Loss Tomography in General Topologies with Network Coding

Minas Gjoka, Christina Fragouli*, Pegah Sattari, Athina Markopoulou
UC Irvine and EPFL, Lausanne*

Problem Context

- An old problem:
  - Network Monitoring/Inference
- A new set of techniques:
  - Network Coding
- Goal of this work:
  - Link loss inference in *general topologies* with network coding capabilities?

Outline

- Problem statement
- Prior Work
- Challenges with general topologies
  - Dealing with cycles
  - Link Identifiability
  - Estimation Algorithm
- Simulations
- Conclusions

What is Network Tomography

- Goal: obtain detailed picture of a network from end-to-end views

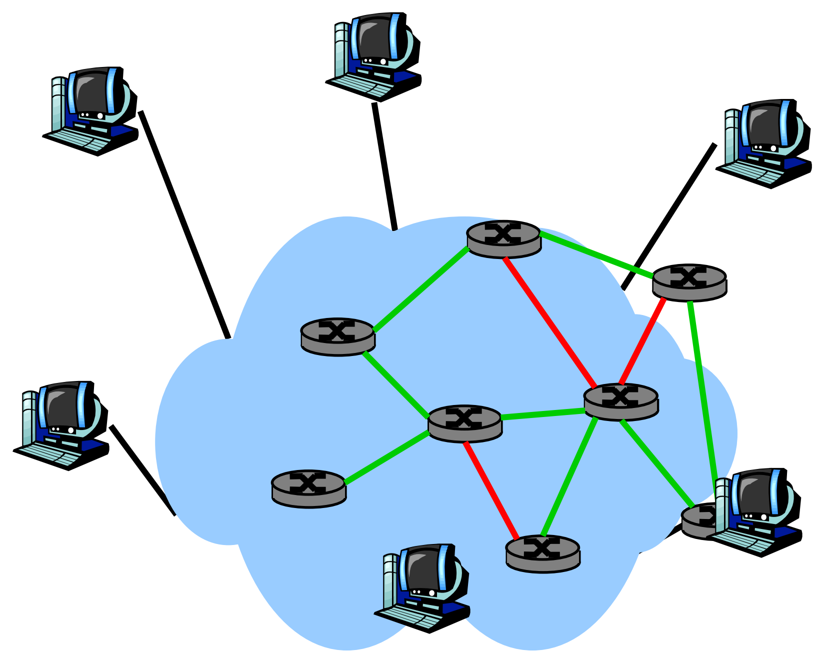

The Network Coding Paradigm

- Allow intermediate nodes to perform operations on incoming packets before forwarding them

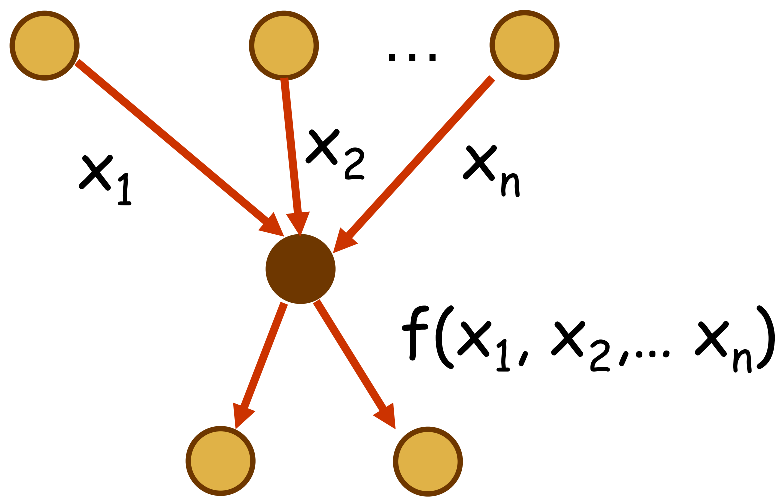

- Potential benefits at the cost of processing at intermediate nodes

Problem Statement

- *Assuming that network coding is already deployed in a general topology network, how to exploit it for network inference ?*
- New opportunities by network coding (NC):
  - NC introduces correlation → exploit for inference
  - NC combined flows into one→ bandwidth savings
- New challenges introduced in general topologies

Challenges of General Topologies

- More than one path for every Source, Receiver pair.

- Challenges:
  - Graphs with cycles
  - Link Identifiability
  - Estimation Algorithm

Link Loss Inference
Problem Statement

- Setup:
  - Topology is known (typically a general graph)
  - [Bernoulli] losses $\{a_k, k \in L\}$, independent links
  - Can send and observe end-to-end probes
- Given:
  - Observations at all receivers $X=\{X_{n,}\ n=1..N\}$
- Goal:
  - estimate link loss probabilities $\alpha = \{\alpha_k : k \in L\}$

Outline

- Problem statement
- Prior Work
- Challenges with general topologies
  - Dealing with cycles
  - Link Identifiability
  - Estimation Algorithm
- Simulations
- Conclusions

Relevant Work

- [Lo Presti & Duffield]
  - General topologies: covering the network with several trees
- [Rabbat, Nowak '04]
  - Multiple sources, unicast probes share fate
  - Joint topology and link loss inference
- [Yan et al., Sigcomm 04-06]
  - Overlay network measurements

Computationally efficient, suboptimal algorithms vs. MLE

Loss Inference

Basic Example

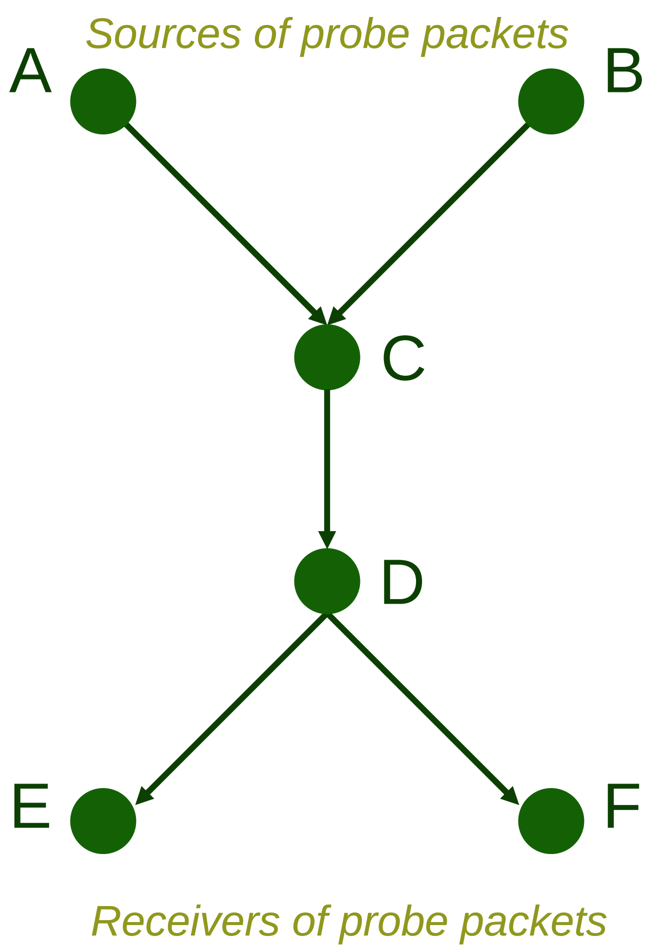

- We want to infer the link loss rates a_k on all links $k \in$ {BC, AC, **CD**, DE, DF}
- using end-to-end probes from {A,B}, to {E,F}

Covering the graph with trees

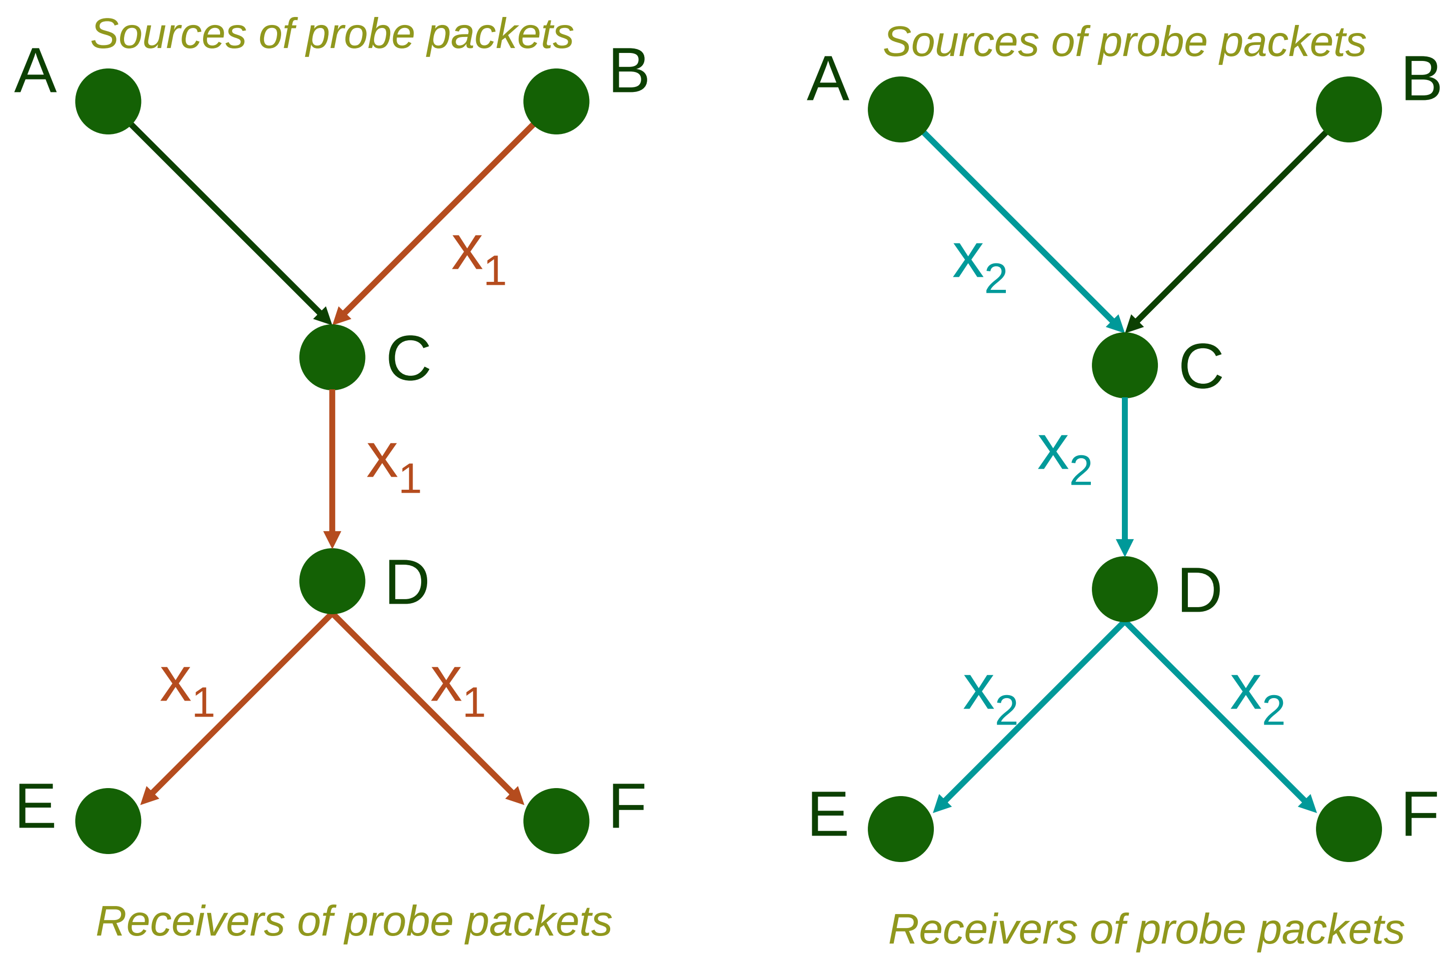

Drawbacks - Multiple trees

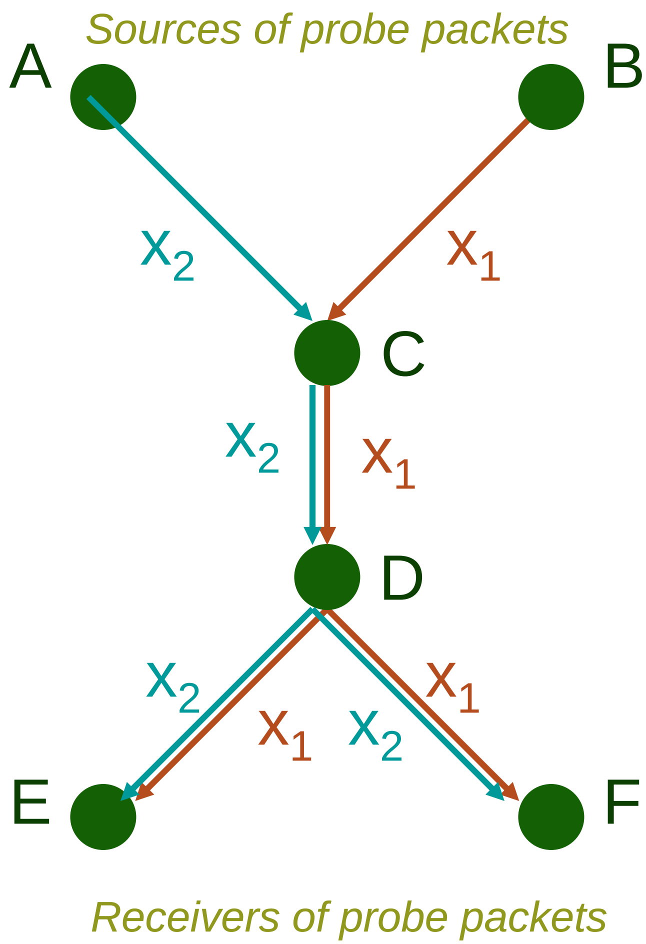

1. We cannot infer the loss rate for edge CD
2. Paths overlap from C and downstream
3. Minimum cost covering with multicast trees is NP-hard
4. Combining observations from 2 trees leads to suboptimal estimation

Network coding approach

[Fragouli, Markopoulou], Allerton 05
[Fragouli, Markopoulou, Srinivasan, Diggavi], ITA 07

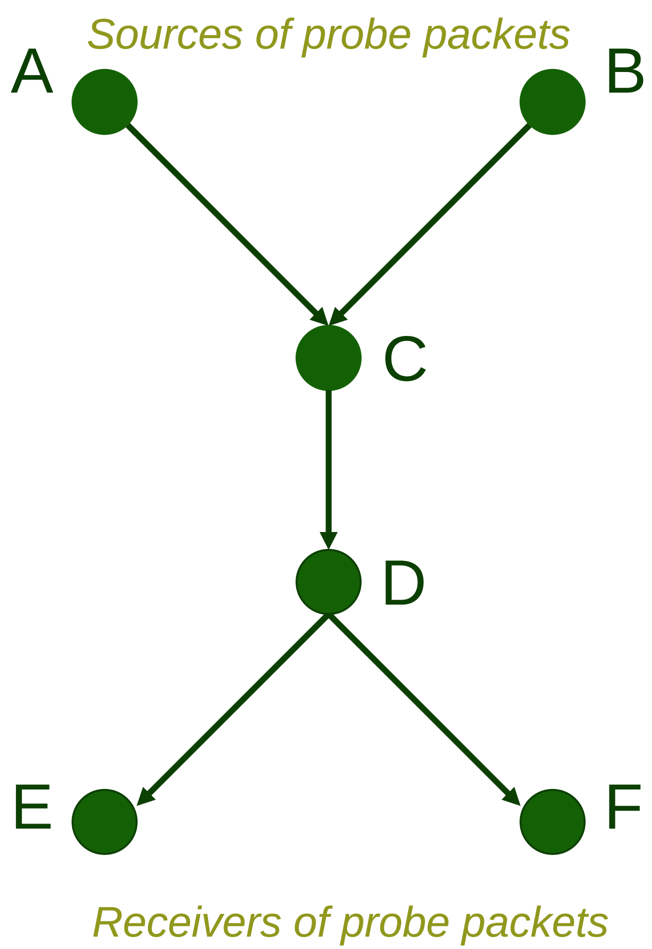

Intermediate node (C):

Within a time window

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

Our prior work covers the network coding approach on tree topologies

New challenges need to be addressed for general topologies

Outline

- Problem statement
- Prior Work
- Challenges with general topologies
  - Dealing with cycles
  - Link Identifiability
  - Estimation Algorithm
- Simulations
- Conclusions

Dealing with cycles

Start from a set of nodes that act as sources. Select an orientation of the graph and a set of receivers so that we have:

- acyclic graph
- minimize the number of receivers
- allows the maximum number of links to be identifiable

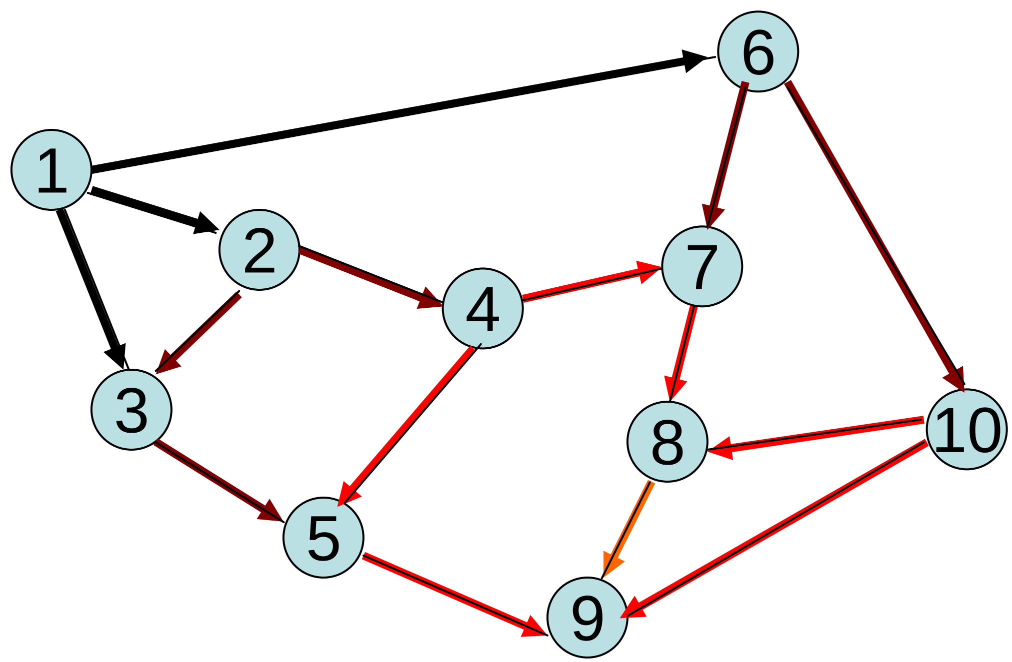

Dealing with cycles (2)

- Acyclic orientation for an undirected graph is possible since the receivers are one of the outputs of the algorithm.
- Breath-first progression creates paths with roughly the same length.
- Reversing orientation allows measuring opposite direction.

Identifiability
Field Size of Coding Scheme

Example with XOR operations at coding **points** - F_2

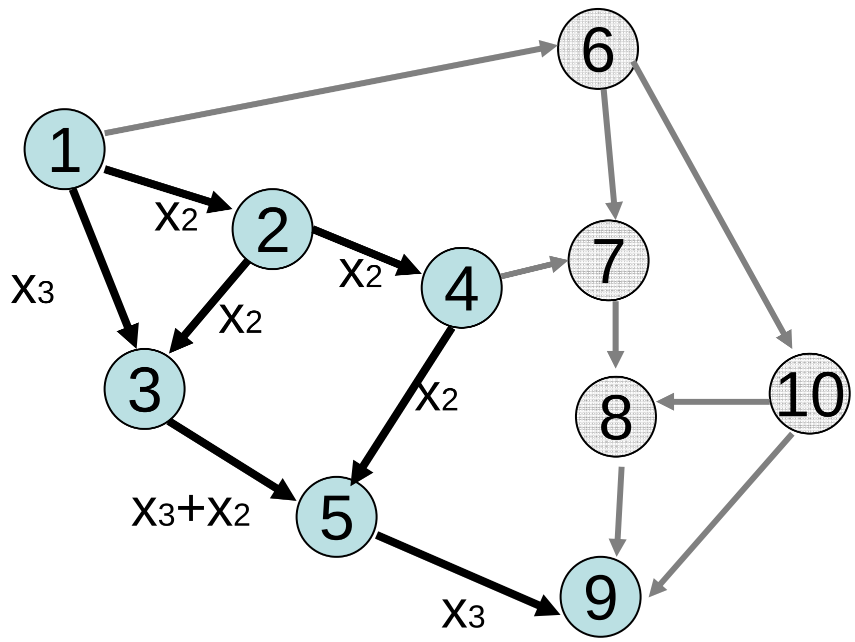

All links functional

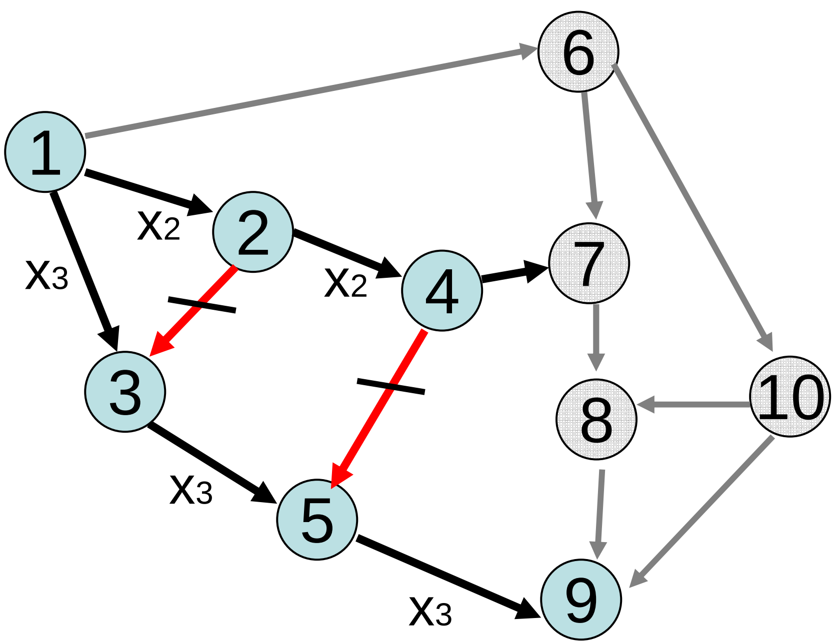

Links 2-3 & 4-5 nonfunctional

Identifiability
Field Size of Coding Scheme

Let $P(e_R)$ be set of paths that connect sources to receiver R and have e_R as their last edge.

Probe coding scheme valid if we can determine which of the $P(e_R)$ paths were functioning from received probes at edge e_R

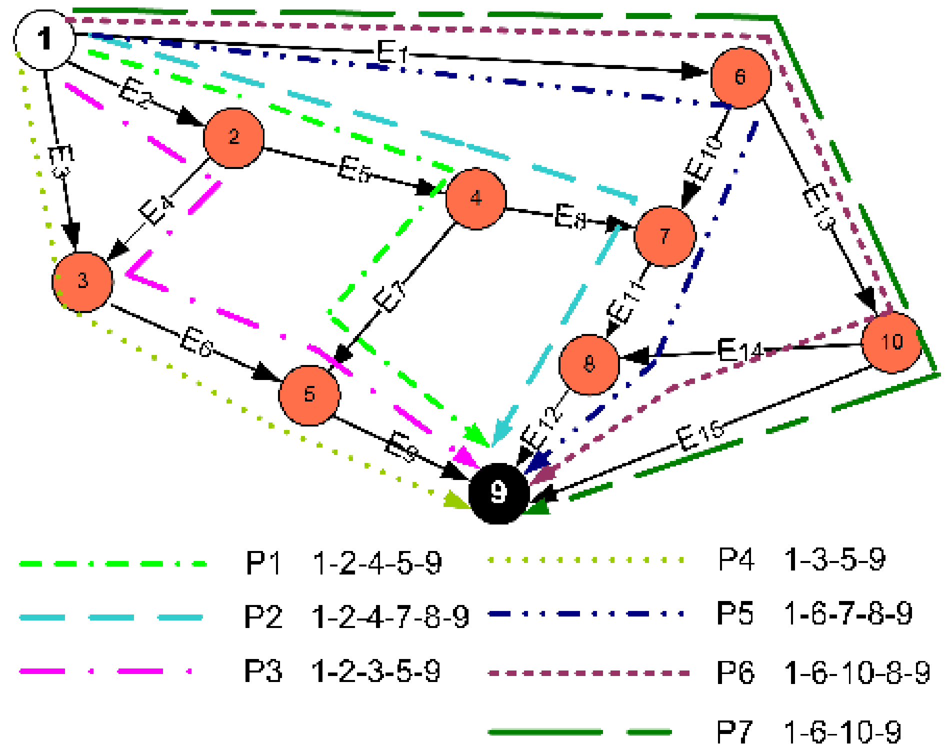

Field size F_q where $q \geq \max_{e_R} P(e_R).$

In practise…

Belief Propagation Estimation Algorithm

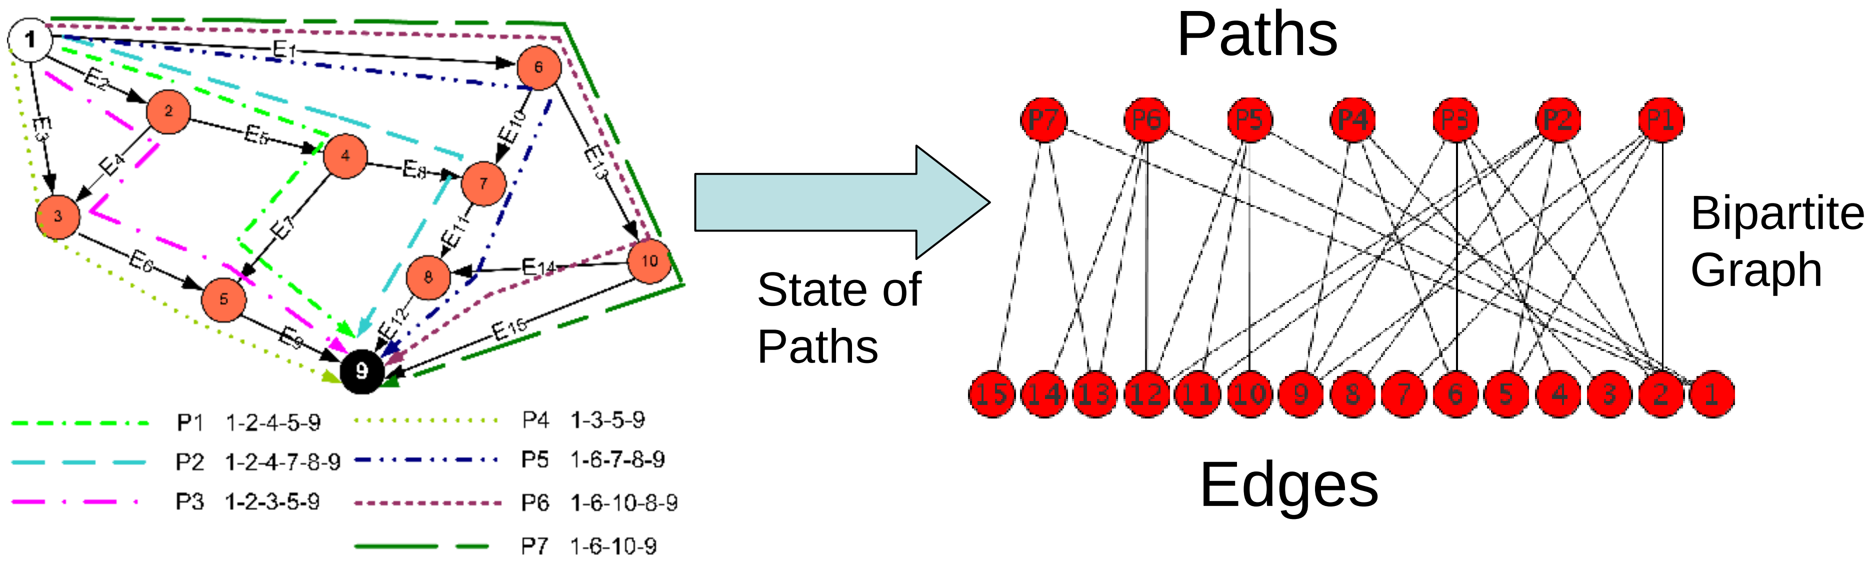

Belief propagation is a sum-product type of algorithm. We use it to infer state of links given the state of the paths for each probe

Loss Estimation at a glance

1. Choose set of sources
2. Produce orientation and prepare coding scheme
3. Send probes at the sources. From the observation at the receivers deduce state of paths.
4. Use Belief Propagation to infer state of links from state of paths.

Outline

- Problem statement
- Prior Work
- Challenges with general topologies
  - Dealing with cycles
  - Link Identifiability
  - Estimation Algorithm
- Simulations
- Conclusions

Simulations

- Abilene-like Topology – 10 nodes & 15 links
- Exodus Topology- 48 nodes & 105 links

- Properties of Oriented Topologies
- Belief Propagation Estimation Accuracy

Orientation Algorithm

- Number of receivers
- Number of distinct paths per receiver
- Number of paths per link and links per path

Orientation Algorithm
Exodus

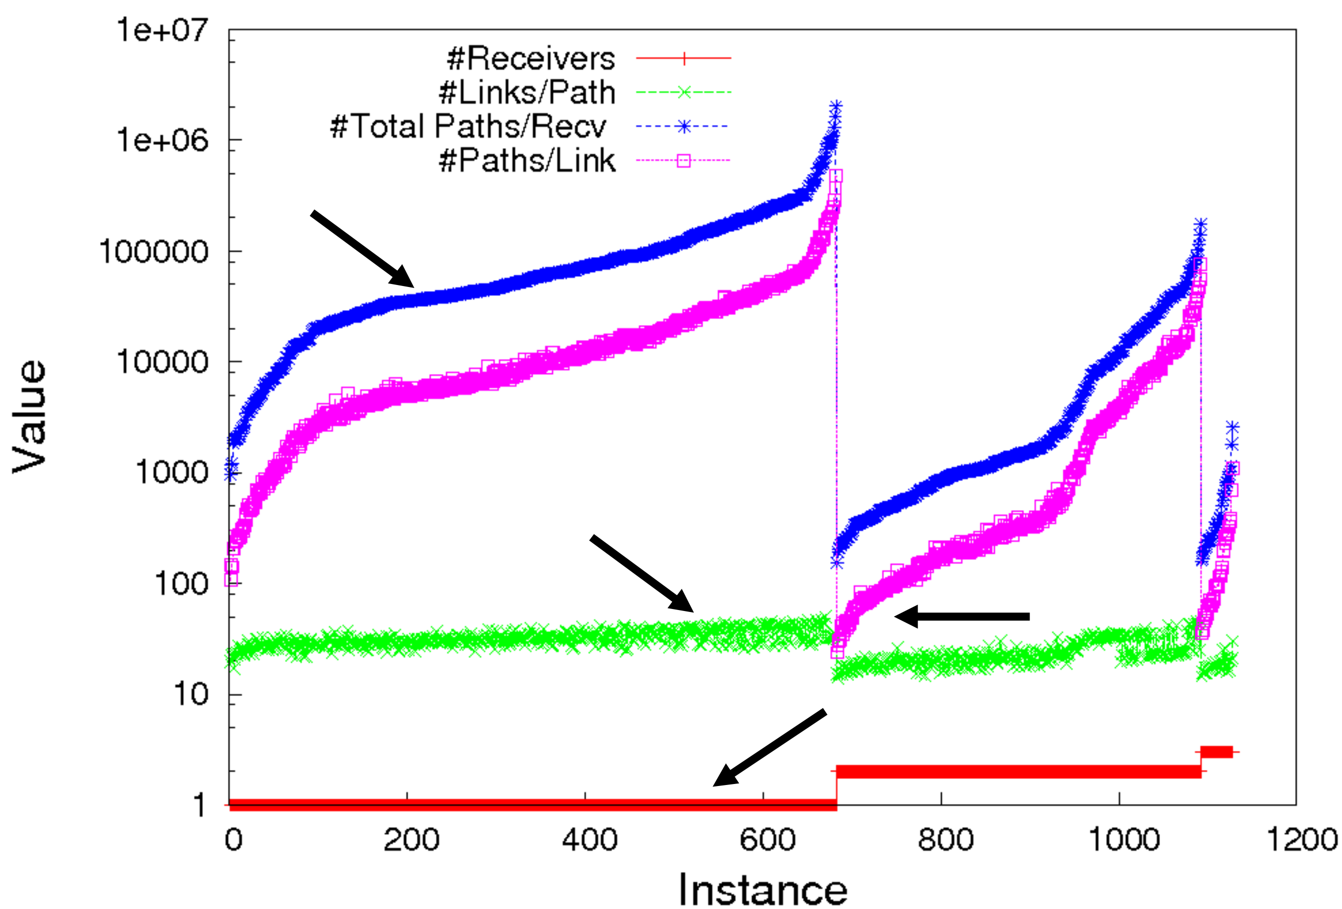

All possible placements of two sources

Estimation Accuracy Metrics

- Mean Square Error for one link (e):

$$MSE = E\left[\left|\alpha_e - \hat{\alpha}_e\right|^2\right]$$

- Entropy Measure for all links:

$$ENT = \sum_e \log MSE(e)$$

where

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

Belief Propagation simulations Abilene-like Topology

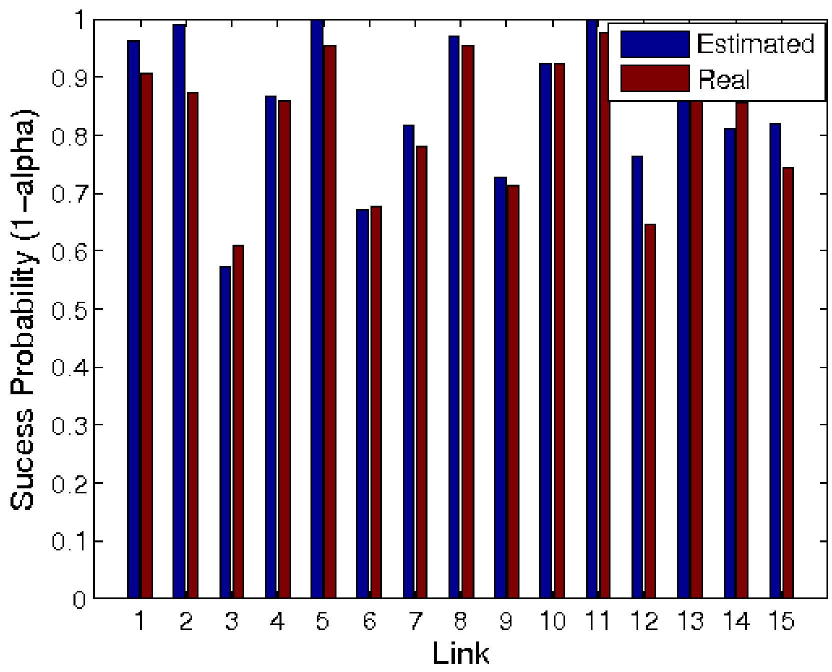

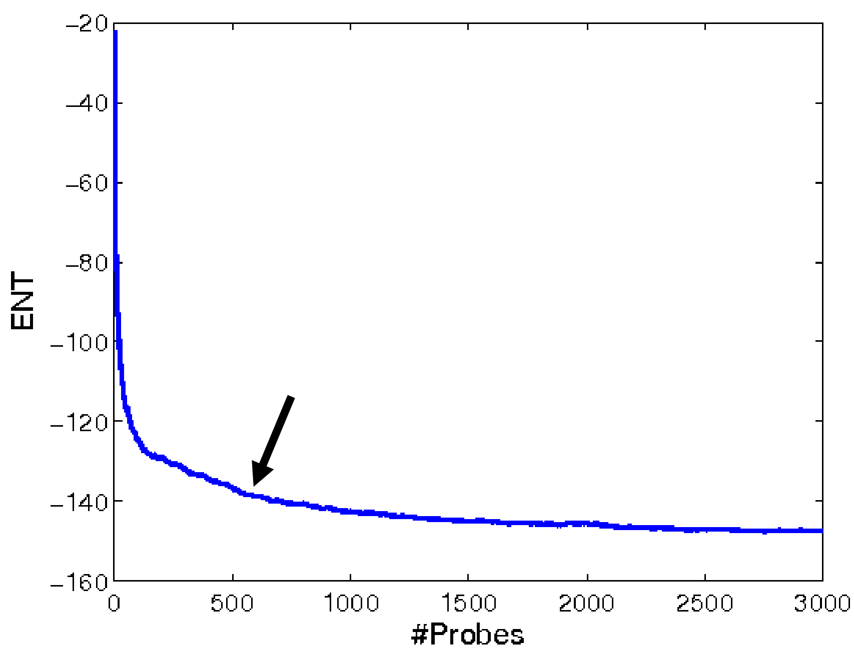

Real loss rates assumed inversely proportional to the link bandwidth with 17% average loss rate

10 Realizations

Belief Propagation simulations Abilene-like Topology (2)

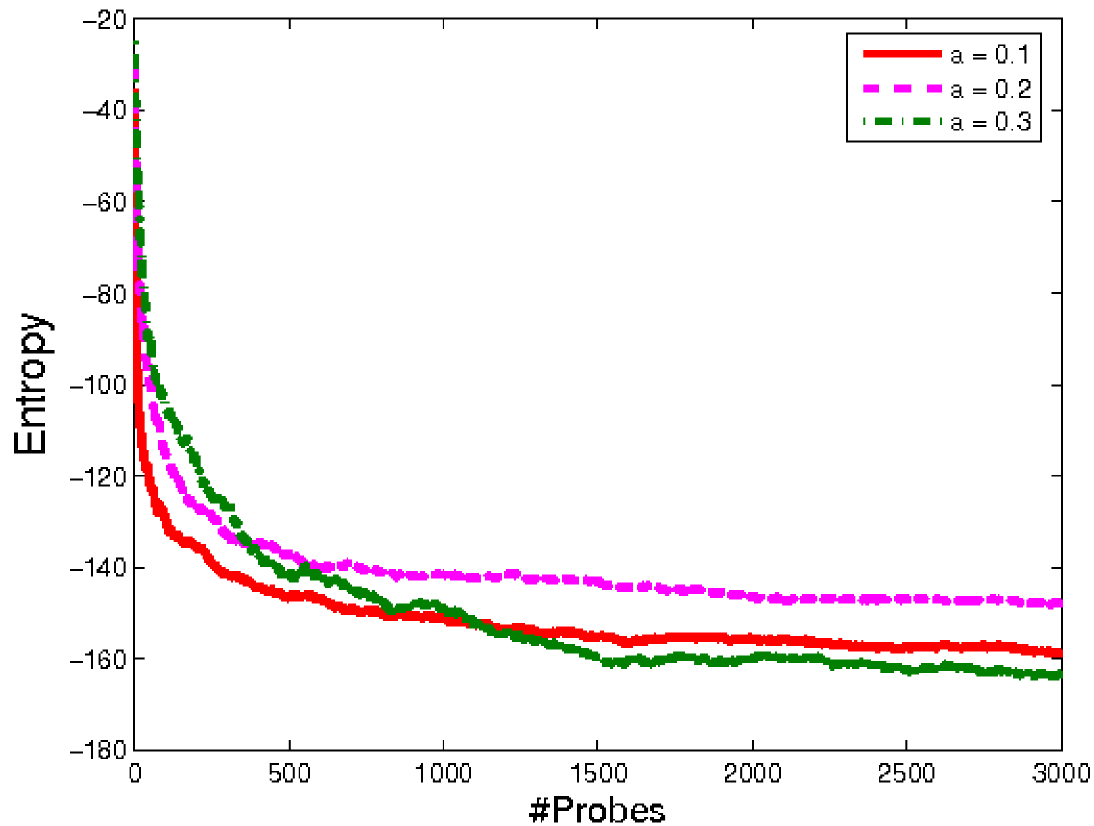

Loss rates same for all links
10 Realizations

How does the # of sources and point of view affect inference?

Src-Rcv	Entropy		
	a=0.05	a=0.15	0.25
{1}-{9}	-178.6	-147.9	-161.6
{9}-{2}	-176.1	-155.8	-166.6
{1,9}-{7}	-189.3	-166.5	-171.7
{9,6}-{4}	-186.9	-169.5	-173.2
{1,5,9}-{7}	-199.8	180.9	-172.3

Link Loss Inference using NC

Summary

- Link loss inference improves with multiple sources and Network Coding
  - Intuition: every probe has more information + multiple sources come for free thanks to NC (1 probe/link)
- A new framework for link loss inference in general graphs with unicast, multicast and network coding
  - selection of sources/receivers, avoiding cycles, code design, estimation algorithm

Future work

- Code design. Improve identifiability and estimation
- Comparison with traditional tree/path-packing approaches in general graphs
- Applications

Thank you