Network Tomography

Christos Gkantsidis

Introduction

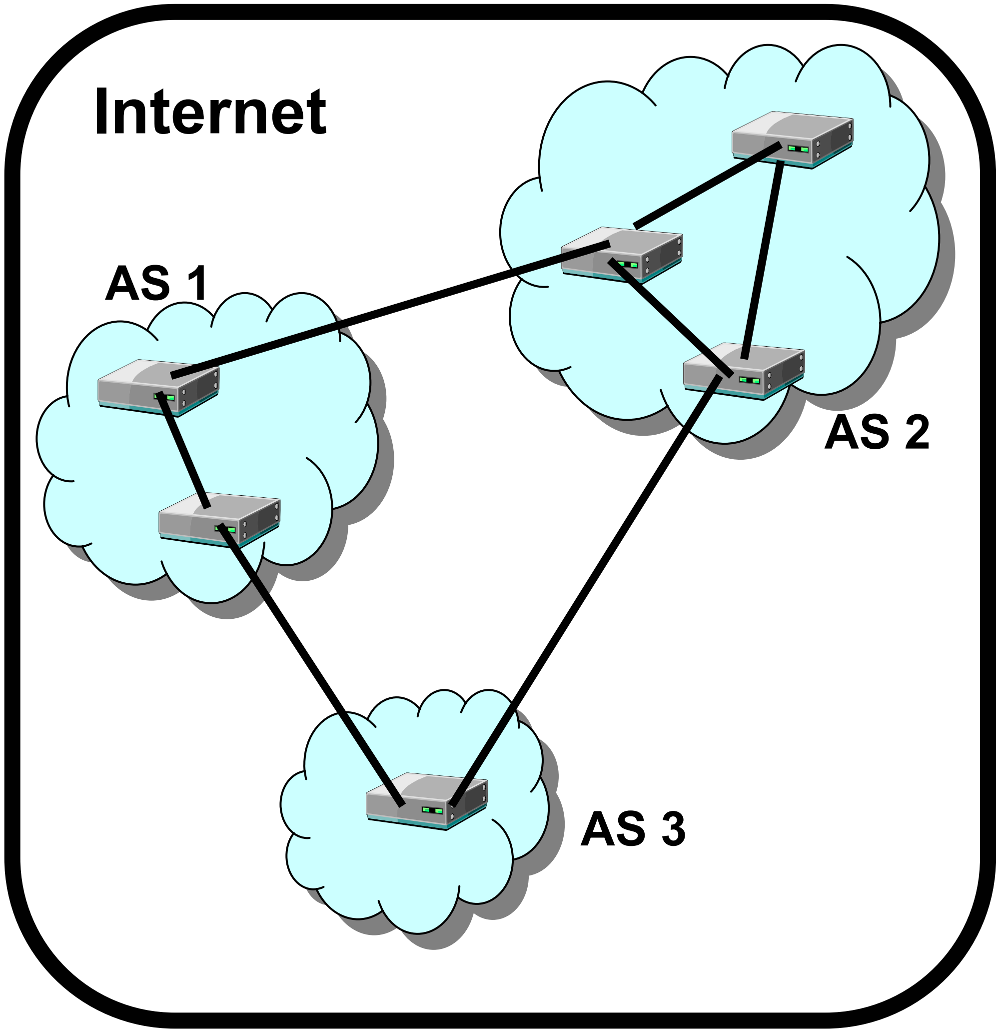

What is the:

- Bandwidth
- Loss Rate
- Connectivity

of the **links** of the network?

Using **only end-to-end** measurements.

Introduction

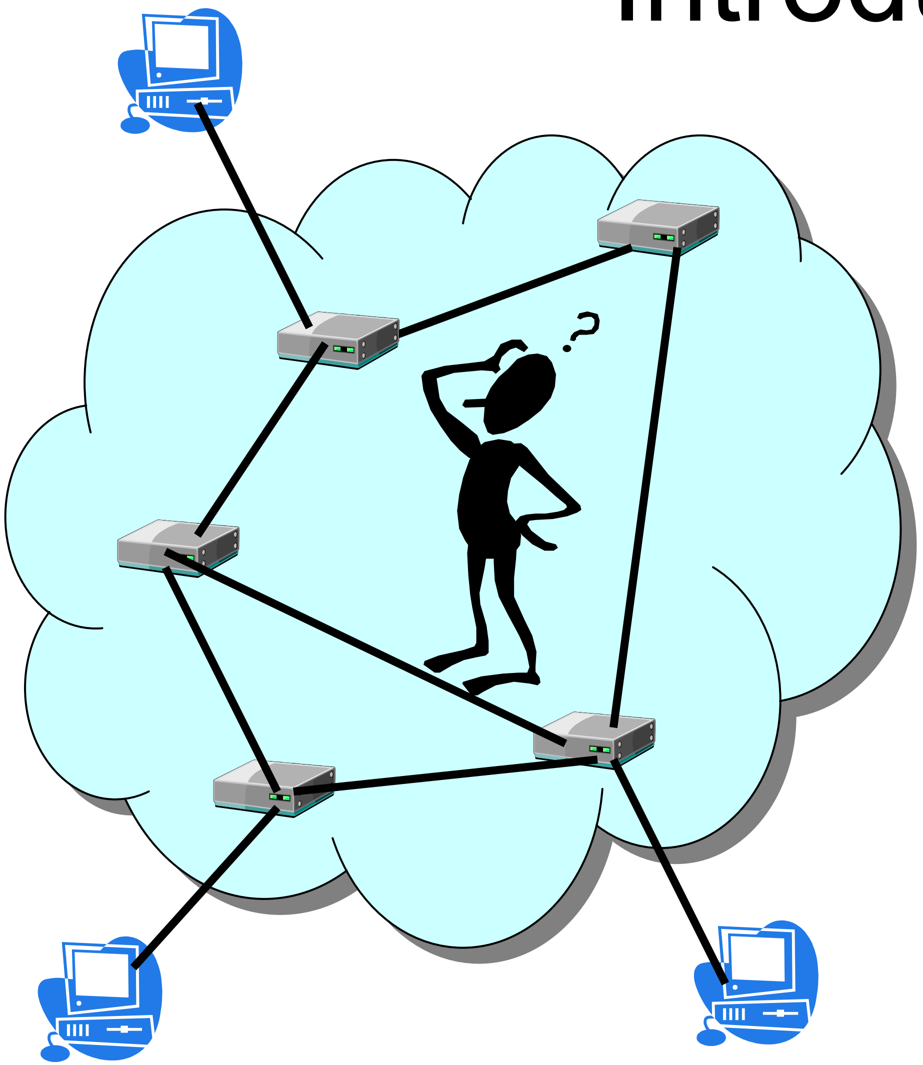

What is the:

• Traffic demands

between **users** of the network?

Using **only limited link** measurements.

Network Tomography

Use a **limited** number of measurements to **infer** network (**link**) performance parameters, using:

- Maximum Likelihood Estimator.
- Bayesian Inference.

and assuming a **prior model**.

Categories of problems:

- Link level parameter estimation.
- Topology Inference.
- Sender-Receiver traffic intensity.

Sender-Receiver Traffic Intensity

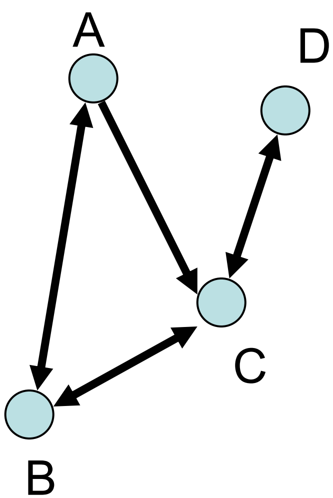

Routing Matrix

Links

	1	2	3	4	5	6	7	8	9	10	11	12
	ab	ac	ad	ba	bc	bd	ca	cb	cd	da	db	dc
1 ($a \rightarrow b$)	1	0	0	0	0	0	0	0	0	0	0	0
2 ($a \rightarrow c$)	0	1	1	0	0	1	0	0	0	0	0	0
3 ($b \rightarrow a$)	0	0	0	1	0	1	1	0	0	1	0	0
4 ($b \rightarrow c$)	0	0	0	0	1	0	0	0	0	0	0	0
5 ($c \rightarrow b$)	0	0	0	0	0	0	1	1	0	1	1	0
6 ($c \rightarrow d$)	0	0	1	0	0	1	0	0	1	0	0	0
7 ($d \rightarrow c$)	0	0	0	0	0	0	0	0	0	1	1	1

Source-Destination Pairs

Sender-Receiver Traffic Intensity

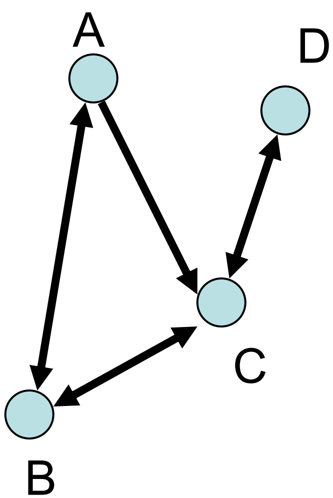

Routing Matrix

Links

	1	2	3	4	5	6	7	8	9	10	11	12
	ab	ac	ad	ba	bc	bd	ca	cb	cd	da	db	dc
1 (a→b)	1	0	0	0	0	0	0	0	0	0	0	0
2 (a→c)	0	1	1	0	0	1	0	0	0	0	0	0
3 (b→a)	0	0	0	1	0	1	1	0	0	1	0	0
4 (b→c)	0	0	0	0	1	0	0	0	0	0	0	0
5 (c→b)	0	0	0	0	0	0	1	1	0	1	1	0
6 (c→d)	0	0	1	0	0	1	0	0	1	0	0	0
7 (d→c)	0	0	0	0	0	0	0	0	0	1	1	1

Source-Destination Pairs

$$Y_{rx1} = A_{rxp} * X_{px1} \quad \text{with } p >> r$$

A = Routing matrix.

X = Source – Destination transmission vector **[Unknown]**.

Y = Link traffic.

Origin-Destination Literature

- **Source-Destination Traffic Estimation.**
  [Vardi, J. of the Amer. Statist. Assoc., 1996].
- Bayesian Inference on Network Traffic Using Link Count Data.
  [Tabaldi and West, J. of the Amer. Statist. Assoc., 1998].
- Time-Varying Network Tomography.
  [Cao et al., J. of the Amer. Statist. Assoc., 2000].
- Traffic Matrix Estimation: Existing Techniques and New Directions.
  [Medina et al., ACM SigComm 2002].
- An Information-Theoretic Approach to Traffic Matrix Estimation.
  [Zhang et al., ACM SigComm 2003].

Link Perf. Inference Literature

- **Multicast-based Inference of Network-internal Characteristics (MINC Project).**
  [Caceres, Duffield, LoPresti, Horowitz, Kurose, Towsley, Paxson].
- Network Loss Inference using Unicast End-to-End Measurement.
  [Coates and Nowak, ITC Seminar on IP Traffic, Measurement and Modelling , 2000].
- Unicast inference of network link delay distributions from edge measurements
  [Shih and Hero, IEEE Int. Conf. on Acoust. Speech and Sig. Proc., 2001].
- Nonparametric Internet Tomography.
  [Tsang, Coates, and Nowak, IEEE Intl. Conf. on Acc., Speech and Signal Proc., 2002].
- Simple Network Performance Tomography.
  [Nick Duffield, ACM IMC 2003].
- Tomography-based Overlay Network Monitoring.
  [Chen, Bindel, Katz, ACM IMC 2003].

Topology Inference Literature

- Multicast Topology Inference from Measured End-to-End Loss.
  [Duffield et al., IEEE Trans. on Info. Theory, 2002]
- **Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements.**
  [Coates et al., ACM Sigmetrics, 2002].
- Merging Logical Topologies Using End-to-end Measurements.
  [Coates et al., ACM IMC, 2003].

Outline

- Origin-Destination.

  Source-Destination Traffic Estimation.

  [Vardi, J. of the American Statistical Association, 1996].

- Link-Level Network Inference.

  Multicast-Based Inference of Network-Internal Loss Characteristics.

  [Cáceres, Duffield, Horowitz, Towsley, IEEE Trans. In Information Theory, 1999].

- Topology Inference.

  Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements.

  [Coates et al., ACM Sigmetrics, 2002].

Outline

- **Origin-Destination.**

  Source-Destination Traffic Estimation.

  [Vardi, J. of the American Statistical Association, 1996].

- Link-Level Network Inference.

  Multicast-Based Inference of Network-Internal Loss Characteristics.

  [Cáceres, Duffield, Horowitz, Towsley, IEEE Trans. In Information Theory, 1999].

- Topology Inference.

  Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements.

  [Coates et al., ACM Sigmetrics, 2002].

Source-Destination Traffic Estimation

- ***Goal:*** Given Y_{rx1}, estimate X_{px1} s.t.

$$Y_{rx1} = A_{rxp} * X_{px1}$$

  with $r << p$.

- ***Idea:***
  - Use a model for X_{px1}, eg Poisson with rates **λ.**
  - Estimate the parameters of the model (**λ**) that maximize the probability of observing Y_{rx1}.

- ***Other Approaches:*** Direct measurements with NetFlow, passive monitoring, etc.

A Simple Example.

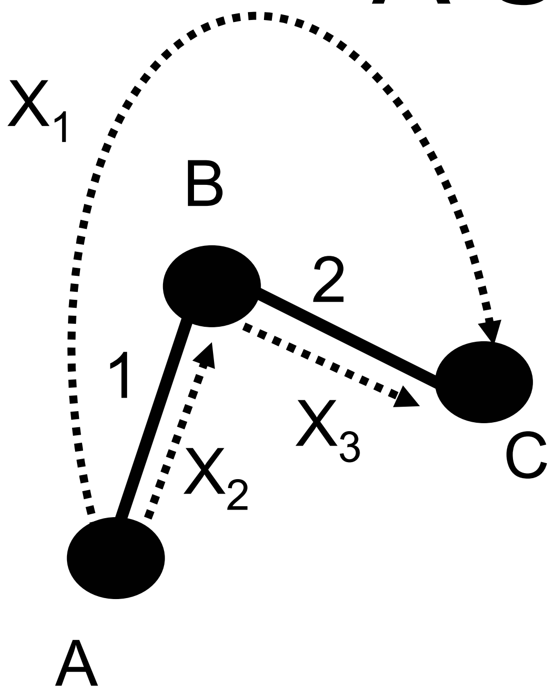

$$\mathbf{X} = (X_1, X_2, X_3)^T$$

$$\mathbf{Y} = (Y_1, Y_2)^T = (1, 2)^T$$

$$\mathbf{Y} = \begin{pmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \cdot \mathbf{X}$$

Model: $X_i \sim \text{Poisson}(\lambda_i)$

1) Find possible $\mathbf{X}$: Very expensive!

$$\mathbf{X} = (1,0,1)' \text{ or } \mathbf{X} = (0,1,2)'$$

2) Find likelihood of $\mathbf{Y}$:

$$L(\lambda) = \left(\lambda_1\lambda_3 + \lambda_2\lambda_3^2/2\right) \exp\left(-\lambda_1 - \lambda_2 - \lambda_3\right)$$

3) Find λ: $\max_\lambda L(\lambda)$. Maybe corner solution!

Expectation-Maximization (EM)

Likelihood maximization when:

$$\lambda = E_\lambda \left[\mathbf{X} | \mathbf{Y} = A\mathbf{X} \right]$$

Algorithm for finding $\boldsymbol{\lambda}$:

1. Pick initial $\boldsymbol{\lambda}^{(0)}$.
2. $\boldsymbol{\lambda}^{(n+1)}$ = E[**X** | **Y**, $\boldsymbol{\lambda}^{(n)}$].

Problems:

1. Difficult to evaluate E[**X** | **Y**, $\boldsymbol{\lambda}$].
2. May converge to non-MLE point.

Normal Approximation

1. Measure **Y** in K periods.
2. Approximate:

   $\mathbf{Y} \sim N_r(\mathbf{A\lambda}, K^{-1}\, \mathbf{A\Lambda A'})$ with $\mathbf{\Lambda} = \mathrm{diag}(\boldsymbol{\lambda})$
3. Compute sample average $\bar{\mathbf{Y}}$ and sample covariance **S**.
4. Equate sample moments to theoretical moments:
   - $\bar{\mathbf{Y}} = \mathbf{A\lambda}$
   - $\mathbf{S} = K^{-1}\, \mathbf{A\Lambda A'}$

Outline

- Origin-Destination.

  Source-Destination Traffic Estimation.

  [Vardi, J. of the American Statistical Association, 1996].

- **Link-Level Network Inference.**

  Multicast-Based Inference of Network-Internal Loss Characteristics.

  [Cáceres, Duffield, Horowitz, Towsley, IEEE Trans. In Information Theory, 1999].

- Topology Inference.

  Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements.

  [Coates et al., ACM Sigmetrics, 2002].

Link-Level Network Inference.

- ***Goal:*** Infer network link characteristics, like loss rate, delay distribution, etc.

- ***Idea***:
  - Collect end-to-end measurements.
  - Assume a) known topology, b) ***model*** for network behavior.
  - Identify *network parameters* that maximize the probability of the observed measurements.

- ***Other Approaches:*** pathchar, traceroute, clink, etc.

Model

Logical Multicast Tree

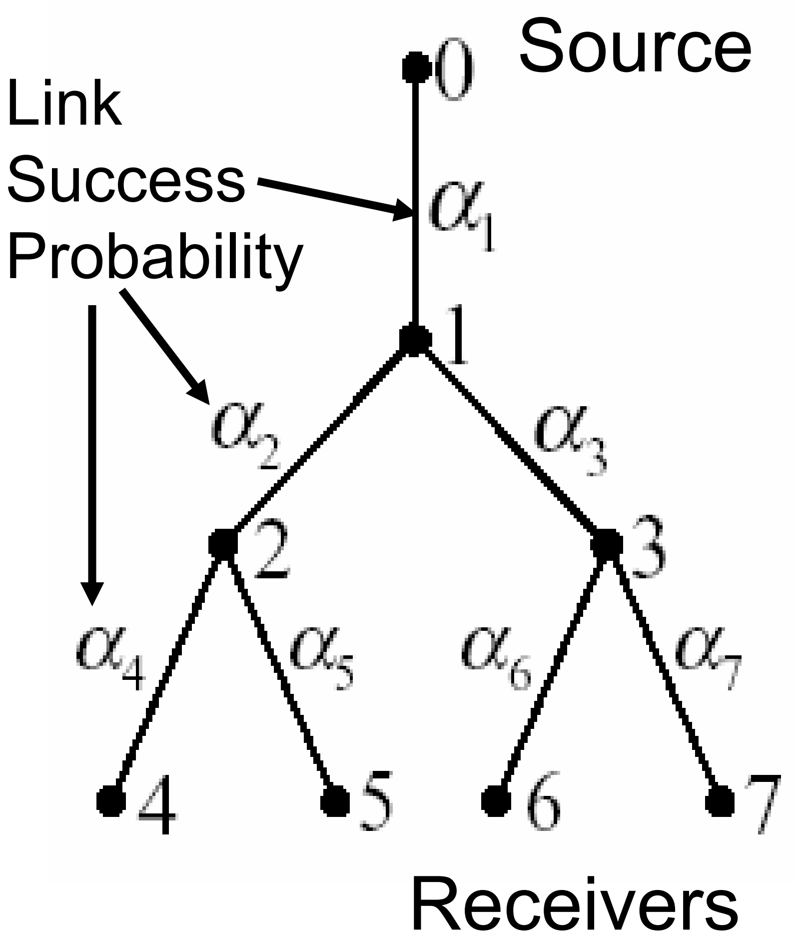

- Bernoulli losses with probabilities **α**:
  - Temporal dependence $\Rightarrow$ Slow convergence.
  - Spatial dependence $\Rightarrow$ Error proportional to dependence.

Model

Logical Multicast Tree

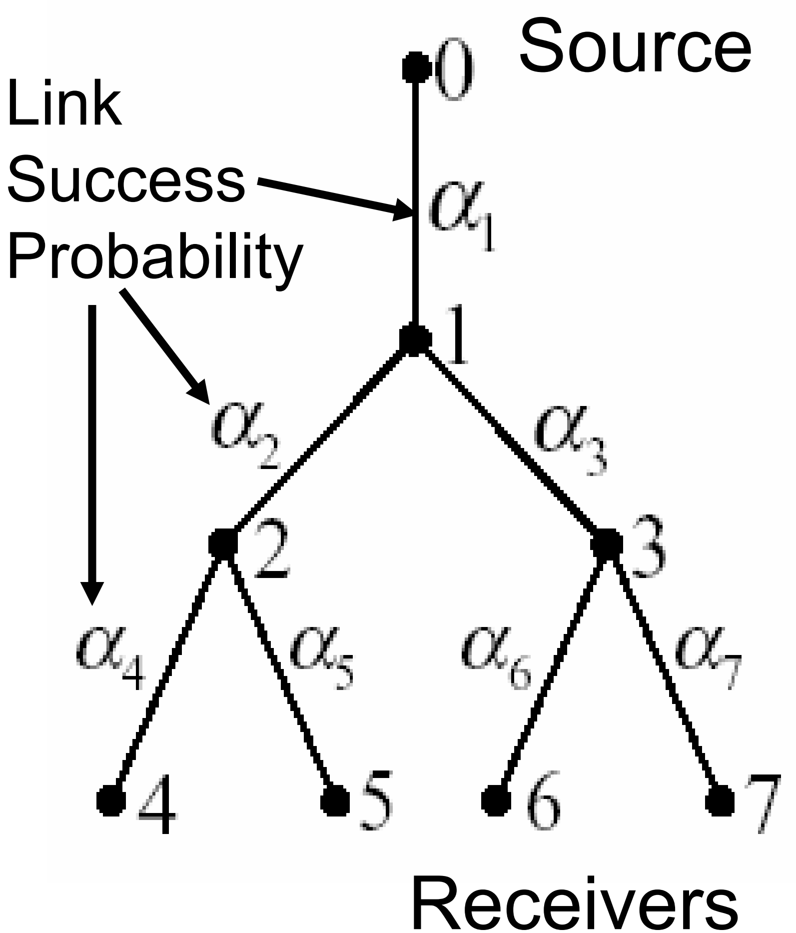

- **Ω** = set of outcomes, i.e. subsets of receivers received a probe packet.
- **n**: number of probes.
- **n(x)**: # of probes with outcome **x∈Ω**.

Find **α** to maximize:

$$p\left(x^1,\ldots,x^n;\alpha\right)=\prod_{x\in\Omega}p\left(x;\alpha\right)^{n(x)}$$

Solution Methodology

Compute **α** to maximize:

$$p\left(x^1, \ldots, x^n; \alpha\right) = \prod_{x \in \Omega} p\left(x; \alpha\right)^{n(x)}$$

Using Maximum Likelihood Estimators.

Properties:

- Strong consistency.
- Asymptotic normality.
- Asymptotic unbiasedness.
- Asymptotic efficiency.

Convergence of Inferred Loss Probabilities

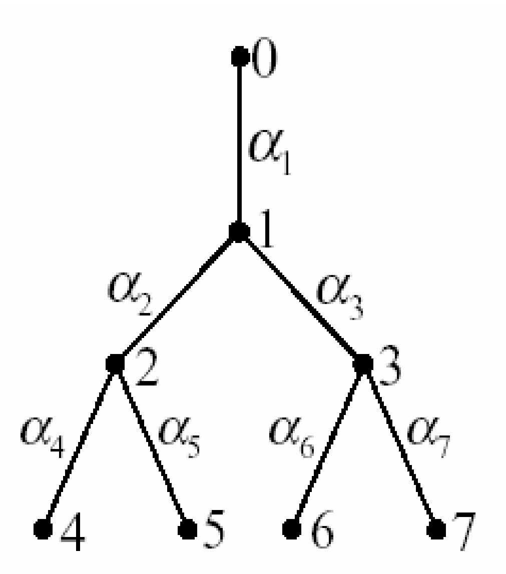

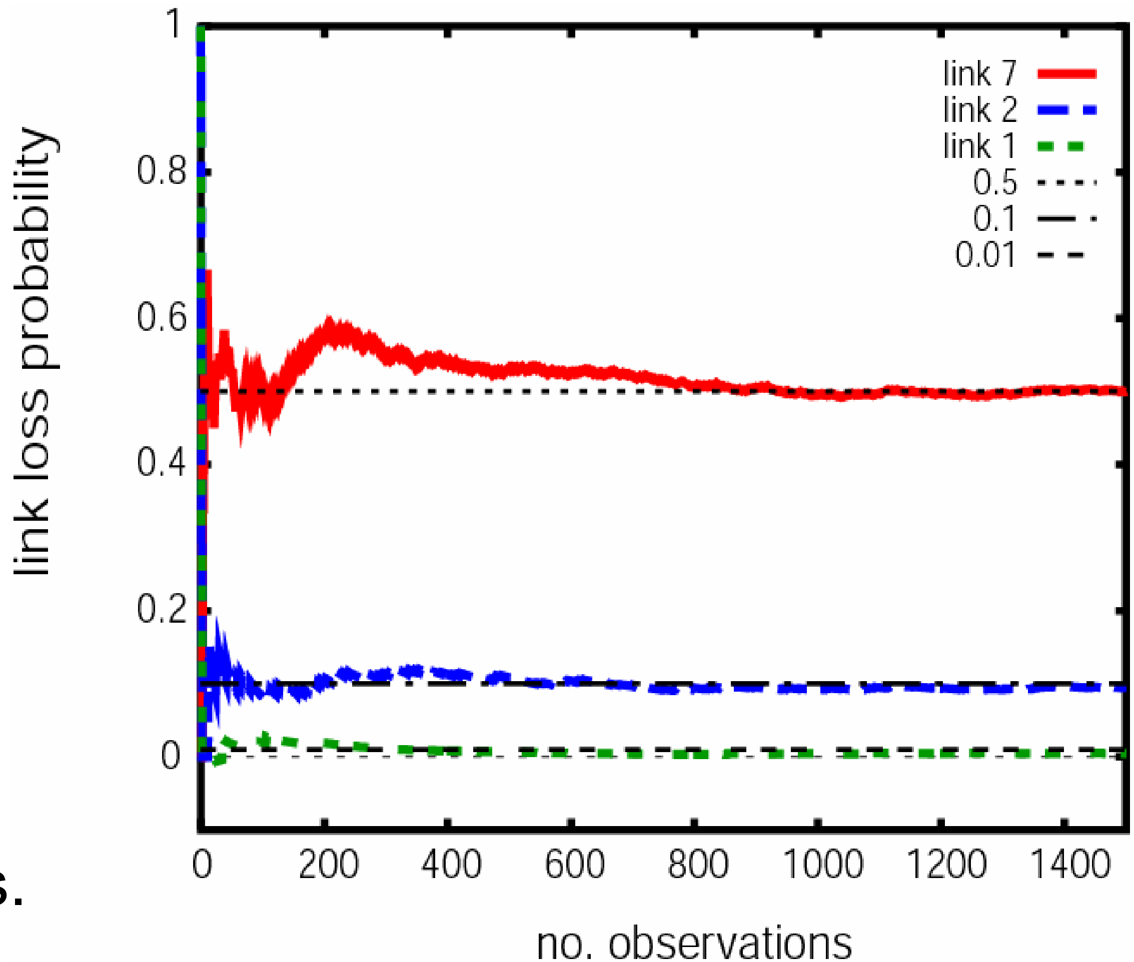

Error ≤0.001 after 2000 observations.

Effect of Topology

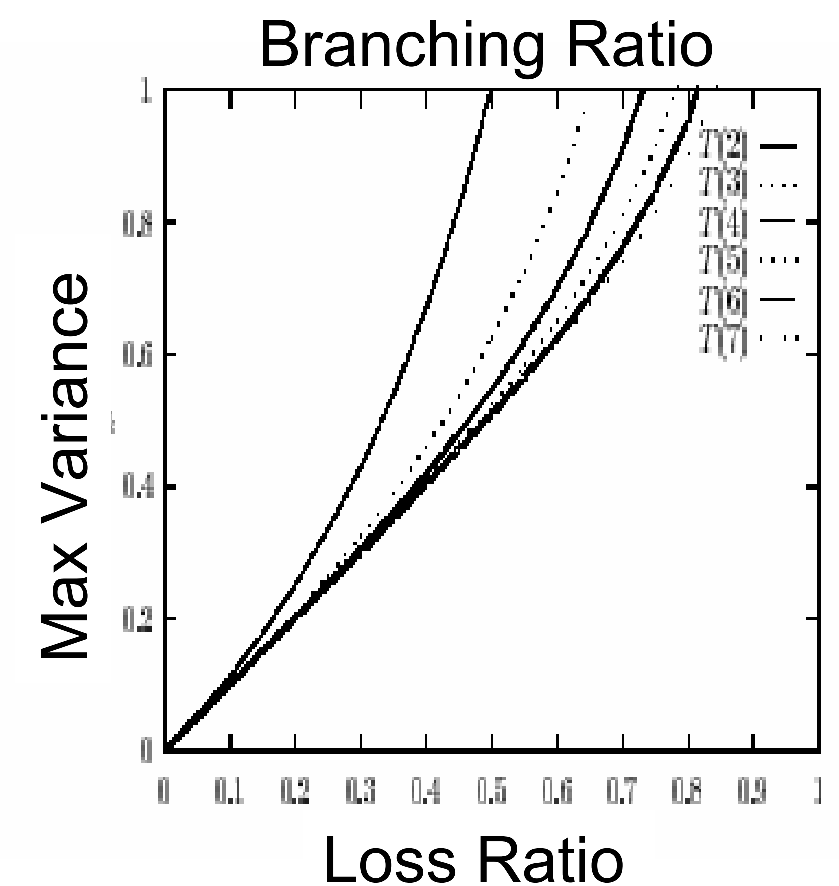

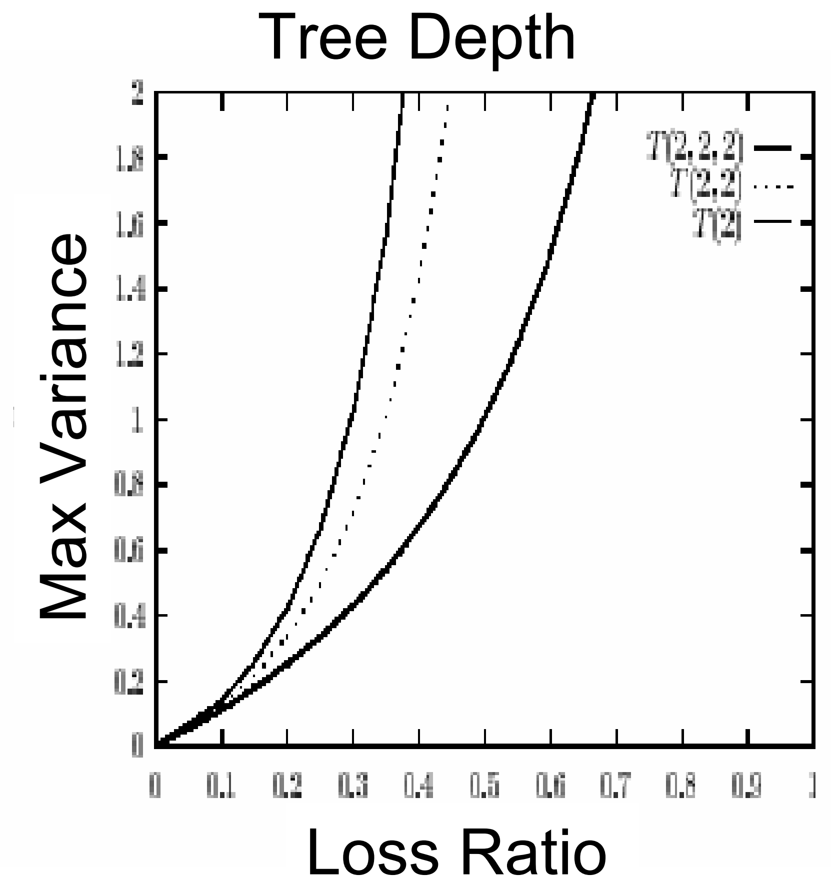

Variance *decreases* with branching ratio.

Variance *increases* with tree depth.

Outline

- Origin-Destination.

  Source-Destination Traffic Estimation.

  [Vardi, J. of the American Statistical Association, 1996].

- Link-Level Network Inference.

  Multicast-Based Inference of Network-Internal Loss Characteristics.

  [Cáceres, Duffield, Horowitz, Towsley, IEEE Trans. In Information Theory, 1999].

- **Topology Inference.**

  Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements.

  [Coates et al., ACM Sigmetrics, 2002].

Topology Inference

- ***Goal:*** Identify the *tree* topology connecting a *single* server to *multiple* receivers.

- ***Idea:***
  – Use a monotonic, increasing function of the number of *shared* links to two receivers, e.g. delay, loss, etc.
  – MLE for topology identification:

$$\mathcal{T}^* = \text{argmax}_{\mathcal{T}\in\mathcal{F}} \text{max}_{\gamma\in\mathcal{G}} \left[p\left(\mathbf{x}|\gamma,\mathcal{T}\right)\right)\right]$$

- ***Other Approaches:*** traceroute, AS map, mtrace, etc.
  – Topology inference more expensive.
  – But, works without router support.
  – “Can” identify **Layer-2** devices.

Sandwich Probe

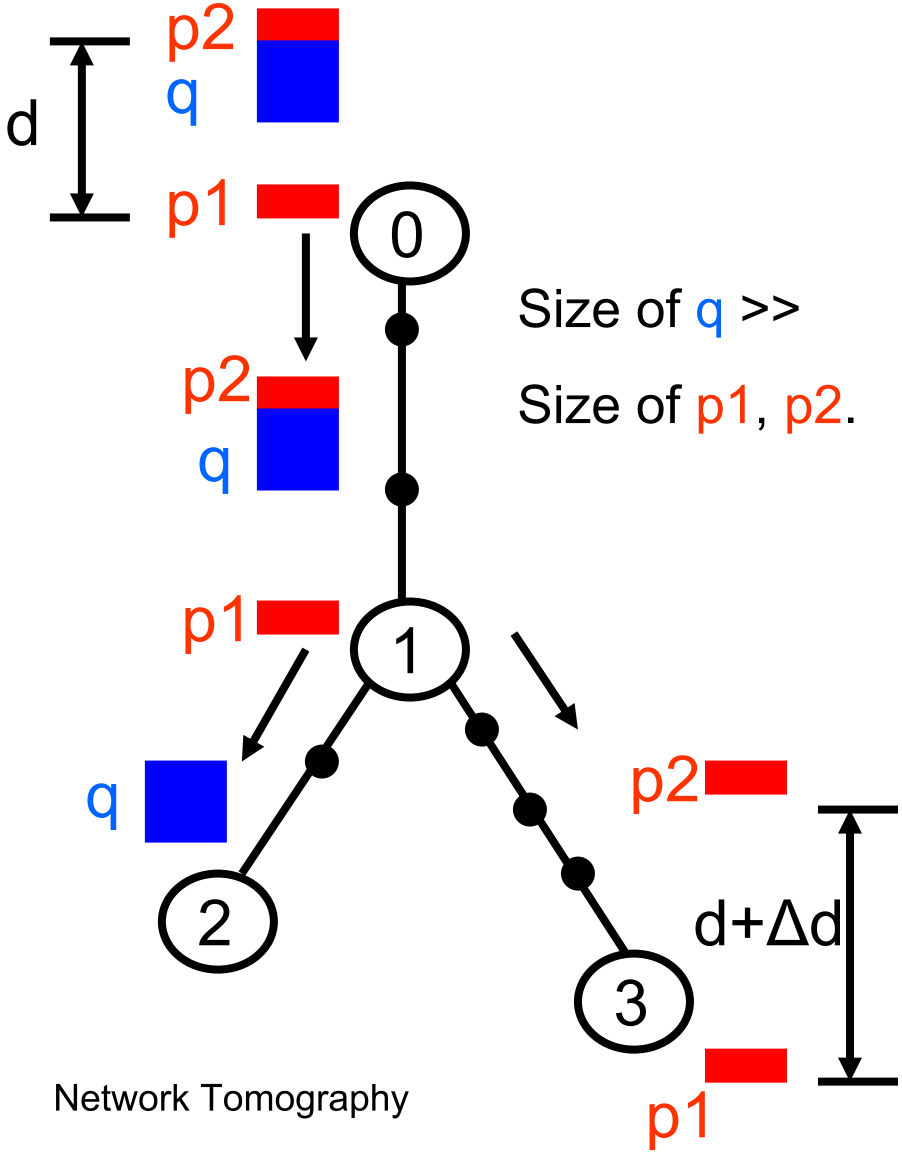

Idea:

- Packet q introduces delay Δd between p1and p2.
- Δd ∝ shared path 0→1.
- Node 3 measures **Δd**.

Advantages:

- Every measurement is important ⇒ Fast.
- No clock synchronization.

Measurement Collection

For each **ordered** pair of nodes (i,j):

- K measurements of delay difference:

$$\Delta d_{i,j}^{k} \text{ for } k = 1, \ldots, K$$

- Compute sample mean and sample variance:

$$x_{i,j} = \frac{1}{K} \sum_{k=1}^{K} \Delta_{i,j}^{k}, \qquad \sigma_{i,j}^{2} = \frac{1}{n} \sum_{k=1}^{K} \left(\Delta_{i,j}^{k} - x_{i,j} \right)^{2}$$

- Asymptotically $x_{i,j}$ is Normal.

Maximum Likelihood Topology Identification.

Find $\mathcal{T}*$:

$$\mathcal{T}^* = \text{argmax}_{\mathcal{T}\in\mathcal{F}}\left[p\left(\mathbf{x}|\mathcal{T}\right)\right)]$$

Assume delay at link l of $\mathcal{T}$ is μ_l. Maximize:

$$L\left(x, \mathcal{T}\right) \equiv \log p\left(x|\mathcal{T}, \widehat{\mu}(\mathcal{T})\right)$$

Prior model:

$$x_{i,j} \sim \mathcal{N}\left(\sum_{l\in S_{i,j}} \mu_l, \sigma_{i,j}^2\right)$$

Penalize trees with many links:

$$L_\lambda\left(x, \mathcal{T}\right) \equiv \log p\left(x|\mathcal{T}, \widehat{\mu}(\mathcal{T})\right) - \lambda n\left(\mathcal{T}\right)$$

with λ a user-defined parameter.

Finding the best tree

Optimize:

$$L_{\lambda}(x, \mathcal{T}) \equiv \log p(x|\mathcal{T}, \widehat{\mu}(\mathcal{T})) - \lambda n(\mathcal{T})$$

Number of trees $\mathcal{T}$ is HUGE.

Instead of exhaustive search for $\mathcal{T}$, use a reversible jump Markov Chain Monte Carlo:

1. Start at state $(\mathcal{T}_0, \mu_0)$.
2. Propose a move to a new state $\left(\mathcal{T}_i, \mu_{i+1}\right)$.
3. Accept the proposal with probability proportional to the ratio of the likelihoods of the two states.
4. Continue and store for each state visited, its likelihood.

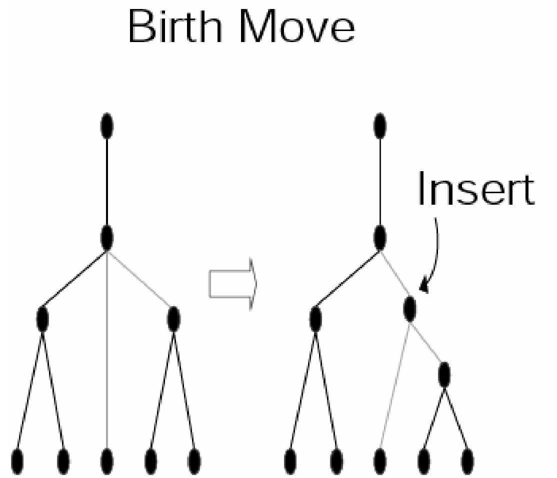

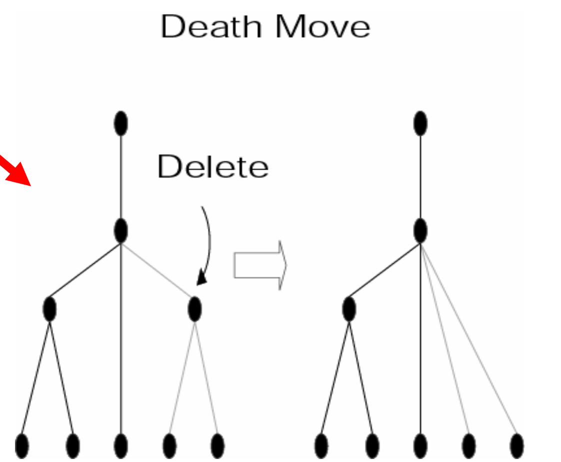

Topology Inference in ns-2

Link Capacity

Access capacity low

Leaf nodes -
Receivers

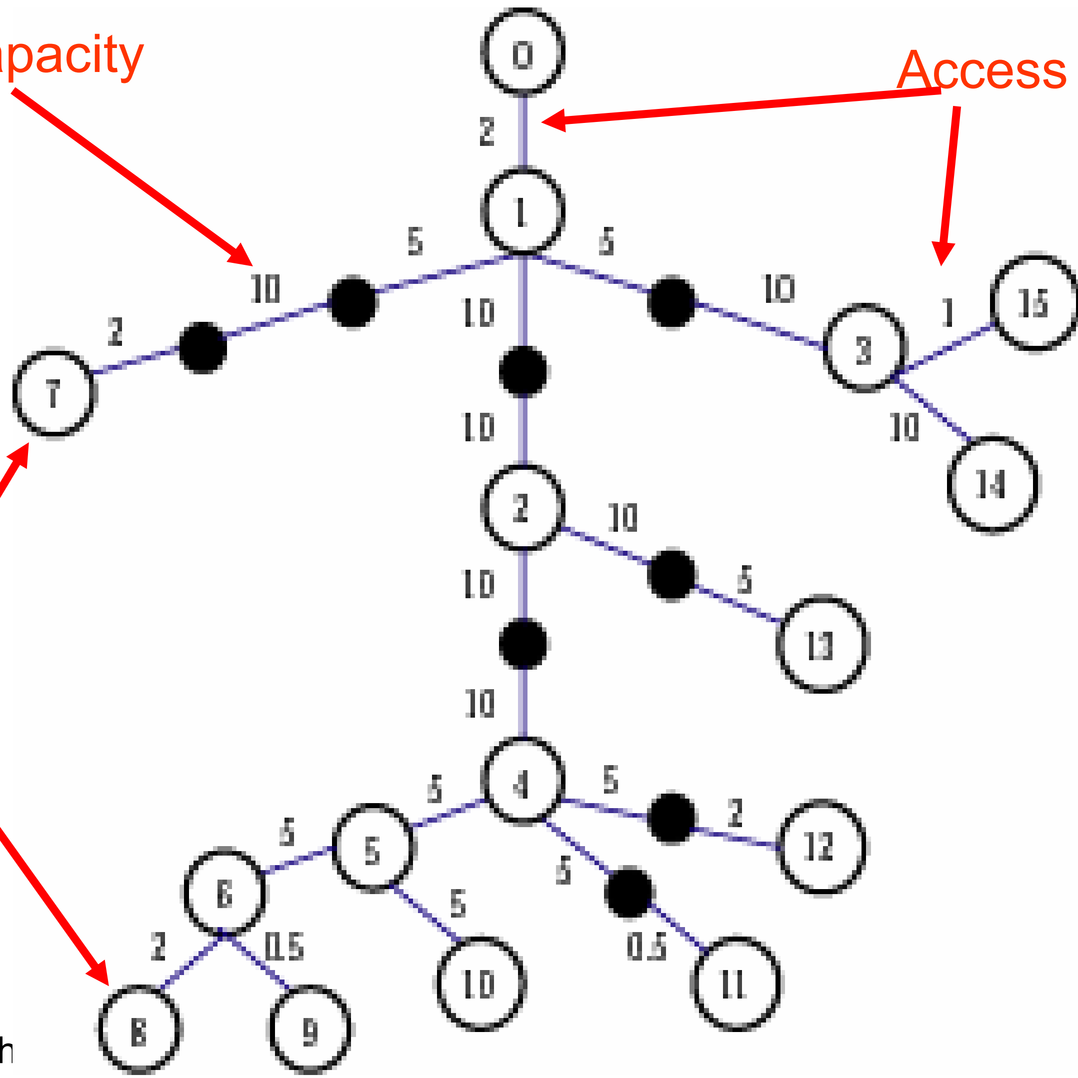

Topology Inference in ns-2

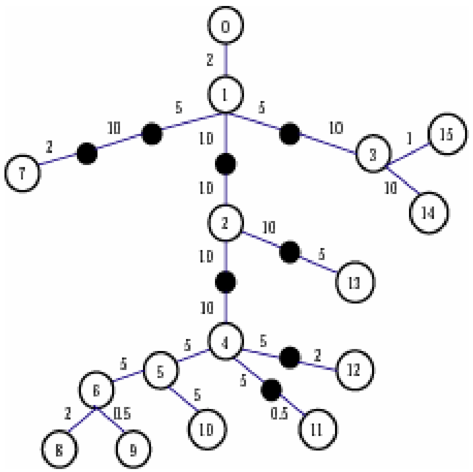

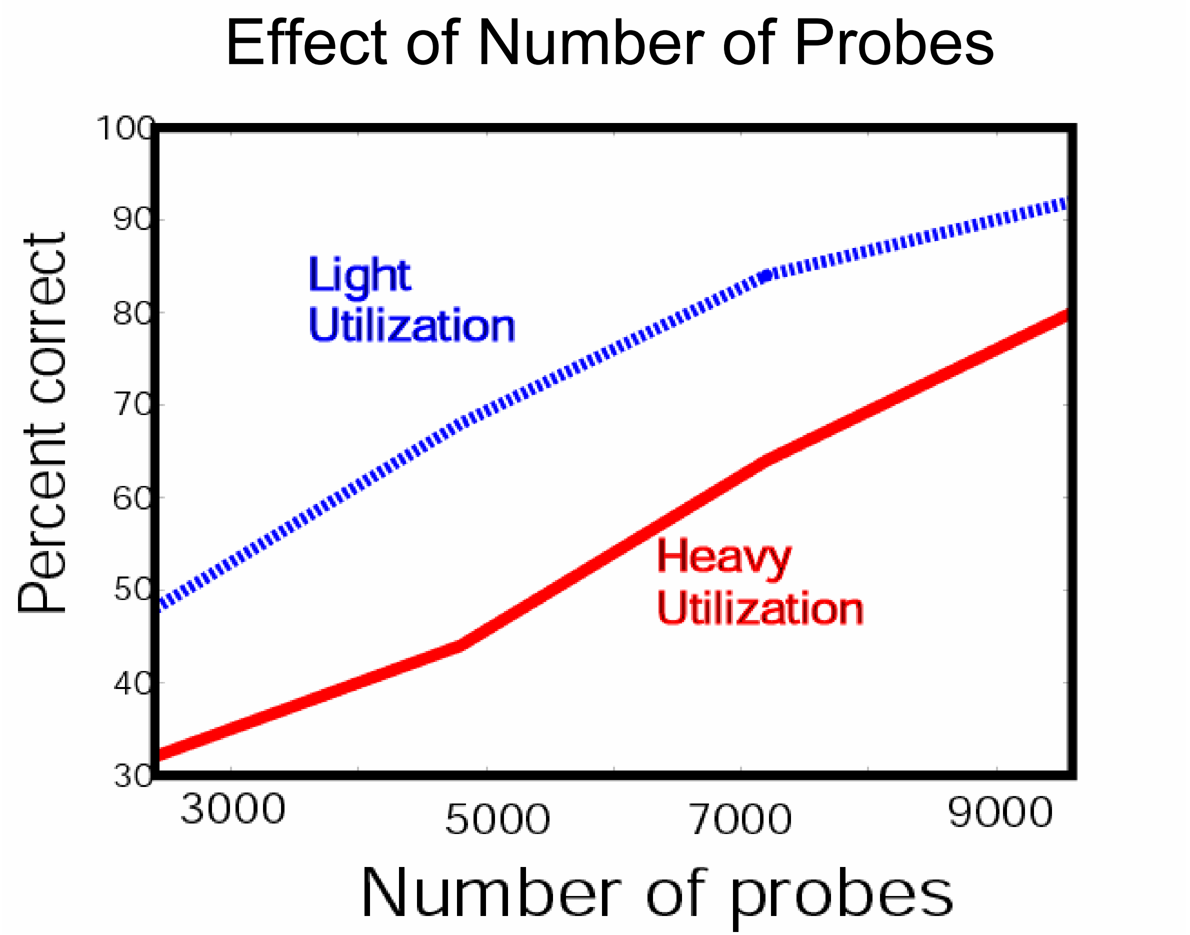

Topology Inference in ns-2

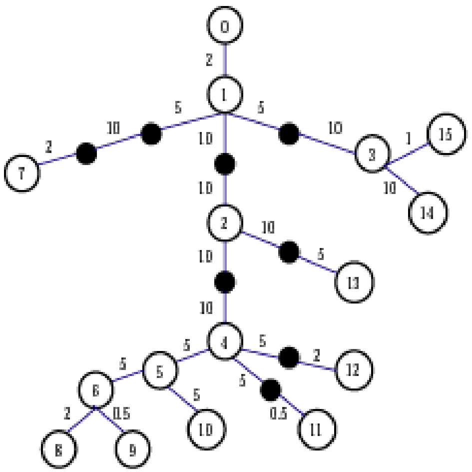

Effect of Penalty λ

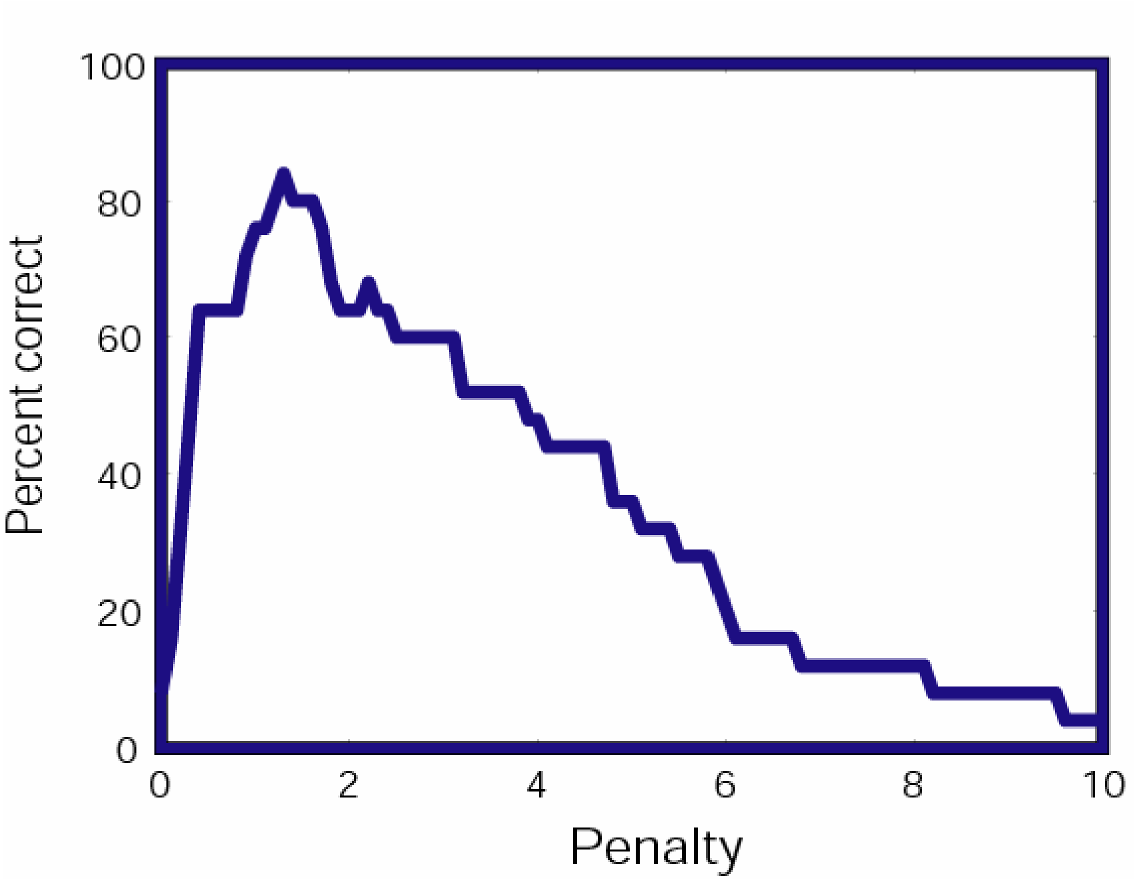

Topology Inference in the Internet

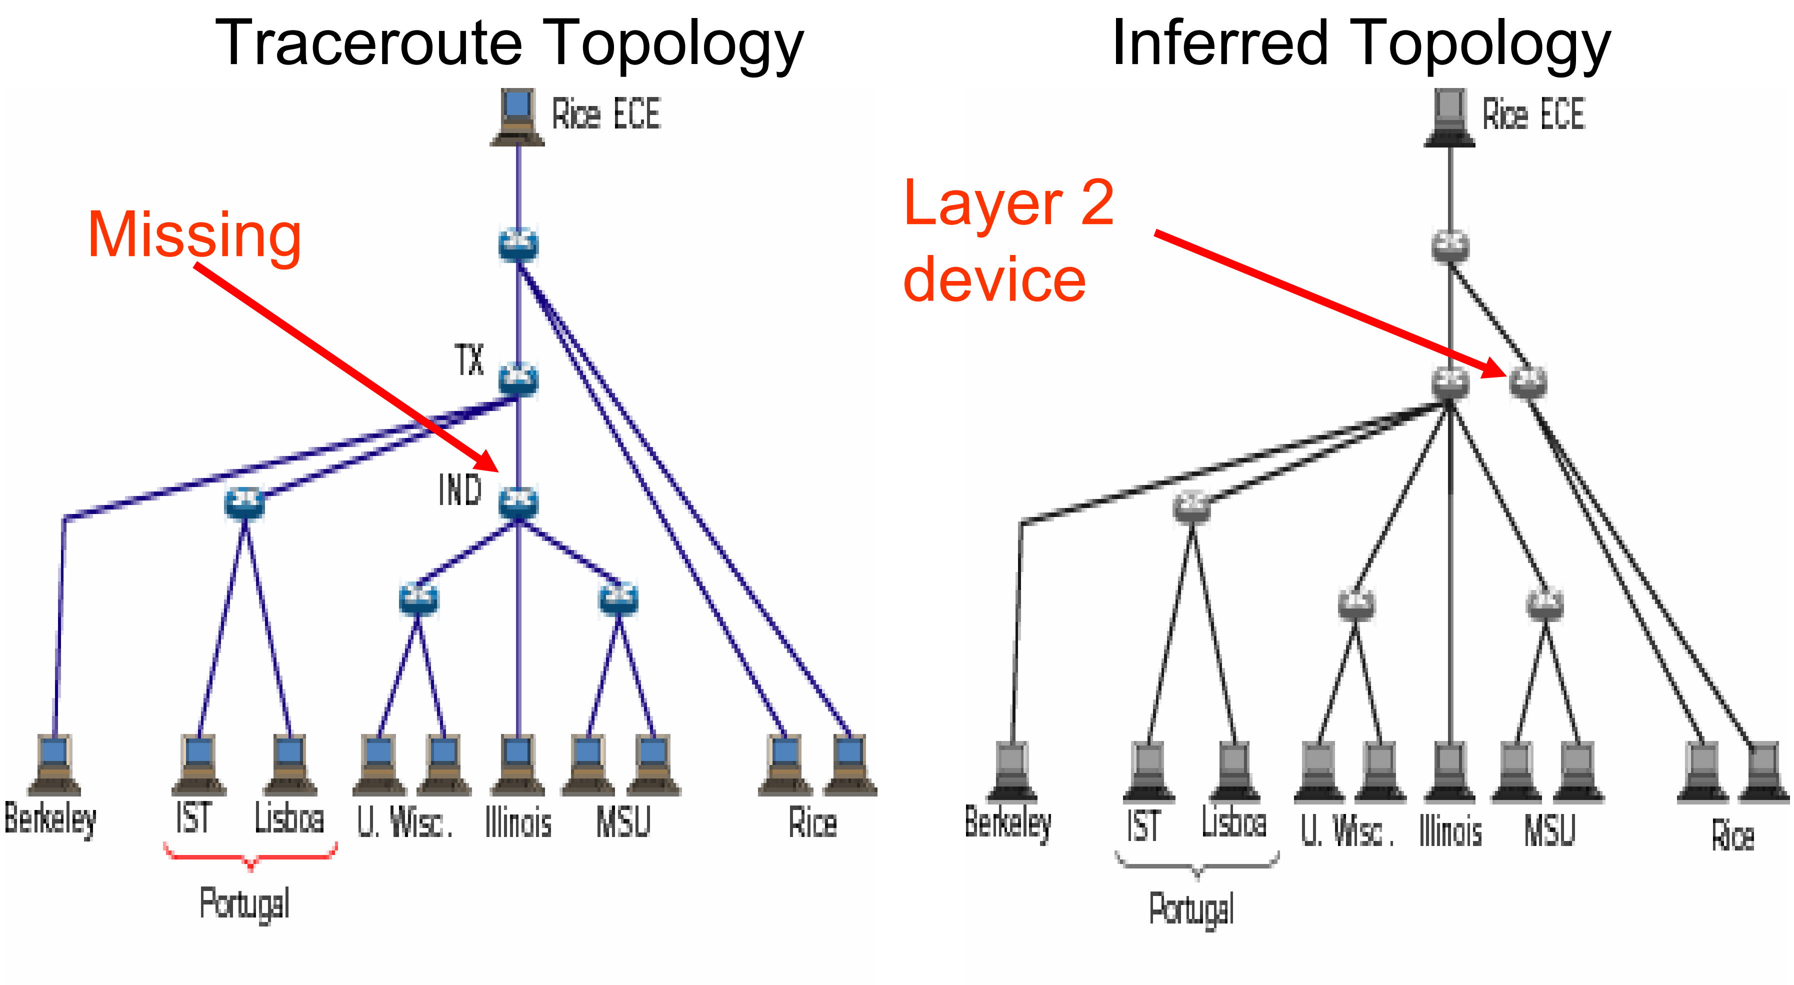

Summary

- Areas of interest:
  - Origin-Destination Traffic Matrix.
  - Link-Level Network Inference.
  - Topology Inference.
- Pick solution with **highest likelihood** according to a predefined **model**.
- Numerically difficult problems.
- Standard tools: MLE, Bayesian Inference, EM, MCMC.

Future Work

- Spatial and temporal dependencies.
- Time-varying, non-stationary OD traffic matrices.
- Traffic models with long range dependencies.
- Identification of anomalous behavior, instead of detailed statistics.
- Better passive traffic monitoring.