Delay tomography for large scale networks

MENG-FU SHIH ALFRED O. HERO III

Communications and Signal Processing Laboratory

Electrical Engineering and Computer Science Department

University of Michigan, 1301 Beal. Ave., Ann Arbor, MI 48109-2122

URSI General Assembly 2002

Network Monigoring and Diagnosis

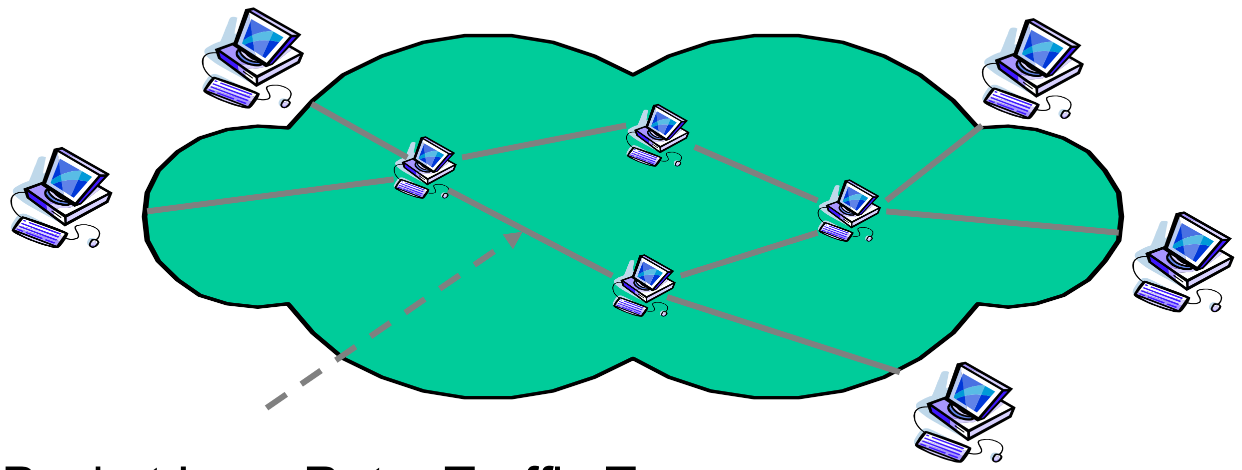

Delay, Packet Loss Rate, Traffic Type, ...

Problems with direct measurement (rmon):

- Diagnosis unavailable or disabled at internal nodes.
- Non-cooperative internal nodes.
- All internal nodes must be synchronized

Network Tomography Problem

- End-to-End Measurements

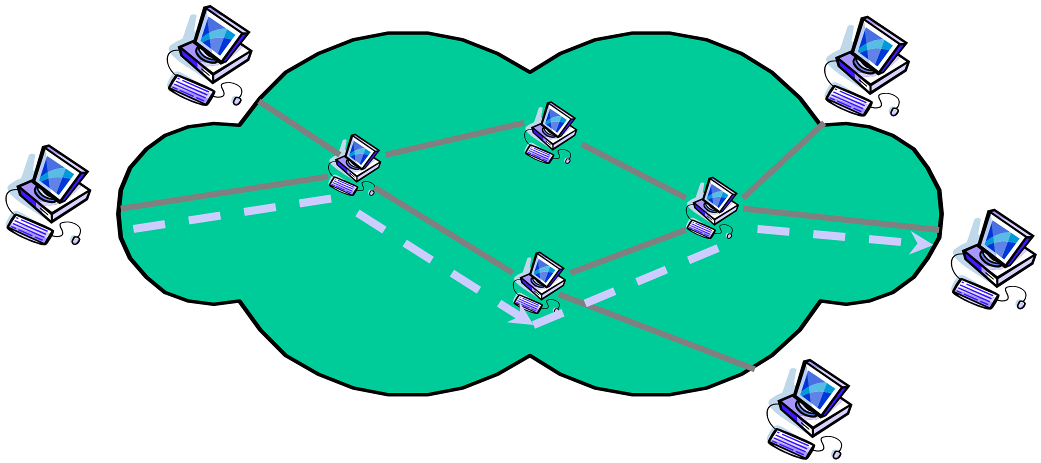

- Active vs. Passive Method
  - Active Method: Send probe packets
  - Passive Method: Monitor existing flows

Importance of Link Delay Statistics

- Assessment and updating of routing/flow control
- QoS assurance, especially for video/audio streaming
- Network upgrade/maintenance planning
- Security, e.g., distributed Denial-of-Service (DoS) attacks

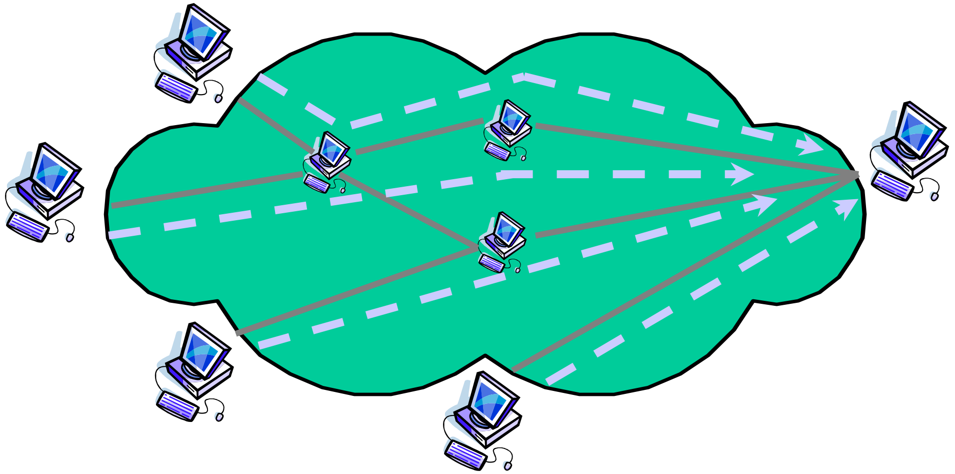

Problem Formulation: General Notations

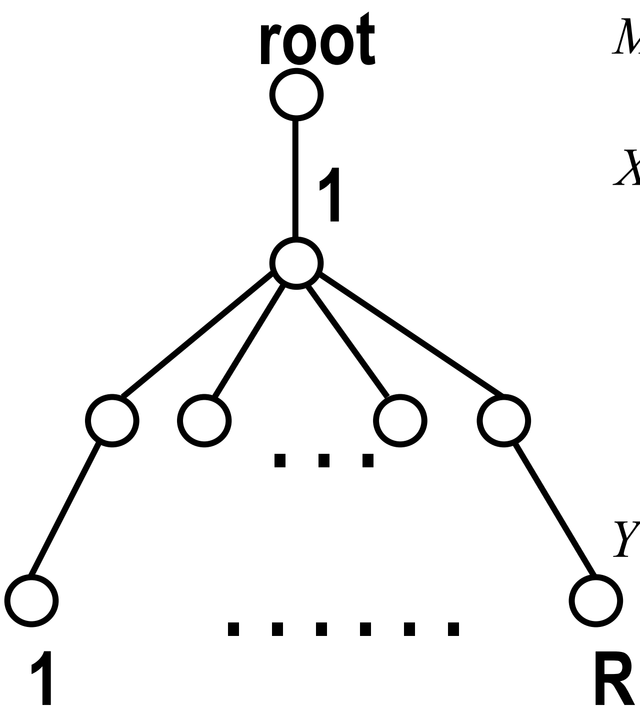

Logical Tree $T=(V,E)$ V : Nodes, E : Links

L links, R leaf nodes/probe paths

N_i : # of packets sent from root to leaf i

M_i : The set of links in probe path i.

$X_l^{(i,n)}$: nth probe packet delay at link l along path i, including queueing delay, retransmission delay, and possibly propagation delay.

$Y^{(i,n)} = \sum_{l \in M_i} X_l^{(i,n)}$: nth End-to-end probe delay along path i.

Problem Formulation: General Assumptions

- Network Assumptions

N1) Network topology known.

N2) Probe paths (routing table) known.

N3) Cooperating edge nodes are synchronized

- Statistical Assumptions

S1) Spatial Independence

For a given packet along path i, $\left\{X_l^{(i,n)}\right\}_{l\in M_i}$ mutually independent.

S2) Temporal Independence and Stationarity

If path i and k both contain link l, $X_l^{(i,n_1)}$ and $X_l^{(j,n_2)}$ i.i.d.

Discrete Delay Model

- Link delays are discretized with bin size q
- Link delay values $X_l^{(i,n)} \in \{0, q, 2q, \cdots, qD\}$
- Link Delay P.M.F. $p_{l,d} = P\left(X_l^{(i,n)} = d\right)$
- Lemma 1.

The delay p.m.f. with two bins at each link is uniquely identifiable from end-to-end packet delays, except when the delay p.m.f.'s at all links are identical.

$$\mathbf{A} = n\begin{bmatrix} \frac{p_2+p_3}{p_1} + \frac{2-p_2-p_3}{1-p_1} + \frac{(1-2p_2)^2}{Q_1} + \frac{(1-2p_3)^2}{Q_2} & \frac{1}{Q_1} & \frac{1}{Q_2} \\ \frac{1}{Q_1} & \frac{p_1}{p_2} + \frac{1-p_1}{1-p_2} + \frac{(1-2p_1)^2}{Q_1} & 0 \\ \frac{1}{Q_2} & 0 & \frac{p_1}{p_3} + \frac{1-p_1}{1-p_3} + \frac{(1-2p_1)^2}{Q_2} \end{bmatrix}$$

$Q_1 = p_1(1-p_2) + p_2(1-p_1) \qquad Q_2 = p_1(1-p_3) + p_3(1-p_1)$

Continuous Delay Model: Gaussian Mixture

- Arbitrary shapes of link delay distributions
- Let $f_l(x) = \sum_{m=1}^{k_l} \alpha_{l,m} \phi(x; \theta_{l,m})$ be the link delay p.d.f at link l.

k_l : the number of mixture components.

$\alpha_{l,m}$: mixing probability for the mth component.

$$0 \le \alpha_{l,m} \le 1, \quad \sum_{m=1}^{k_l} \alpha_{l,m} = 1$$

$\phi(x; \theta_{l,m})$: Gaussian density function with mean and

variance $\theta_{l,m} = \left\{ \mu_{l,m}, \sigma_{l,m}^2 \right\}$

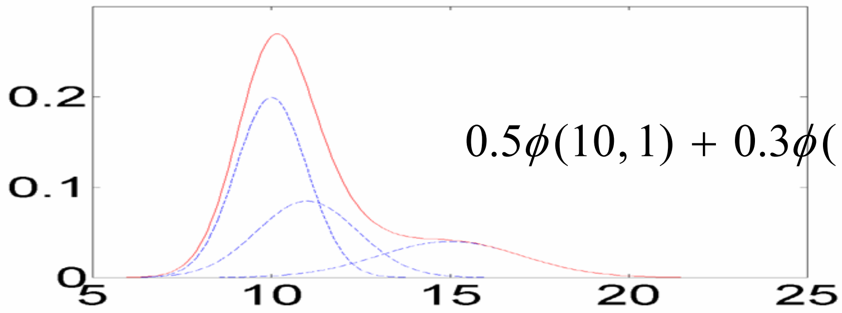

$0.5\phi(10, 1) + 0.3\phi(11, 2) + 0.2\phi(15, 4)$

Continuous Delay Model: Identifiability Problem

- Example: Two leaf tree. Let $k_1, k_2, k_3 = 1$

$$f(y_1, \mathrm{y}_2) = \phi(y_1; \{\mu_1 + \mu_2, \sigma_1^2 + \sigma_2^2\}) \cdot$$

$$\phi(y_2; \{\mu_1 + \mu_3, \sigma_1^2 + \sigma_3^2\})$$

$$\mu_{Y_1} = \mu_1 + \mu_2$$

$$\mu_{Y_2} = \mu_1 + \mu_3$$

$$\sigma_{Y_1} = \sigma_1 + \sigma_2$$

$$\sigma_{Y_2} = \sigma_1 + \sigma_3$$

4 equations with 6 unknowns!

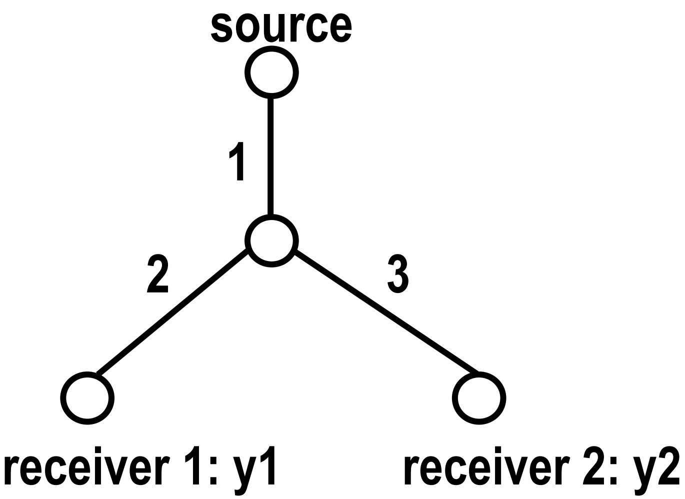

Mixed Finite Mixture Model

- Utilization factor of a queueing system ρ

$0 \le \rho < 1$ for stable system. $\Rightarrow P(\text{Queue is empty}) = 1 - \rho = \alpha_0$

- Introduce a delta component at (or near) 0 with probability mass α_0
- Link delay p.d.f. becomes $f_l(x) = \alpha_{l,0}\delta(x) + \sum_{m=1}^{k_l} \alpha_{l,m}\phi(x;\theta_{l,m})$

$$\sum_{m=0}^{k_l} \alpha_{l,m} = 1$$

- Sufficient condition for identifiability (asymptotic)

The delay distribution defined above is identifiable from end-to-end measurements if (1) $\alpha_{l,0} > 0$ for all l *(2) All the Gaussian components in link delay distributions have distinct means and variances.*

Mixed Finite Mixture Model: Example

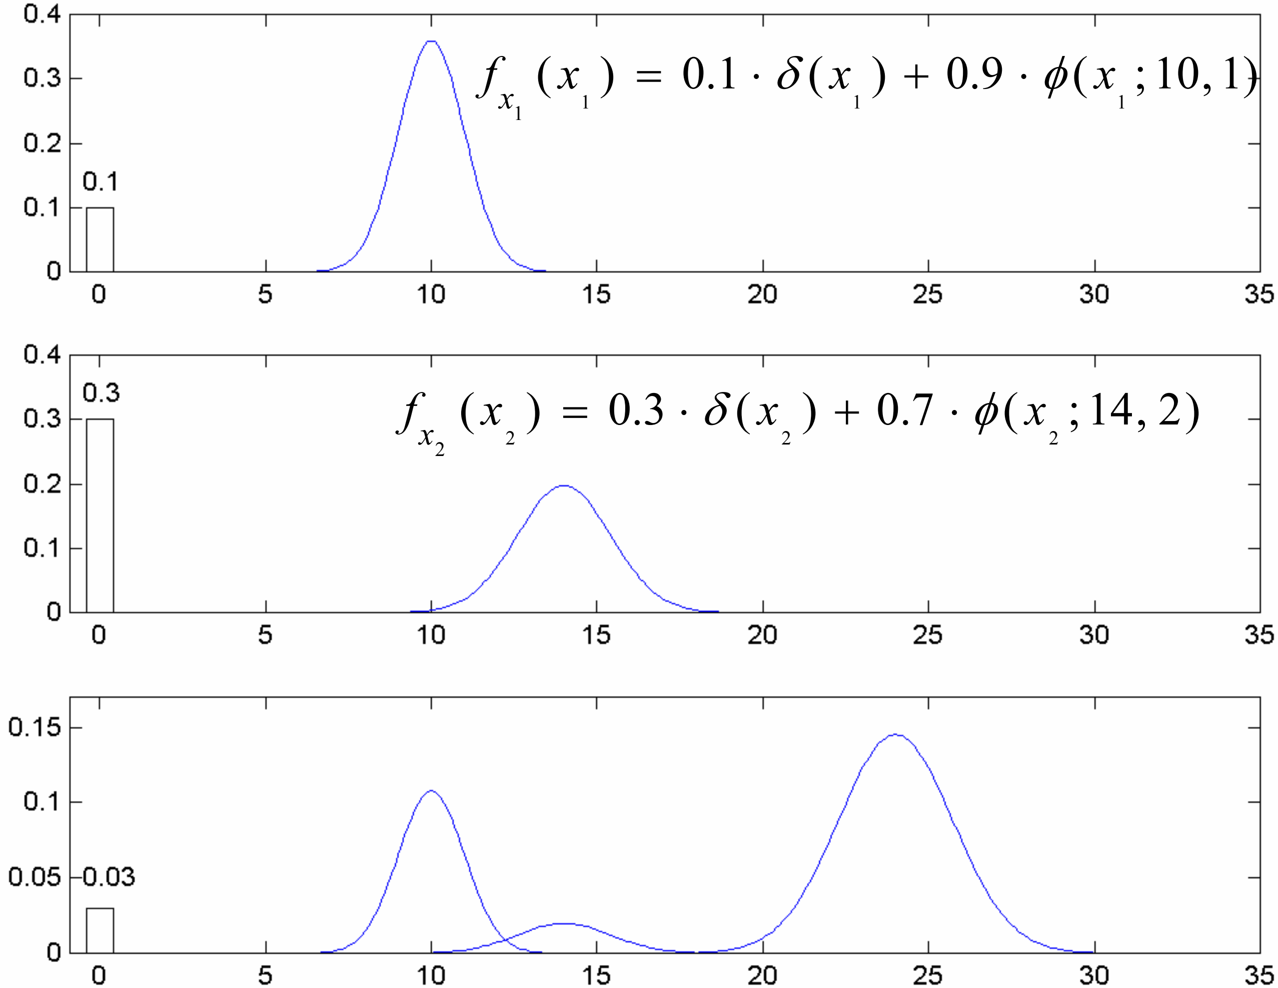

$$f_{x_1+x_2}(x_3) = 0.03 \cdot \delta(x_3) + 0.27 \cdot \phi(x_3; 10, 1) + 0.07 \cdot \phi(x_3; 14, 2) + 0.63 \cdot \phi(x_3; 24, 3)$$

EM Estimation Algorithm: Notations

- Assume prior knowledge of $\{k_l\}$
- Component indicator vector $z_l^{(i,n)} = \left\{z_{l,0}^{(i,n)}, z_{l,1}^{(i,n)}, \cdots, z_{l,k_l}^{(i,n)}\right\}$

$z_{l,m}^{(i,n)} = 1$ if $x_l^{(i,n)}$ is generated by the mth component,

$z_{l,m}^{(i,n)} = 0$ otherwise

- Unobserved data $\{\mathbf{x}, \mathbf{z}\}$, $\mathbf{x} = \left\{x_l^{(i,n)}\right\}$, $\mathbf{z} = \left\{z_l^{(i,n)}\right\}$
- Observed data $\mathbf{y} = \left\{y^{(i,n)}\right\}$

- Complete data $\{\mathbf{x}, \mathbf{y}, \mathbf{z}\}$
- Parameter vector $\Theta = \left\{\alpha_{l,m}, \theta_{l,m}\right\}$

EM Estimation Algorithm

- Complete data likelihood $\propto \log L(\mathbf{x},\mathbf{z}\mid\Theta)$

$$\log L(\mathbf{x},\mathbf{z}\mid\Theta) = \sum_{l=1}^{L}\sum_{i:l\in M_i}\sum_{n=1}^{N_i}\Big\{ z_{l,0}^{(i,n)}\log\alpha_{l,0} + \sum_{m=1}^{k_l} z_{l,m}^{(i,n)}\left(\log\alpha_{l,m} + \log\phi(x_l^{(i,n)};\theta_{l,m})\right)\Big\}$$

- Let

$$\omega_{l,m}^{(i,n)} = E\left[z_{l,m}^{(i,n)} \mid y^{(i,n)};\Theta^t\right]$$

$$Q_{l,m}^{(i,n)}(\theta_{l,m},\Theta^t) = E\left[z_{l,m}^{(i,n)}\log\phi(x_l^{(i,n)};\theta_{l,m}) \mid y^{(i,n)};\Theta^t\right]$$

EM Estimation Algorithm

- E-Step

$$E\left[\log L(\mathbf{x},\mathbf{z}\mid\Theta)\mid\mathbf{y};\Theta^t\right]=\sum_{l=1}^{L}\sum_{i:l\in M_i}\sum_{n=1}^{N_i}$$

$$\left\{\sum_{m=0}^{k_l}\omega_{l,m}^{(i,n)}\log\alpha_{l,m}+\sum_{m=1}^{k_l}Q_{l,m}^{(i,n)}(\theta_{l,m},\Theta^t)\right\}$$

- M-step $\alpha_{l,m}^{t+1}=\dfrac{\sum_{i:l\in M_i}\sum_{n=1}^{N_i}\omega_{l,m}^{(i,n)}}{\sum_{i:l\in M_i}N_i}$

$$\theta_{l,m}^{t+1}=\arg\max_{\theta}\sum_{i:l\in M_i}\sum_{n=1}^{N_i}Q_{l,m}^{(i,n)}(\theta,\Theta^t)$$

Computer Experiment

- Matlab Simulation with 15000 i.i.d. end-to-end delays for each probe path.
- Numbers of Gaussian mixture components and true/estimated delta factor $\alpha_{l,0}$

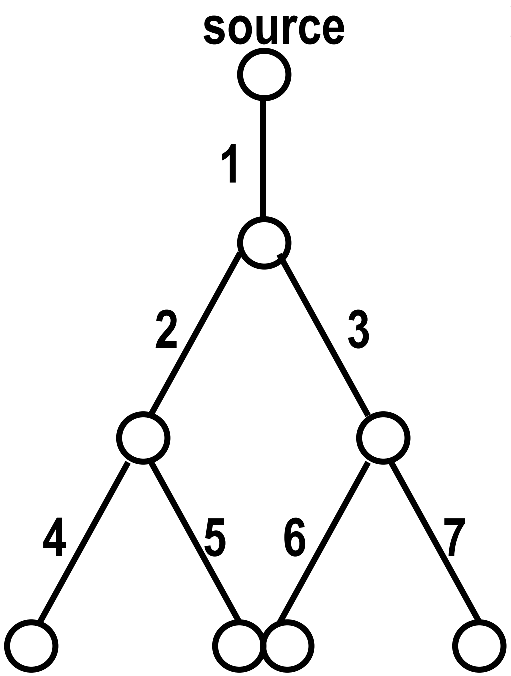

Link	1	2	3	4	5	6	7
k_l	3	2	2	2	2	2	2
$\alpha_{l,0}$	0.25	0.3	0.1	0.2	0.15	0.3	0.2
$\hat{\alpha}_{l,0}$	0.253	0.304	0.099	0.199	0.152	0.313	0.201

- True (solid) and estimated (dotted) Gaussian mixture components.

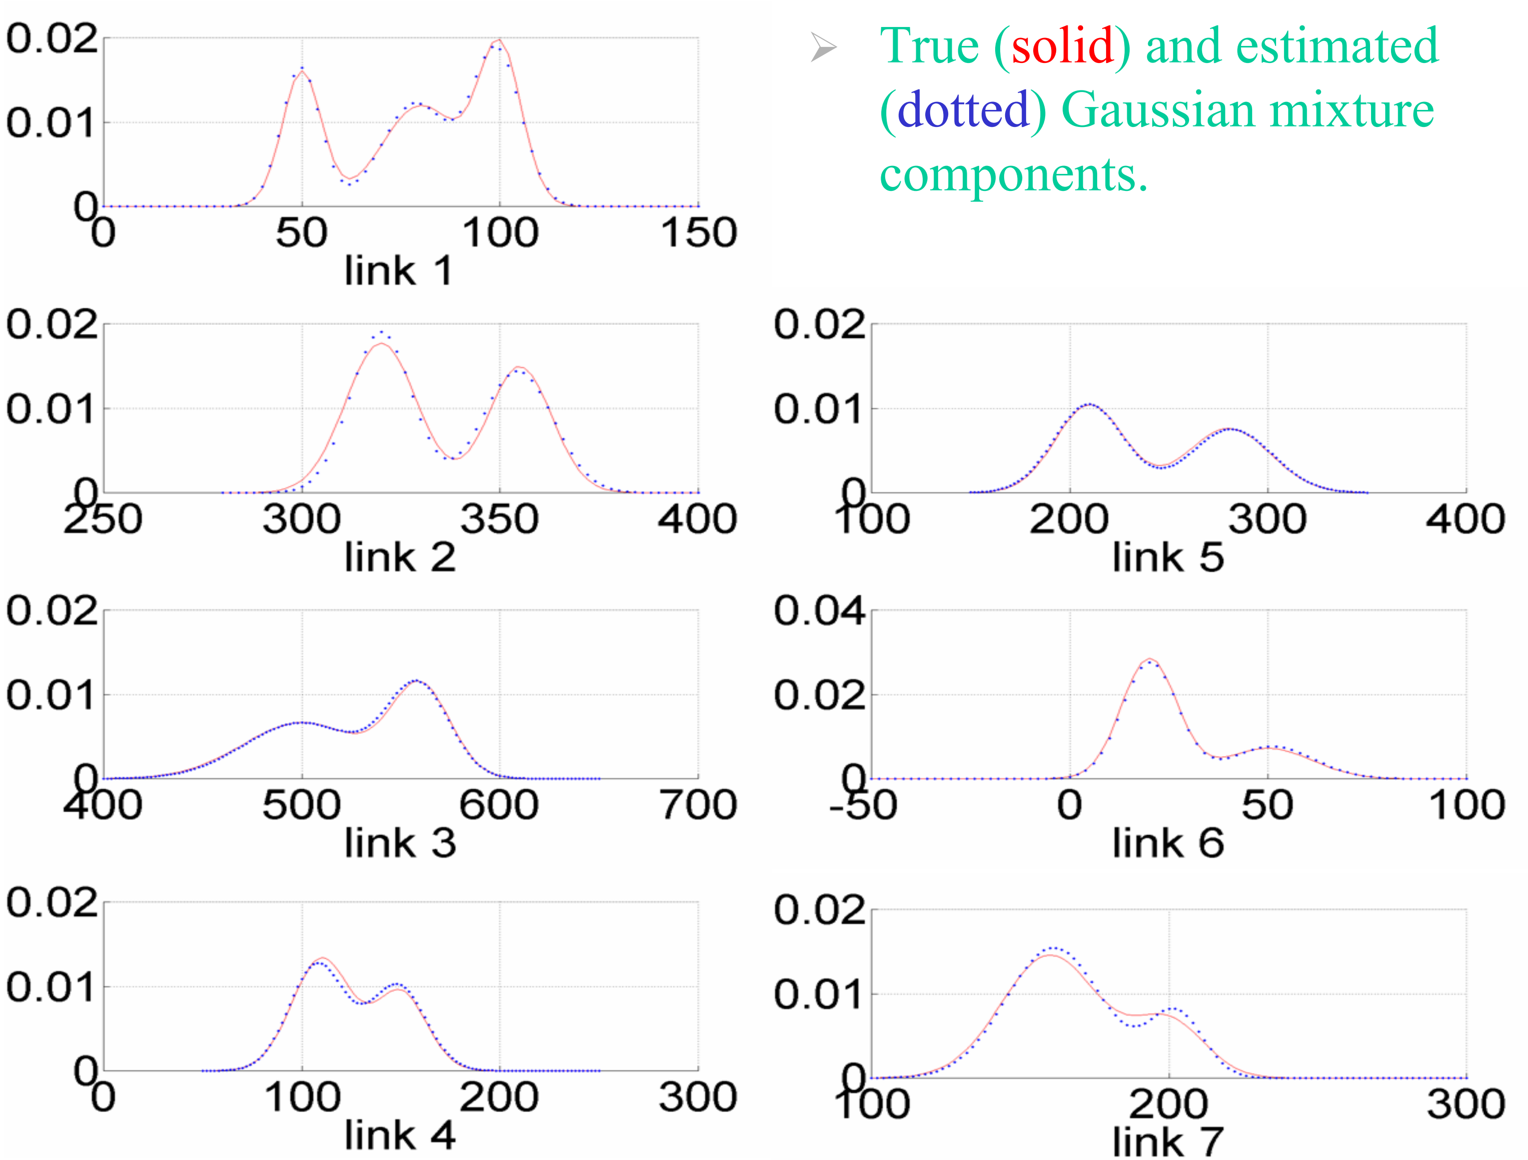

Conclusion and Extensions

- Conclusions
  - Discussion of discrete and continuous delay models.
  - Proposed mixed finite Gaussian mixture model for link delay.
  - EM algorithm implementation with known model orders.
- Extensions
  - Unsupervised model order estimation.
  - Adaptive algorithm for parameter and model order update.

References

- F. L. Presti, N. G. Duffield, J. Horowitz, D. Towsley, "Multicast-based inference of network-internal delay distributions," Umass CMPSCI 99-55, 1999.
    - Logical multicast tree.
    - Discrete link delays with finite levels.
    - Canonical delay tree, i.e., there is a nonzero probability that a probe experiences no delay in traversing each link.
    - Sample-average approach. Identifibility is proved by showing bijection mapping exists from the link delay distributions to the probabilities of the events in which the end-to-end delay is no greater than *iq* for at least one receiver.
    - Continuous model is discussed, but identifiability problem is left open.

- M. Coates and R. Nowak, "Network tomography for internal delay estimation," ICASSP 2001, Salt Lake City, May 2001.
    - Logical unicast tree.
    - Discrete link delays with finite levels.
    - Back-to-back packet pair measurements.
    - MLE using EM-based algorithm.
    - Sequential Monte Carlo tracking of time variation.