicast rate when $T = V \backslash \{s\}$, we have $\mathcal{R}_{tree}(D) = \mathcal{R}_{nc}(D)$ and thereby, $\mathcal{R}_{tree}(B) = \mathcal{R}_{nc}(B)$.

For the conversion, we employ edge splitting to isolate each relay node.

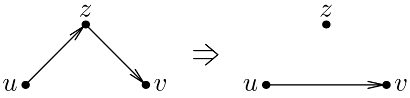

Proof of theorem 2

A relay node can be split off without affecting the edge connectivity among other nodes. $\Rightarrow \mathcal{R}_{nc}(B) = \mathcal{R}_{nc}(D)$

Lemma

[Frank & Jackson 1995] Let $D = (V + z, E)$ be a node balanced directed graph. For each link $\overrightarrow{uz} \in E$, there exists a link $\overrightarrow{zv} \in E$, such that after splitting off $\overrightarrow{uz}, \overrightarrow{zv}$, the edge connectivity between every pair of nodes in V is unchanged.

Each multicast tree in the resulting network D can be converted into a multicast tree in the original network B.
$\Rightarrow \mathcal{R}_{tree}(B) \geq \mathcal{R}_{tree}(D)$

Remarks

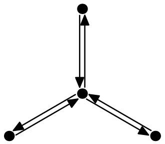

link balanced

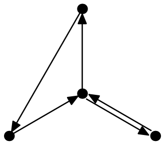

node balanced

- 'link balanced' can be relaxed to 'node balanced'.
- The core of Internet is close to a link balanced network.
- From the proof: neither coding (network coding) or replication (IP multicast) is necessary at interior routers.
  - Each split corresponds to a forwarding at the interior routers.
- A polynomial time algorithm for tree-packing can be extracted from the proof.

α-Balanced Networks

For a general bidirected network, define the link imbalance ratio

$$\alpha = \max_{c(u,v)>0} \frac{c(v,u)}{c(u,v)}$$

Theorem 3

In bidirected networks, $\theta \leq \alpha$.

Proof: Ignoring the excess capacity, we can perform an optimal tree packing in the resulting link balanced network, achieving a multicast throughput no less than $\frac{1}{\alpha}\mathcal{R}_{nc}$.

α-Balanced Networks

Proposition 2

For $\alpha \geq 1$, there exists an α-balanced network, where $\theta \geq \sqrt{\alpha}/4$.

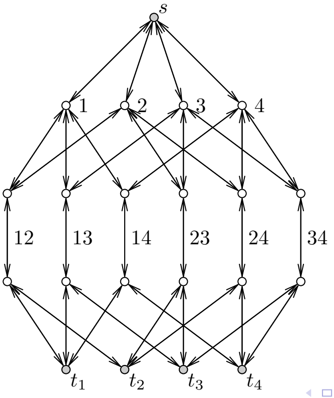

Undirected Network

- Bidirected network: for each pair of adjacent nodes u, v, $c(u, v)$ and $c(v, u)$ are fixed and independent of each other.
- Undirected network: the two directions share a total link capacity $c(\{u, v\})$
  - Let $f(u, v)$ denote the information flow rate from u to v
  - Link capacity constraints: $f(u, v) + f(v, u) \leq c(\{u, v\})$.

Theorem 4 (Li and Li, CISS2004)

In an undirected multicast network, $\theta \leq 2$.

Example

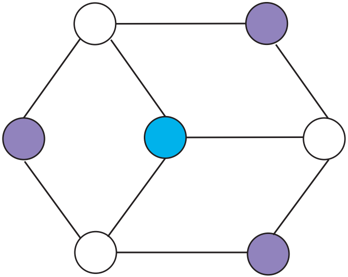

Example

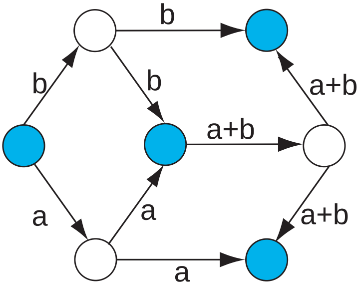

Example

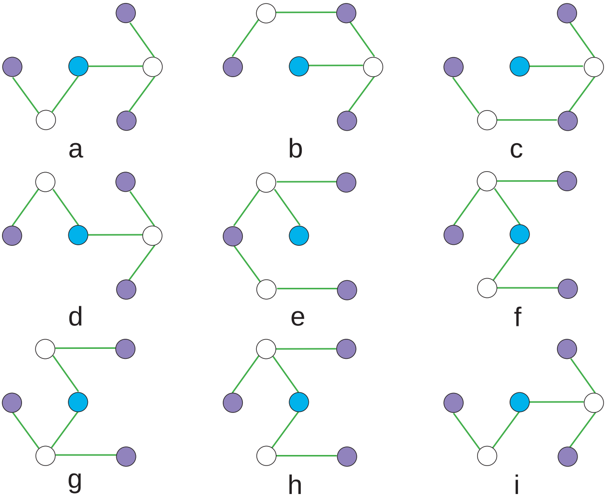

Example

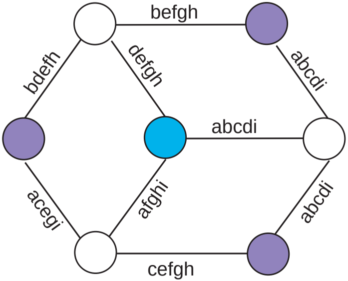

- Total rate: $\mathcal{R}_{tree} = 0.2 \times 9 = 1.8$
- $\theta = \mathcal{R}_{nc}/\mathcal{R}_{tree} = 2/1.8 = 10/9 \in [1, 2]$
- Optimal tree packing in general is NP-hard

Hyper-Network as an Extension

Hyper-network

- A (hyper-)link connect 2 or more nodes.
- When one node transmits through a hyper-link, all other nodes can simultaneously receive.
- For example: wireless link, Ethernet bus.

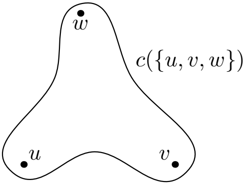

- Let $f(u \to vw)$ denote the transmission rate from u to v, w.
- Link capacity constraint: $f(u \to vw) + f(v \to uw) + f(w \to uv) \leq c(\{u, v, w\})$.

Hyper-networks

The size/cardinality of a hyper-link is defined as the number of nodes it covers.

Theorem 5

In a hyper-network with max edge size β, $\theta \leq \beta$.

An undirected network is a special case with $\beta = 2$.

Proof of theorem 5

Proof sketch:

1. Given Hyper-network H, construct a completely link balanced network B as follows:

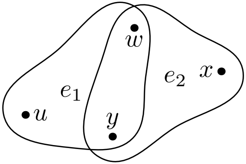

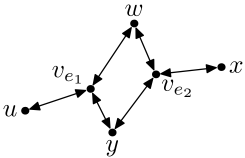

2. According to theorem 2, $\mathcal{R}_{nc}(B) = \mathcal{R}_{tree}(B)$. It suffices to verify (i) $\mathcal{R}_{nc}(H) \leq \mathcal{R}_{nc}(B)$ and (ii) $\mathcal{R}_{tree}(H) \geq \frac{1}{\beta}\mathcal{R}_{tree}(B)$.

Lower bound

Proposition 3

There exists a hyper-network (H, s, T) *with max edge size* β*, such that* $\theta \geq \frac{1}{4} \log \beta$.

Proof sketch: Consider the relay nodes of the combination network as hyper-links connecting the source and the receivers. Coding advantage in this hyper-network is the same as in the directed combination network, while the size of each hyper-link is $\binom{n-1}{k-1} + 1$.

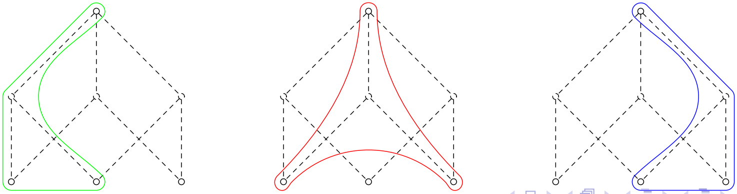

Conclusion

- Completely balanced networks: end-system multicast suffices
- Bidirected networks: coding advantage depends on network symmetry
- Hyper-networks: coding advantage depends on hyper-link size
- Open problem: close gaps between upper-bounds and lower-bounds

Seminars at INC

- 2011.08.17. Space Information Flow
- 2011.12.14. On Benefits of Network Coding in Bidirected Networks and Hyper-networks
- 2012. A Geometric Framework for Studying the Multiple Unicast Network Coding Conjecture
- 2012. Network Coding in Planar Networks

Thank you!

There is nothing more practical than a good theory.

— Kurt Lewin, 1945