Flexicast Delay Network Tomography

George Michailidis
Department of Statistics
The University of Michigan

Joint work with Earl Lawrence (LANL)
and Vijay Nair (UM)

Roadmap

- Background and Motivation.
- Modeling Considerations.
- Estimation.
- VoIP Testing.

Network Modeling Goals

- Assess network performance.
- Monitor the network for anomalous behavior.
- Capacity planning.
- Improve routing protocols.
- Do all of the above quickly.

Challenges

- Networks are decentralized and multilayered.
- Restricted access to measurements.
- Many users with different quality of service requirements.

Important quality of service metrics due to queuing: Loss and Delay

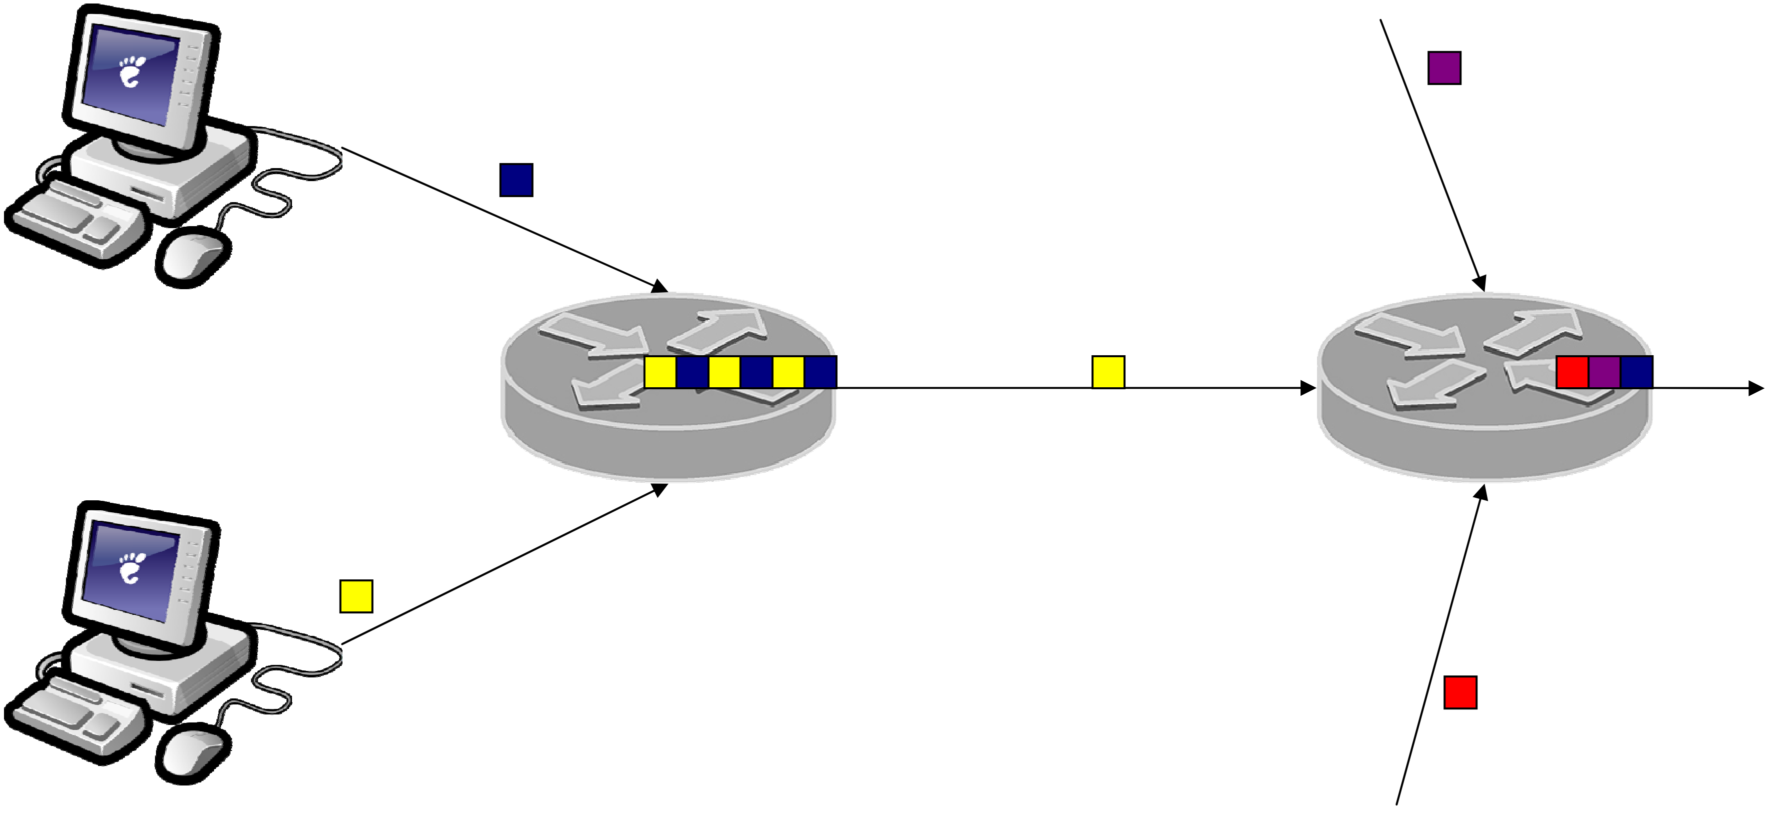

Active Network Tomography

Probe the network with test traffic and measure the end-to-end (path) characteristics. Use these data to obtain information about link characteristics.

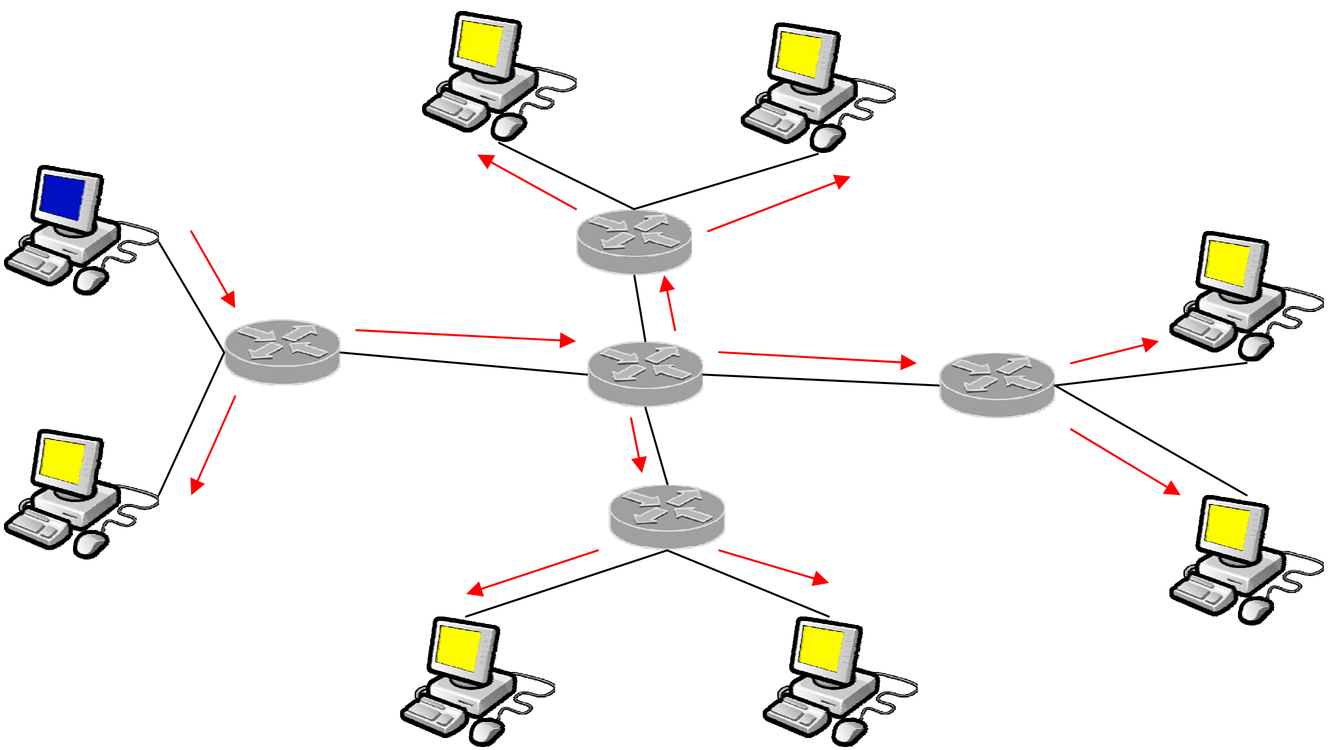

Tree Topology

- Induced by single source probing.
- Logical topology
  - Only includes branching points.

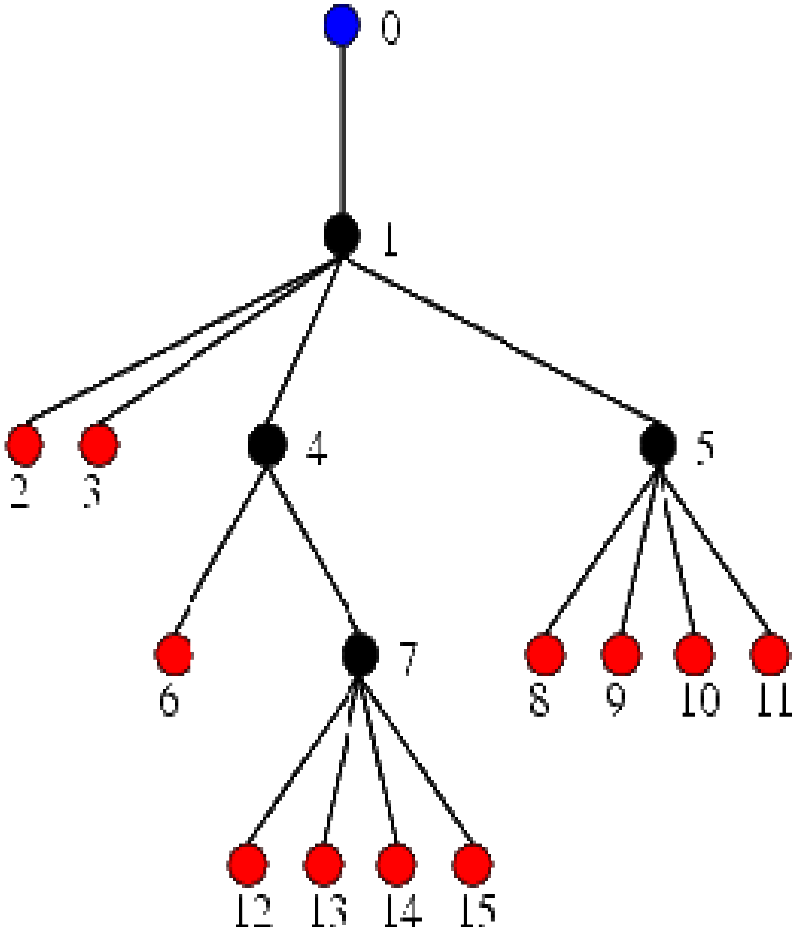

Stochastic Model Choices

- Choice depends on analysis goals
  - Planning
  - Monitoring
  - Detailed Analysis
- Types of models
  - Non-parametric models
  - Parametric models
  - Semi-parametric models
  - Mixture models

Stochastic Model Choice

- General assumption:

  Spatio-temporal independence.

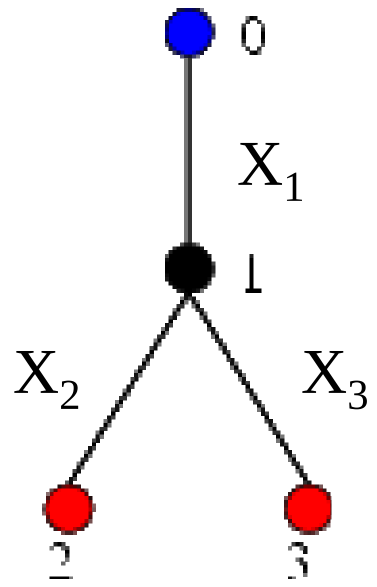

$Y_2 = X_1 + X_2$ $\quad$ $Y_3 = X_1 + X_3$

Probing

- Unicast: Commonly used but has identifiability issues.
- Multicast: Identifiable but intrusive and hard to consider.
- Flexicast: Identifiable and allows intelligent experimentation.

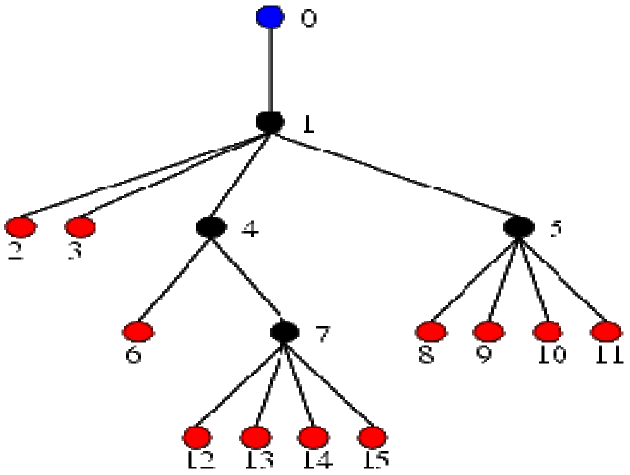

Important Issue: Identifiability of Model Parameters

Probing condition:

1. Every receiver must be covered.
2. Every internal node must be used as a branching point.

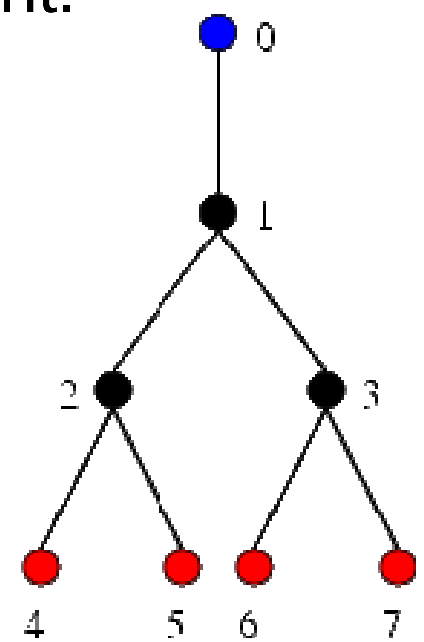

Discrete Delay Model

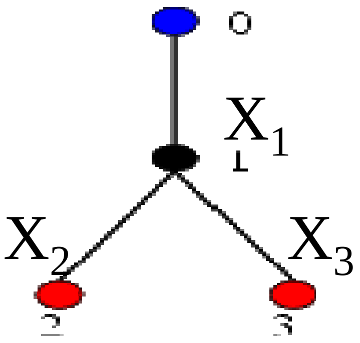

- $X_k \in \{0, q, 2q, \ldots, bq\}$
  $P\{X_k = iq\} = \alpha_k(i)$
- Flexible regarding shape.
- Probing condition is necessary
  and sufficient for identifiability.

- Studied in Liang and Yu (2003), Tsang et al. (2003), Lawrence et al. (2005)

Discrete Delay Model: Estimation through the EM

$$l(\vec{\alpha};\mathbf{Y}) = \sum_{c\in\mathcal{C}} \sum_{\vec{y}\in\mathcal{Y}^c} N^c_{\vec{y}} \log[\sum_{\vec{x}\in\mathcal{X}^c(\vec{y})} \mathrm{P}(\vec{\alpha})\{X^c = \vec{x}\}]$$

$$\mathrm{E-Step}: \quad M_i^{k\ (q+1)} = \sum_{c\in\mathcal{C}:k\in\mathcal{T}^c} \sum_{\vec{x}\in\mathcal{X}^c:x_k=i} \frac{\mathrm{P}\{\vec{X}^c=\vec{x}\}^{(q)}}{\mathrm{P}\{\vec{Y}^c=\vec{y}(\vec{x})\}^{(q)}} N^c_{\vec{y}}$$

$$\mathrm{M-Step}: \quad \alpha_k(i)^{(q+1)} = \frac{1}{m_k} M_i^{k\ (q+1)}$$

Discrete Delay Model

- EM algorithm computationally expensive
  - Scales linearly with number of bins
  - Scales exponentially with tree size
  - Batch processing mode, especially in a multi-source topology

UNC Probing Experiments

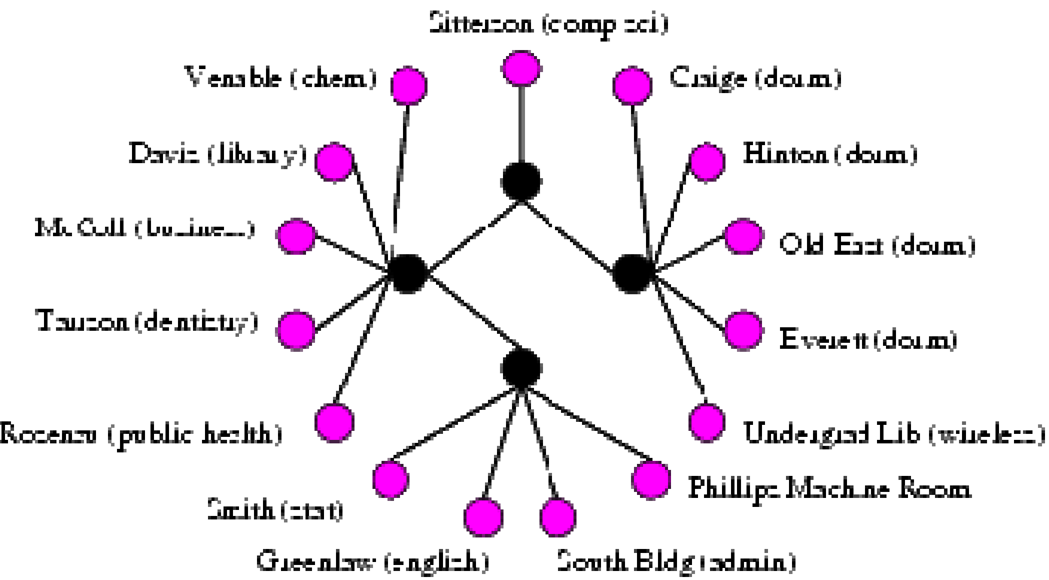

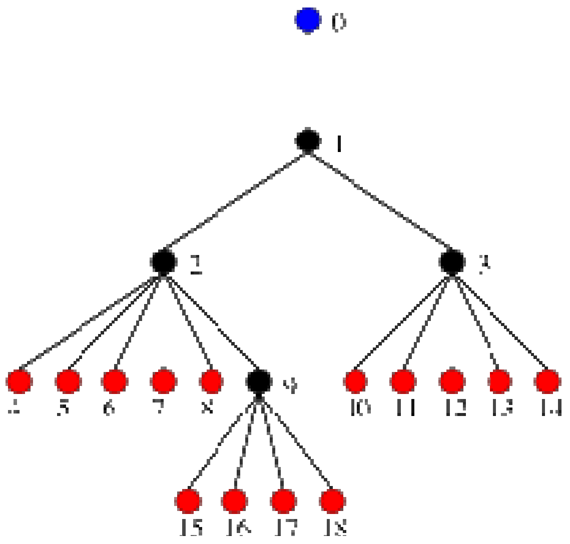

<4,5>, <6,7>, <8,10>, <11,12>, <13,14>, <15,16>, <17,18>

Large Delay Results

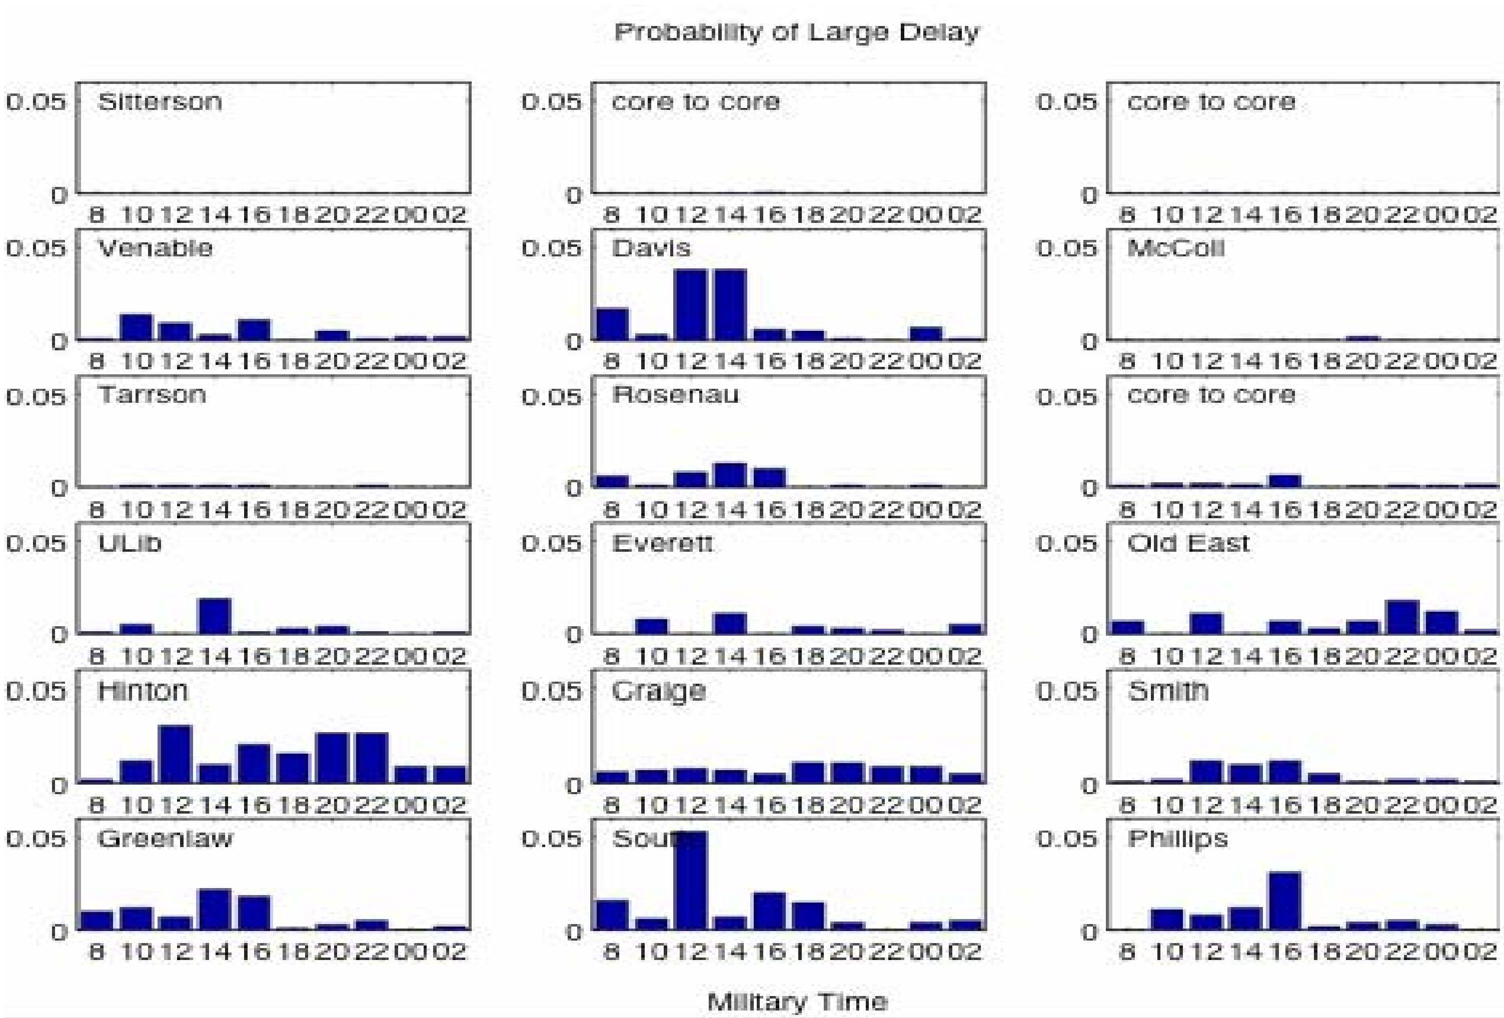

Grafting: Local MLE Combinations through a fixed point algorithm

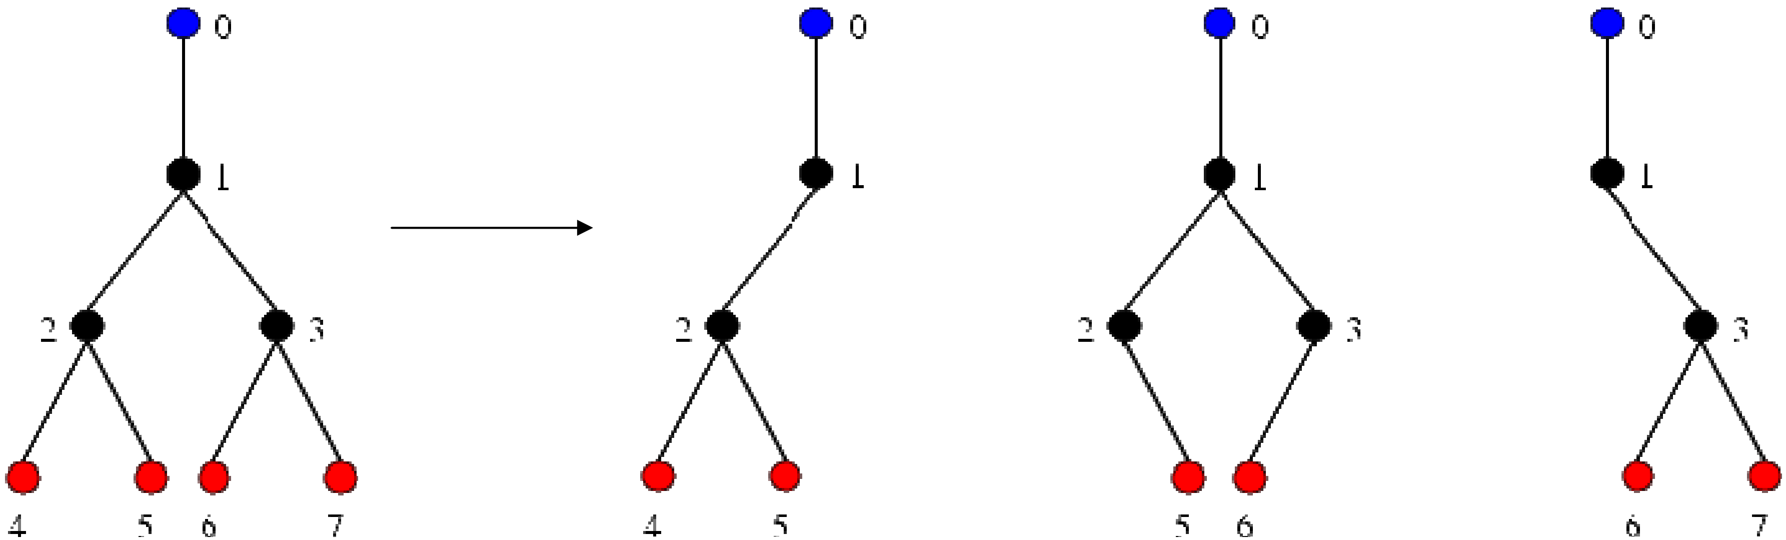

<4,5>	$\vec{\pi}_{0,2}$	$\vec{\alpha}_1$	$\vec{\pi}_{0,3}$
<5,6>	$\vec{\alpha}_4$	$\vec{\pi}_{1,5}$	$\vec{\alpha}_6$
<6,7>	$\vec{\alpha}_5$	$\vec{\pi}_{1,6}$	$\vec{\alpha}_7$

Parametric Models

- Delays are continuous.
- Nonparametric models are computationally expensive; hence, not appropriate for monitoring purposes
- Almost always parametric assumptions are not valid.
- Identifiability is an issue.

$$X_k \sim \mathrm{Exp}(\theta_k)$$

$$\mathrm{E}(\vec{Y}) = \begin{pmatrix} \theta_1 + \theta_2 \\ \theta_1 + \theta_3 \end{pmatrix}$$

$$\mathrm{Cov}(\vec{Y}) = \begin{bmatrix} \theta_1^2 + \theta_2^2 & \theta_1^2 \\ \theta_1^2 & \theta_1^2 + \theta_3^2 \end{bmatrix}$$

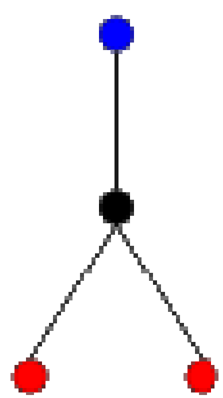

$$X_k \sim \mathrm{N}(\mu_k, v_k)$$

$$\mathrm{E}(\vec{Y}) = \begin{pmatrix} \mu_1 + \mu_2 \\ \mu_1 + \mu_3 \end{pmatrix}$$

$$\mathrm{Cov}(\vec{Y}) = \begin{bmatrix} v_1 + v_2 & v_1 \\ v_1 & v_1 + v_3 \end{bmatrix}$$

Functional Moments Model

- $X_k \sim F_k(x\,;\,\theta_k)$
- If θ_k is estimable from central order moments of order two and higher, then the probing condition is sufficient for identifiability.
- Includes exponential, gamma, log-normal, Weibull, and others.
- Maximum likelihood estimation almost intractable for all but toy topologies.
- Studied in Lawrence et al. (2006).

Extension: Point Mass Model

- $X_k \sim p_k\delta(0) + (1-p_k)F_k(x)$
- For $p_k > 0$, the probing condition is sufficient for identifiability.

(y,y) $y \sim F_1$

(y,0) $y \sim F_2$

(0,y) $y \sim F_3$

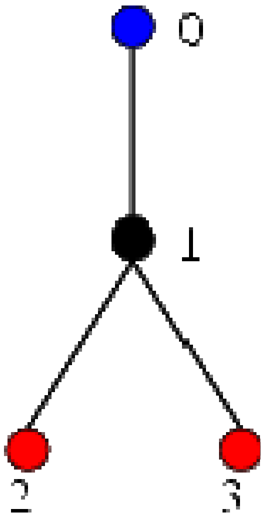

Semi-parametric Modeling

- Specify desired moments:

  $E(X_k) = \mu_k$

  $Var(X_k) = \phi \mu_k^{\gamma}$
- Probing condition is sufficient.
- Simple structure is good for monitoring purposes and some capacity planning (Lawrence et al. 2005).

Moment Estimation

- Match observed moments by minimizing least-squares.
- Gauss-Newton search:

$$
\begin{aligned}
\mathrm{M}(\theta) &\approx \mathrm{M}(\theta_0) + D(\theta - \theta_0) \\
\mathrm{M}(\theta) - \mathrm{M}(\theta_0) &\approx D(\theta - \theta_0) \\
\hat{\mathrm{M}} - \mathrm{M}(\theta_i) &\approx D\beta \\
\theta_{i+1} &= \theta_i + \hat{\beta}
\end{aligned}
$$

Exponential Example

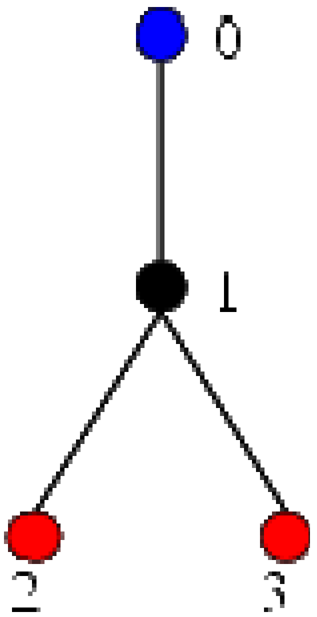

$$\begin{aligned}
\mathrm{E}(Y_2) &= \theta_1 + \theta_2 \\
\mathrm{E}(Y_3) &= \theta_1 + \theta_2 \\
\mathrm{E}(Y_2 - \nu_2)(Y_3 - \nu_3) &= \theta_1^2 \\
\mathrm{E}(Y_2 - \nu_2)^2 &= \theta_1^2 + \theta_2^2 \\
\mathrm{E}(Y_3 - \nu_3)^2 &= \theta_1^2 + \theta_3^2
\end{aligned}$$

$$D = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 2\theta_1 & 0 & 0 \\ 2\theta_1 & 2\theta_2 & 0 \\ 2\theta_1 & 0 & 2\theta_3 \end{bmatrix}$$

Semi-parametric Example

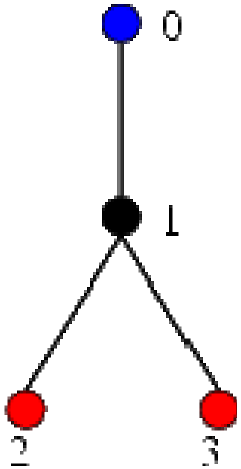

$$D = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ \phi\gamma\mu_1^{\gamma-1} & 0 & 0 & \mu_1^\gamma & \phi log(\mu_1)\mu_1^\gamma \\ \phi\gamma\mu_1^{\gamma-1} & \phi\gamma\mu_2^{\gamma-1} & 0 & \mu_1^\gamma + \mu_2^\gamma & \phi log(\mu_1)\mu_1^\gamma + \phi log(\mu_2)\mu_2^\gamma \\ \phi\gamma\mu_1^{\gamma-1} & 0 & \phi\gamma\mu_3^{\gamma-1} & \mu_1^\gamma + \mu_3^\gamma & \phi log(\mu_1)\mu_1^\gamma + \phi log(\mu_3)\mu_3^\gamma \end{bmatrix}$$

Least Squares Moment Properties

- Asymptotically consistent and normal.
- Moments can include:
  - Observed zero and loss probabilities.
- Fast Estimation.

More UNC Probing Experiments

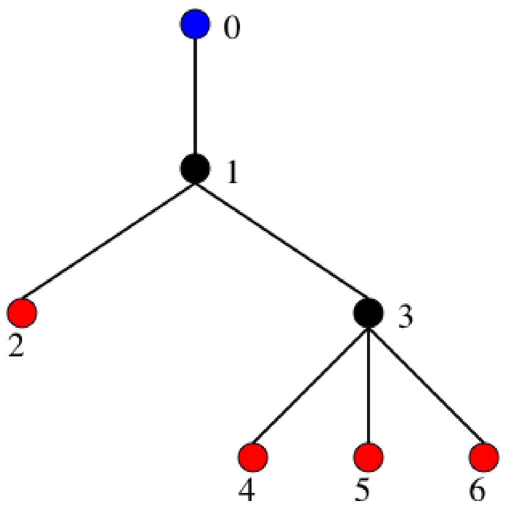

- A small portion of the UNC tree.
- Omnicast: <2,4,5,6>.
- Fit the two moment semiparametric model.

Moment Results

$$\hat{\mu} = \begin{pmatrix} 0.85 \\ 0.50 \\ 0.15 \\ 1.31 \\ 0.48 \\ 1.85 \end{pmatrix}, \quad \hat{\phi} = 53.60, \quad \hat{\gamma} = 0.94.$$

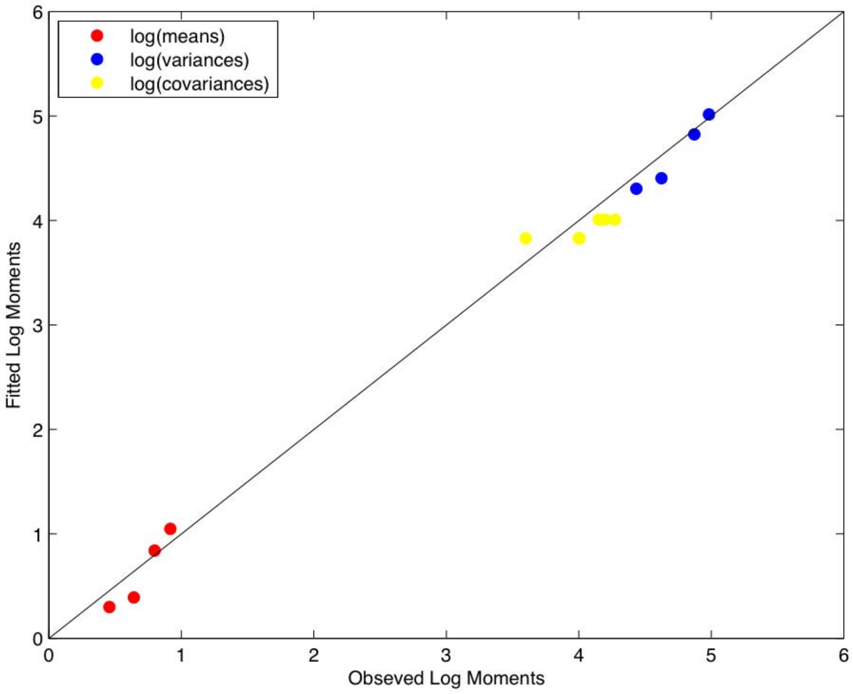

A Monitoring Application

- Ns-2 simulation engine
- Network topology comprised of 6 links with 10 Mb capacity
- Background traffic: combination of TCP and UDP sources
- TCP sources: exponential interarrival times and Pareto durations
- Monitoring period: 60 mins
- Tomography inverse problem solved every 30 secs
- Probing rate: 10 packets/sec
- Monitoring scenario: # of TCP sources doubles on link 3 halfway

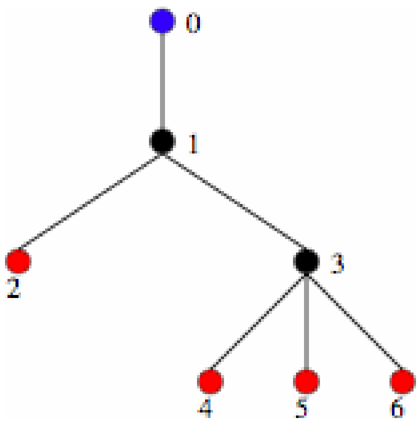

A Monitoring Application

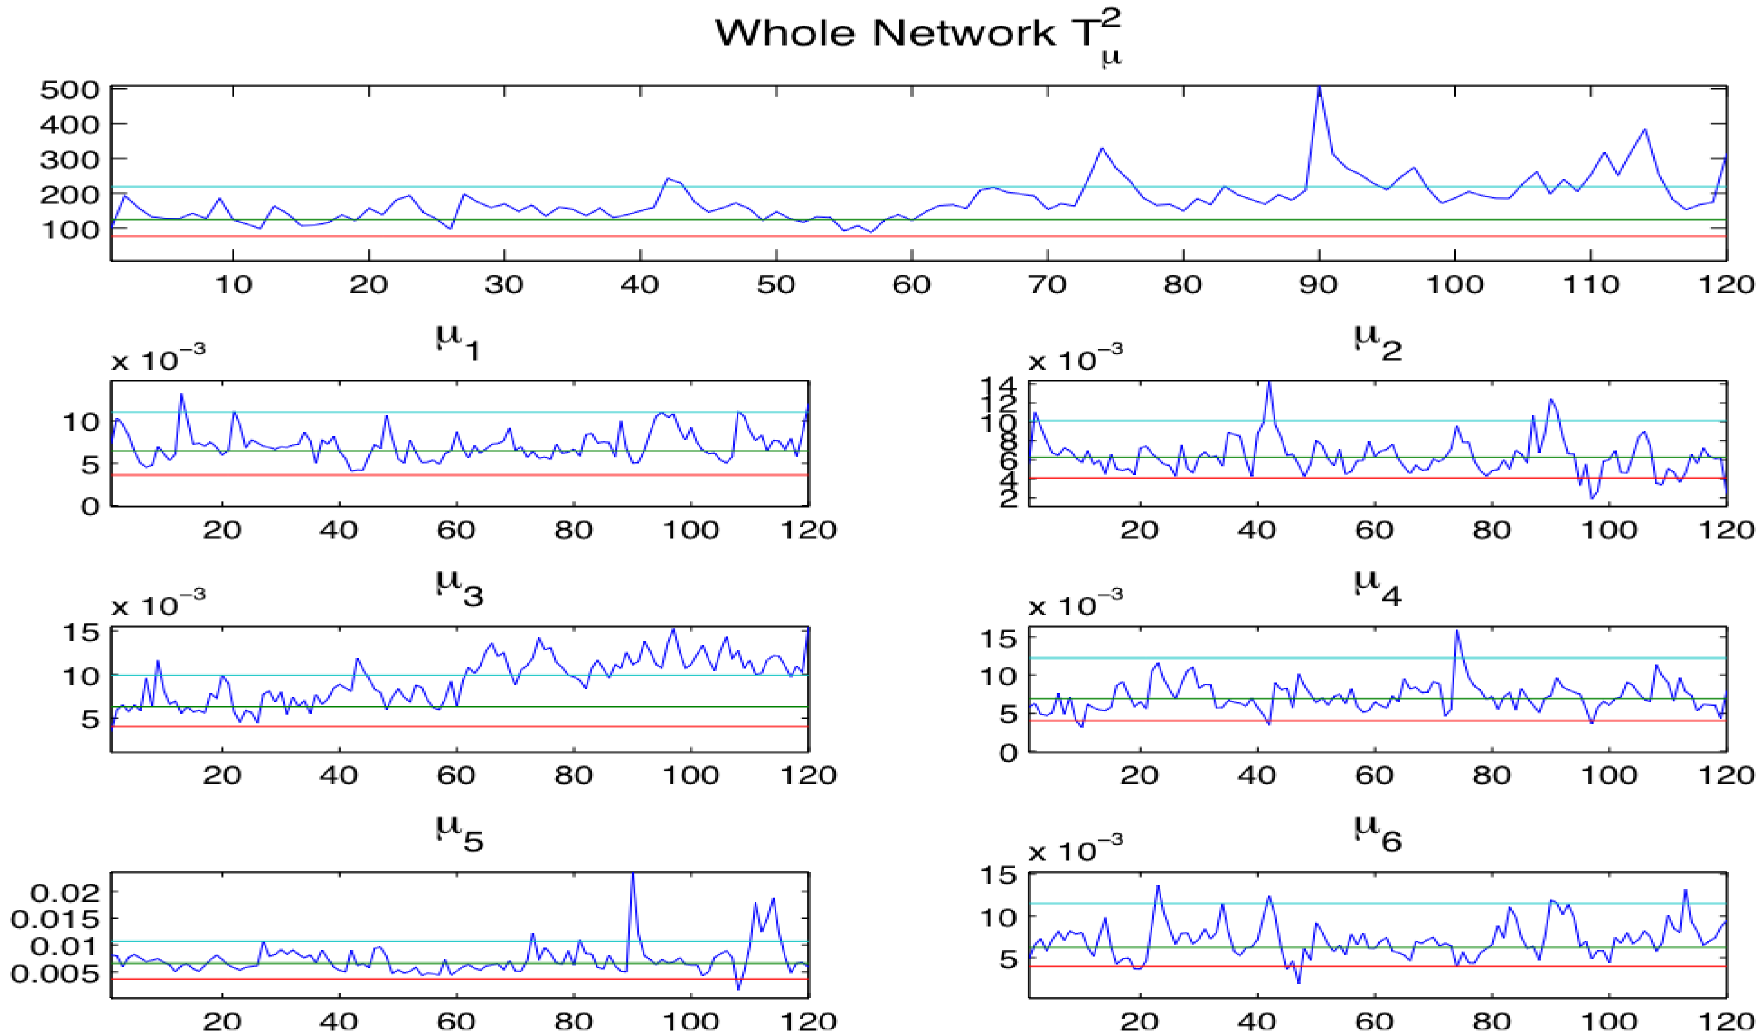

A Monitoring Application

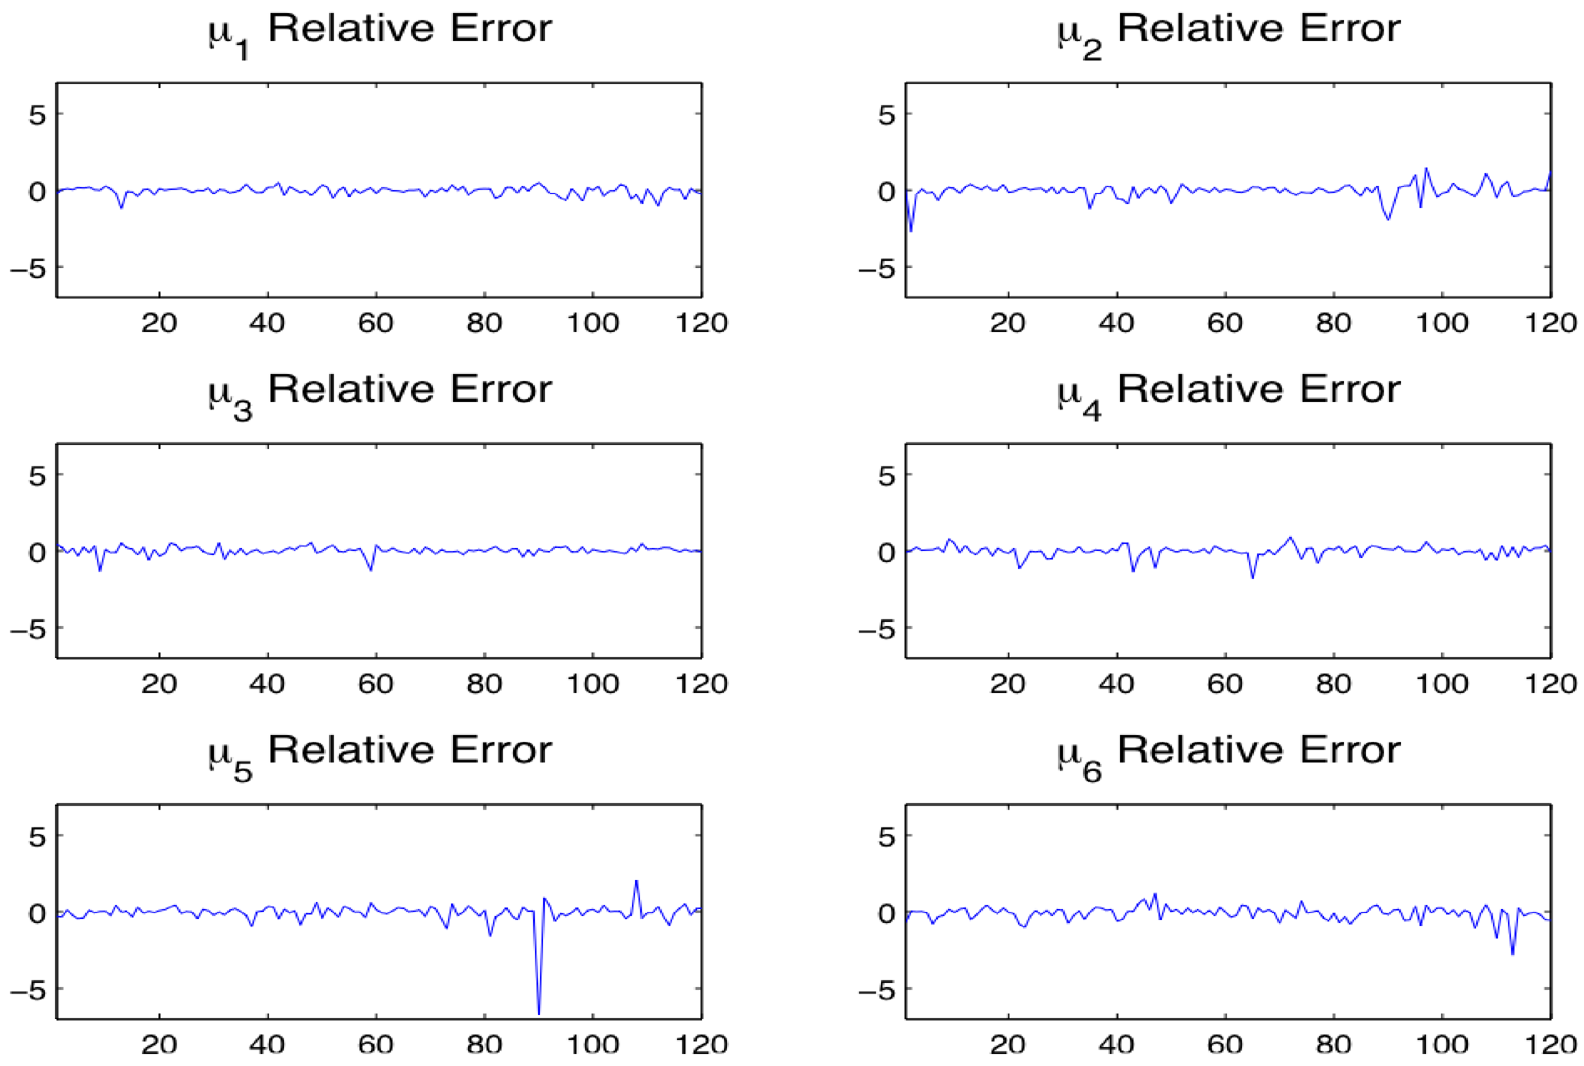

Flexible parametric modeling

- Mixtures of exponential distributions
  - Feldmann, A. and Whitt, W. *Fitting Mixtures of Exponentials to Long-Tail Distributions to Analyze Network Performance Models.* Performance Evaluation (1998).
  - Every completely monotone pdf is a mixture of exponentials.
  - Every completely monotone pdf can be approximated by mixtures of exponentials.

Estimation issues

- Even when the number of mixture components is known maximum likelihood estimation is very involved.
- A better way to proceed is to use a Bayesian approach combined with the Metropolis algorithm

Example

A four component mixture model on each link:
λ=[1,.1,.01,.001],
π=[.4,.3,.2,.1]

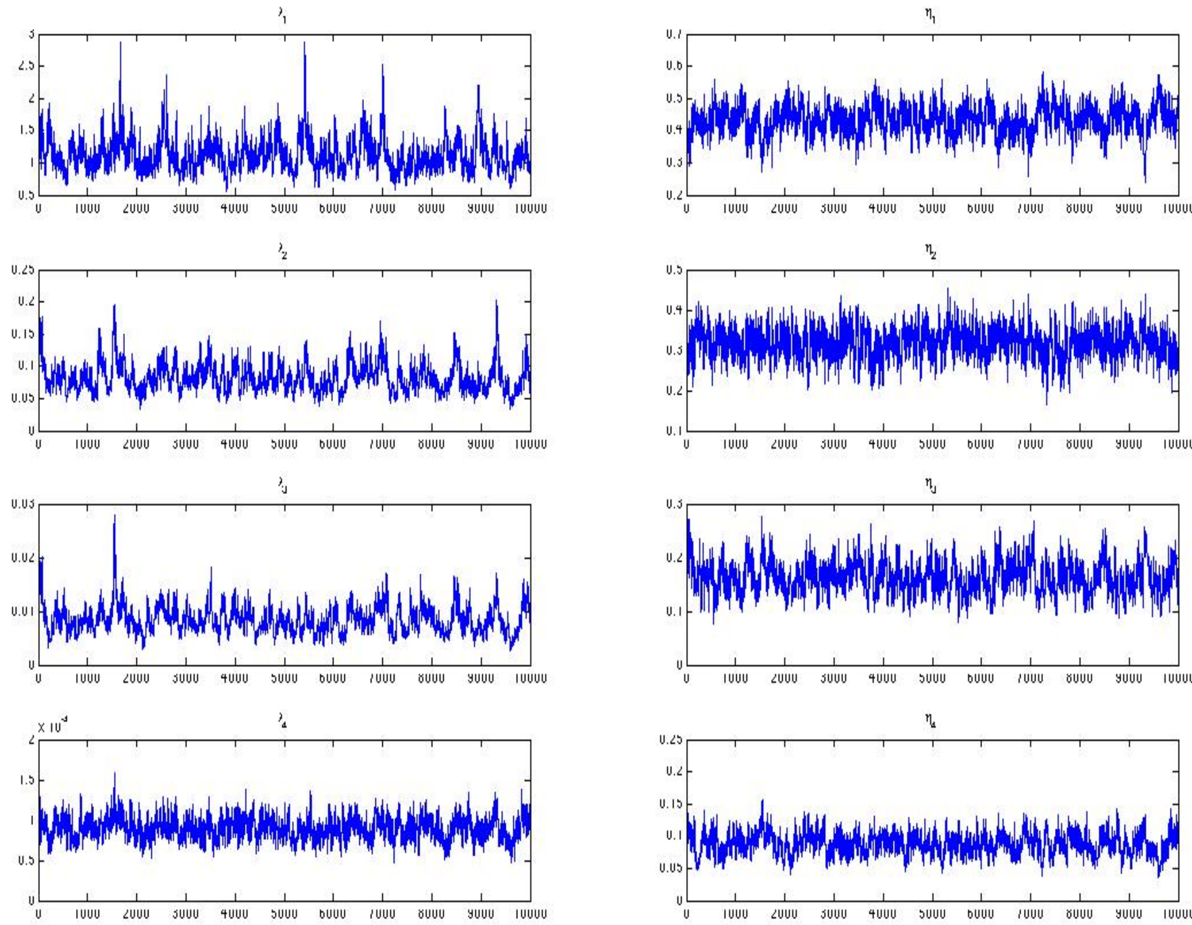

Advantages of the mixture modeling framework

- Tails can be approximated well.
- Complexity is linear in the number of parameters and observations.
- Can include prior information from previous estimates, which can feed into monitoring scenarios.

Concluding Remarks

- Active tomography techniques prove useful in characterizing network performance
- Suitable for network monitoring
- Additional challenges: more realistic models, time-varying parameters, efficient and fast estimation
- Other issues: multi-source topologies, different types of data collection schemes, topology ID problem, network monitoring