Toward A Complete Active Tomography Framework

Earl Lawrence

LANL

Outline

- Background.
- State Of The Art.
- Future Directions.

Internet In Action

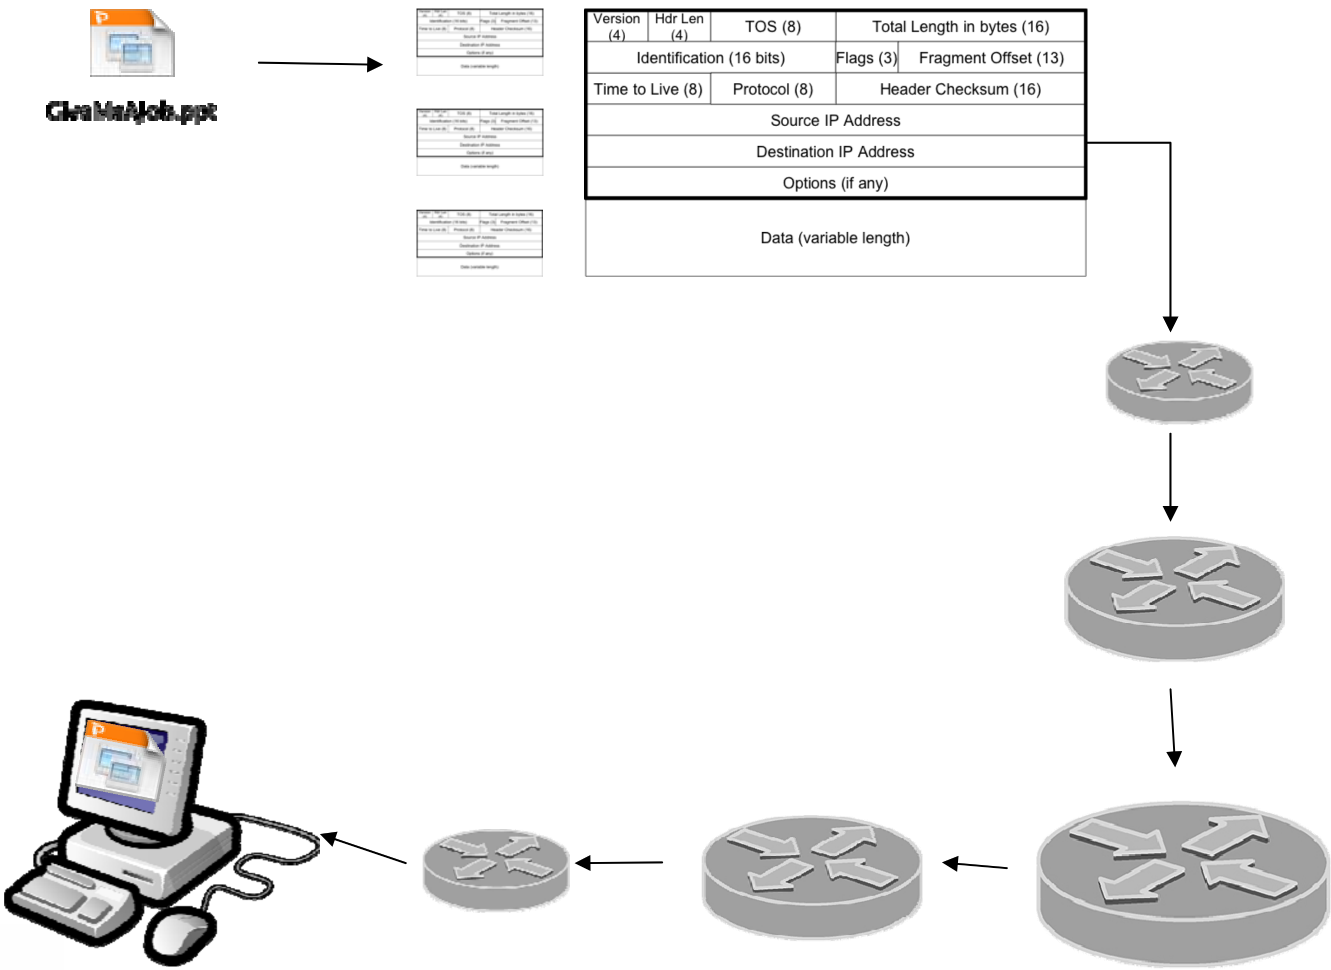

Queueing and Loss

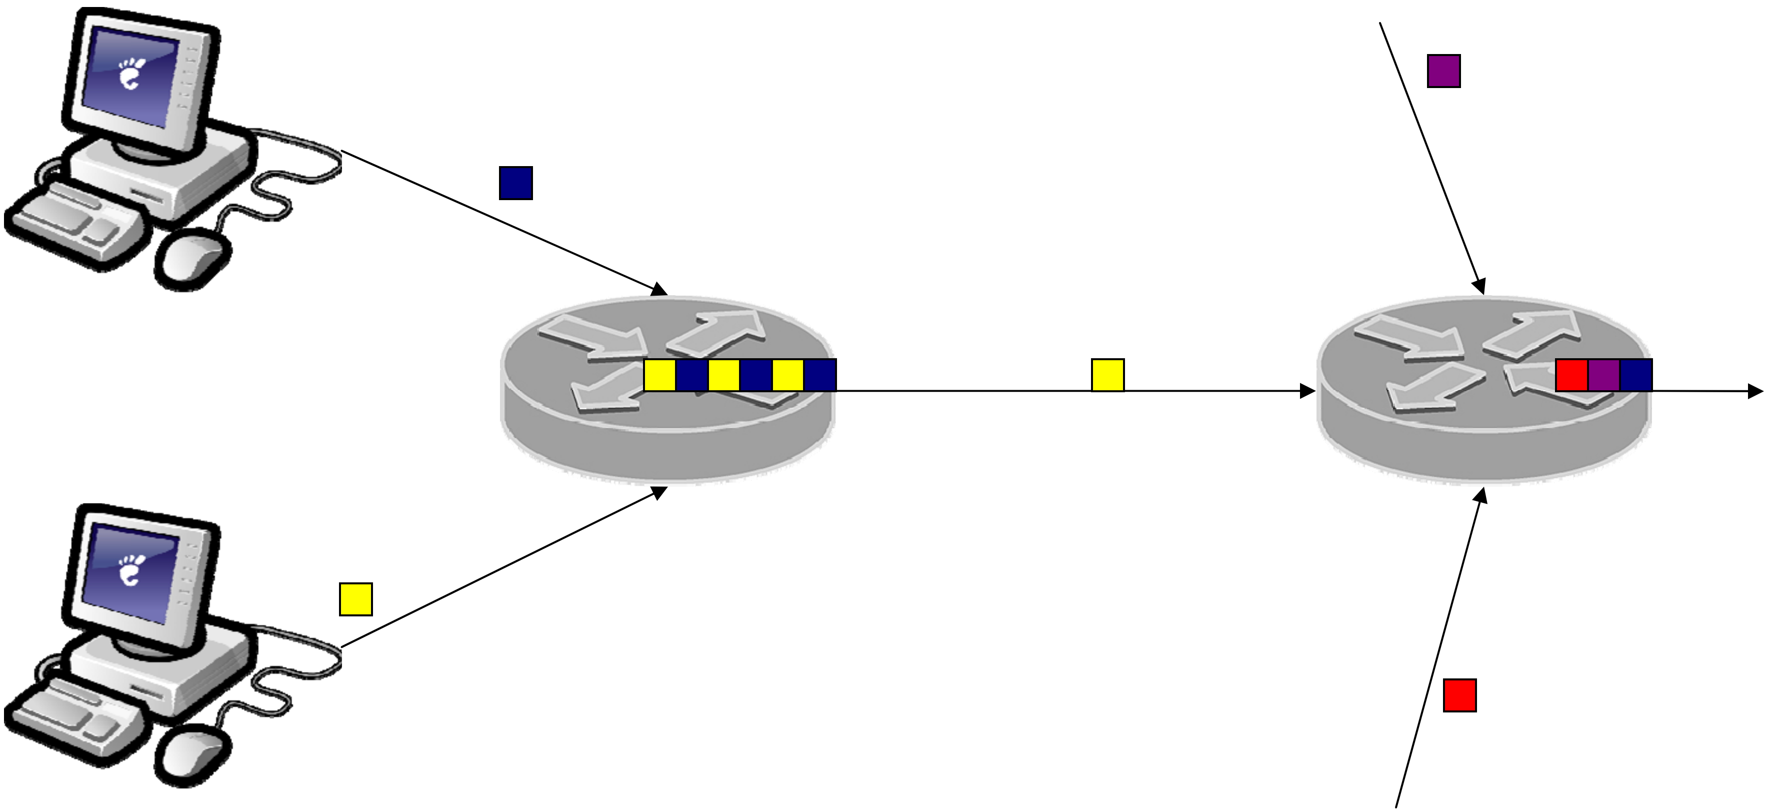

Modeling Goals

- Performance assessment.
- Anomaly detection.
- Capacity planning.
- Dynamic routing.
- All of the above, but faster.

Why is that so hard?

- Networks are decentralized and multilayered.
- Access to measurement is restricted.
- Lots and lots of users/applications.

Voice-Over-IP

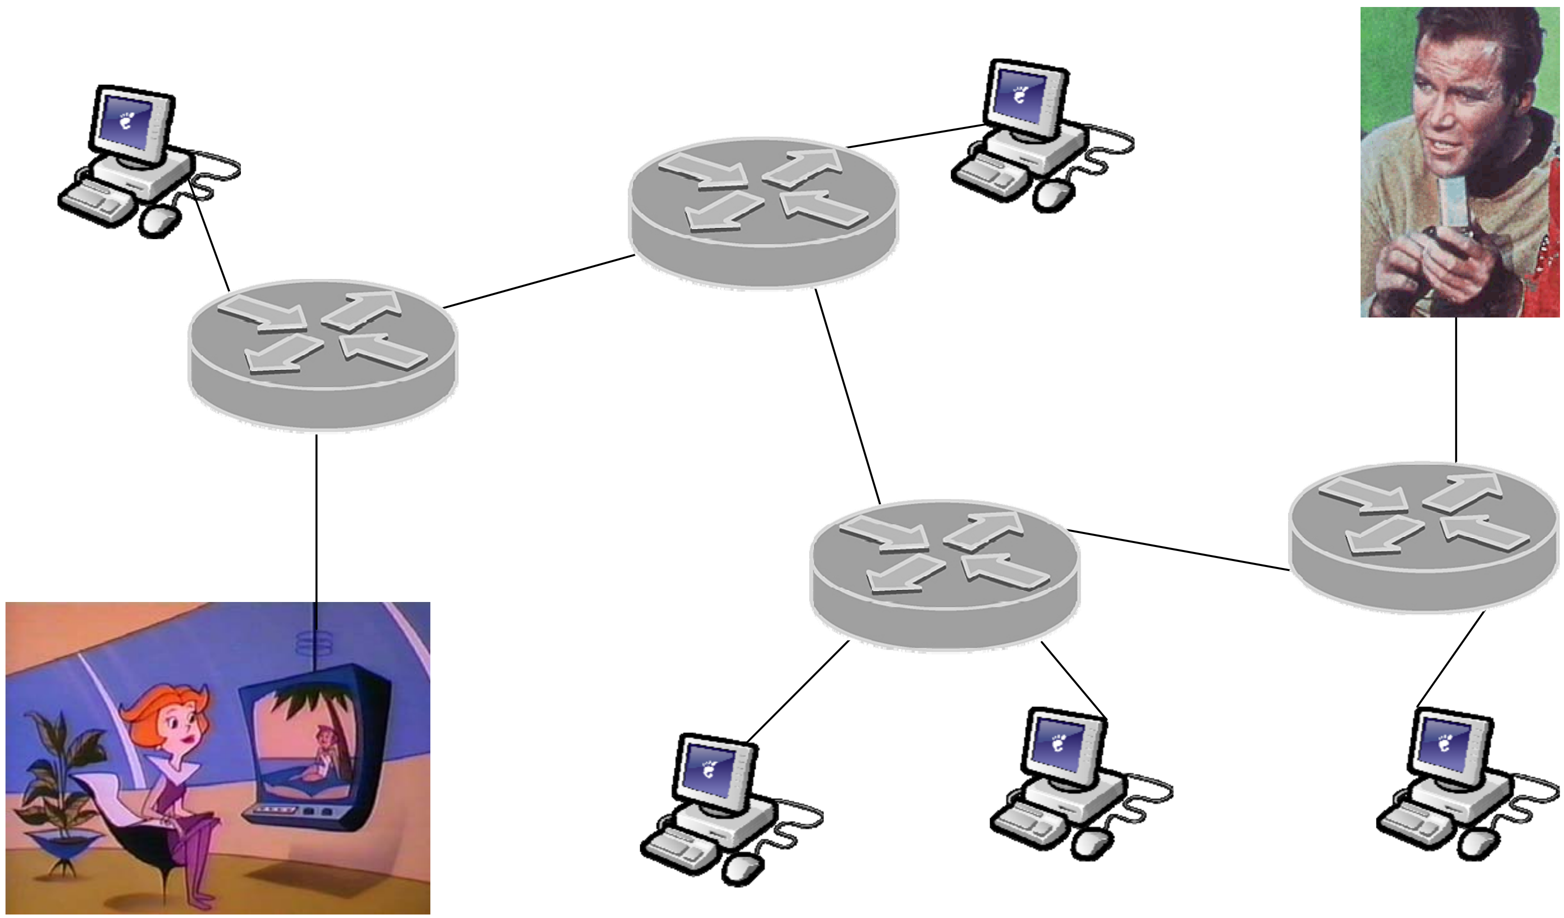

Active Network Tomography

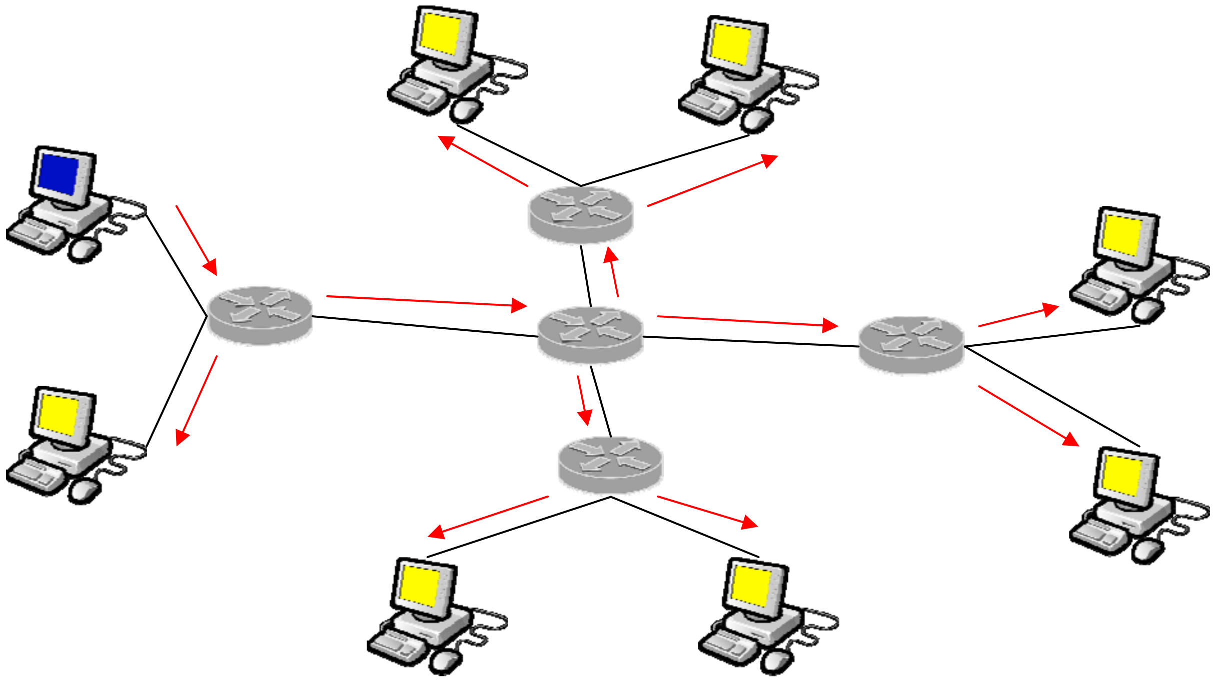

- No internal access.
- Inject traffic and collect end-to-end performance metrics.

Every stat talk should have this equation:

$$Y = AX$$

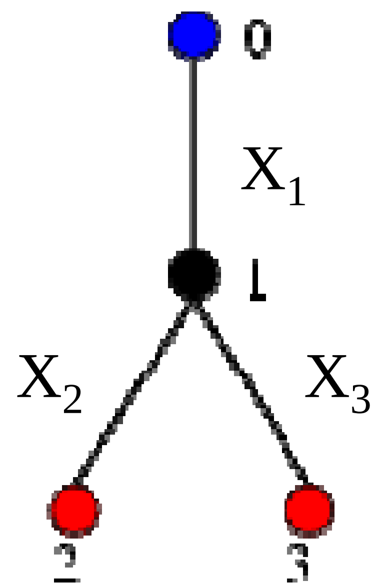

$$Y_2 = X_1 + X_2 \qquad Y_3 = X_1 + X_3$$

Tree Topology

- Induced by single sourcing probing.
- Logical topology: only includes branching points.
- Can be generalized to multi-source by considering overlapping trees.

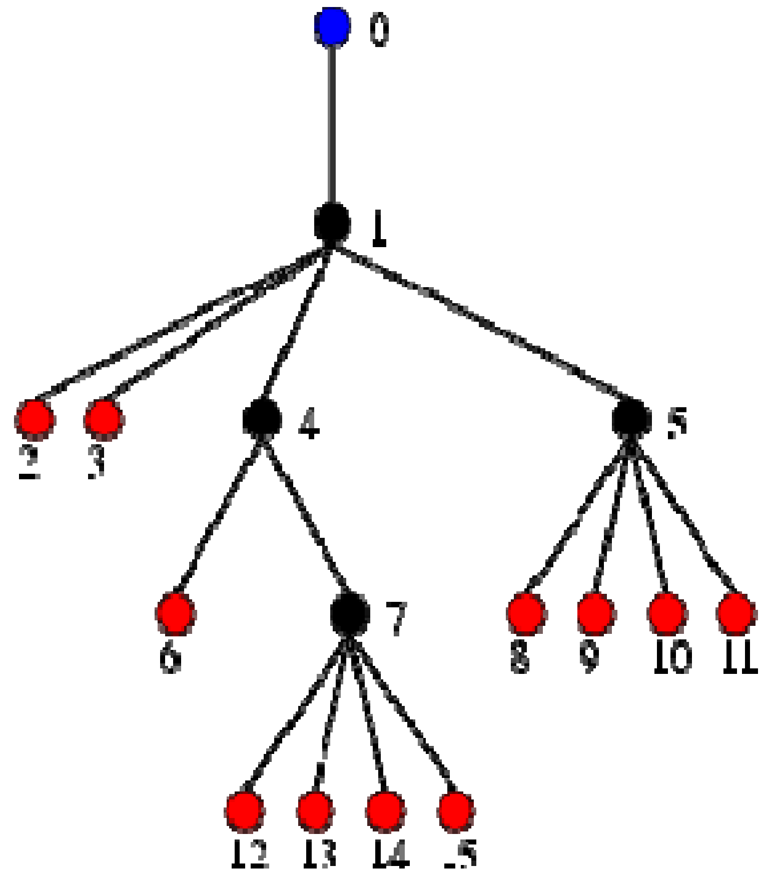

Probing

- Most traffic is unicast -- single source, single receiver.
  – Not sufficient for general estimation.
- Multicast or back-to-back unicast can probe groups of receivers simultaneously.
  – Results in correlated observations that help regularize the estimation problem.
- Flexicast

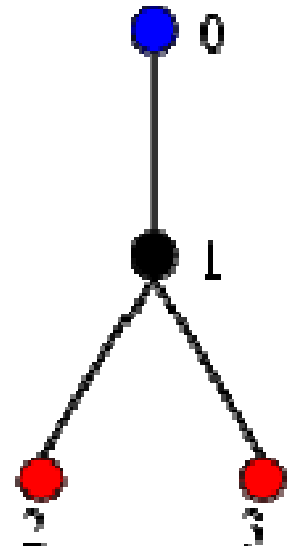

Discrete Delay Modeling

$$X_k \in \{0, q, 2q, \ldots, bq\}$$
$$P\{X_k = iq\} = \alpha_k(i)$$

- Simple structure.
- Flexible shape.
- Can be adapted for varying goals.

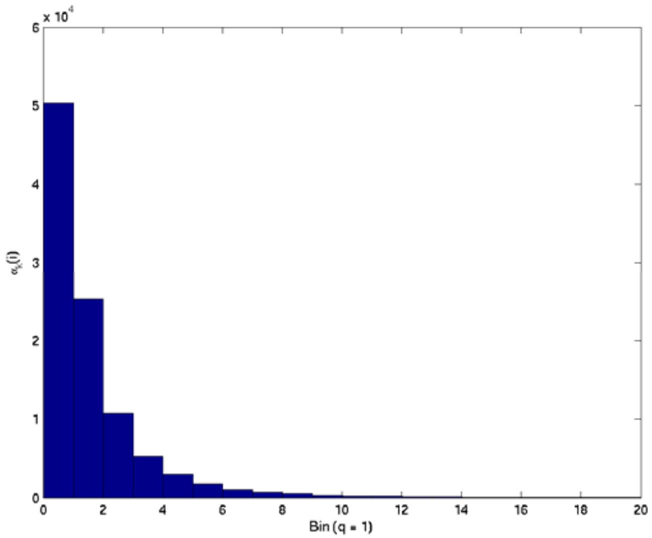

Identifiability

In order for a collection of flexicast probing schemes to identify the discrete link delay distribution at every link in a tree, two conditions are necessary and sufficient:

1. Every receiver must be covered.
2. Every internal node must serve as a branching point for at least one probing scheme.

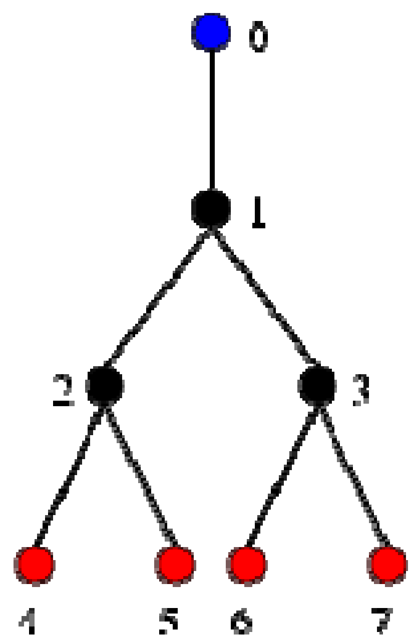

Estimation

- $$l(\vec{\alpha}; \mathbf{Y}) = \sum_{c \in \mathcal{C}} \sum_{\vec{y} \in \mathcal{Y}^c} N^c_{\vec{y}} \log[\sum_{\vec{x} \in \mathcal{X}^c(\vec{y})} \mathrm{P}(\vec{\alpha})\{X^c = \vec{x}\}]$$
- MLE via EM: Impute link counts to compute probabilities.
  - Poor computational complexity.
  - Usual properties hold.
- Grafting: compute local MLEs for probe groups and combine.
  - Lose some statistical efficiency.
  - Maintain some inferential properties.
  - Much better computational efficiency.

General Identifiability

In order for a collection of flexicast probing schemes to identify the link delay distribution at every link in a tree, three conditions are sufficient:

1. α_k, $p_k > 0$
2. Every internal node must serve as a branching point for at least one probing scheme.
3. Every receiver must be covered.

Parametric Modeling

- Simple structure -- good for monitoring.
- Looser identifiability: If the parameters can be identified from second-order and higher moments, then flexicast probing is sufficient.
  - Exponential, Log-Normal, Weibull, Gamma, …
- Usually not very realistic based on knowledge of the traffic.
- Illustrative for a more general fitting procedure.

Semiparametric Modeling

- Specify quantities of interest in terms of a few parameters.

$$\begin{aligned} \mathrm{E}(X_k) &= \mu_k \\ \mathrm{Var}(X_k) &= \phi\mu_k^\gamma \end{aligned}$$

- Fewer assumptions that parametric choices.
- Still appropriate for monitoring.
- Fits within the general framework.
- Amenable to the moment estimation scheme.

Moment Estimation

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

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

- Fast compared with previous methods.

Results

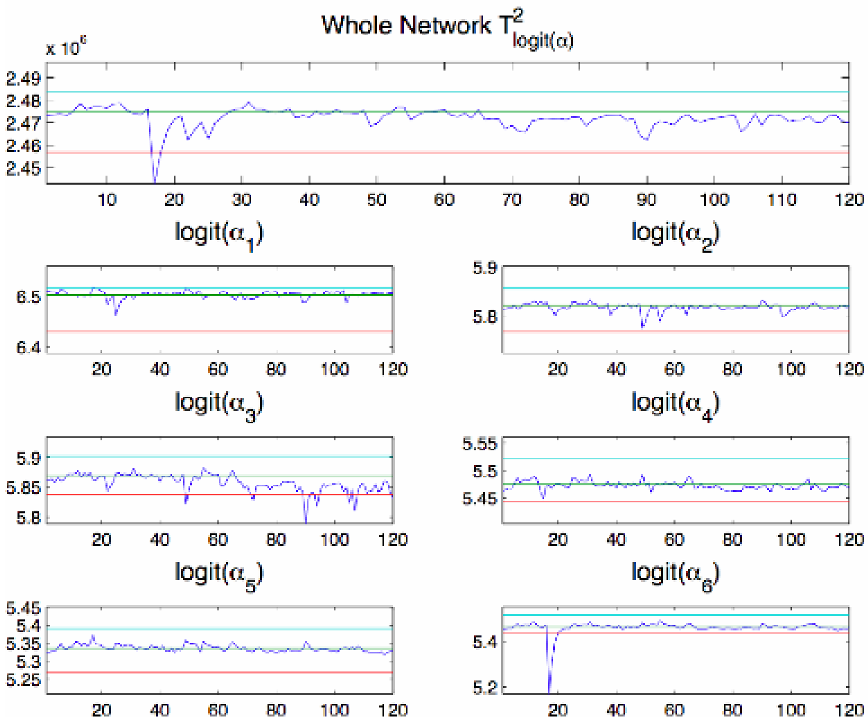
Whole Network $T^2_{logit(\alpha)}$
x 10^6
logit(α_1)
logit(α_2)
logit(α_3)
logit(α_4)
logit(α_5)
logit(α_6)

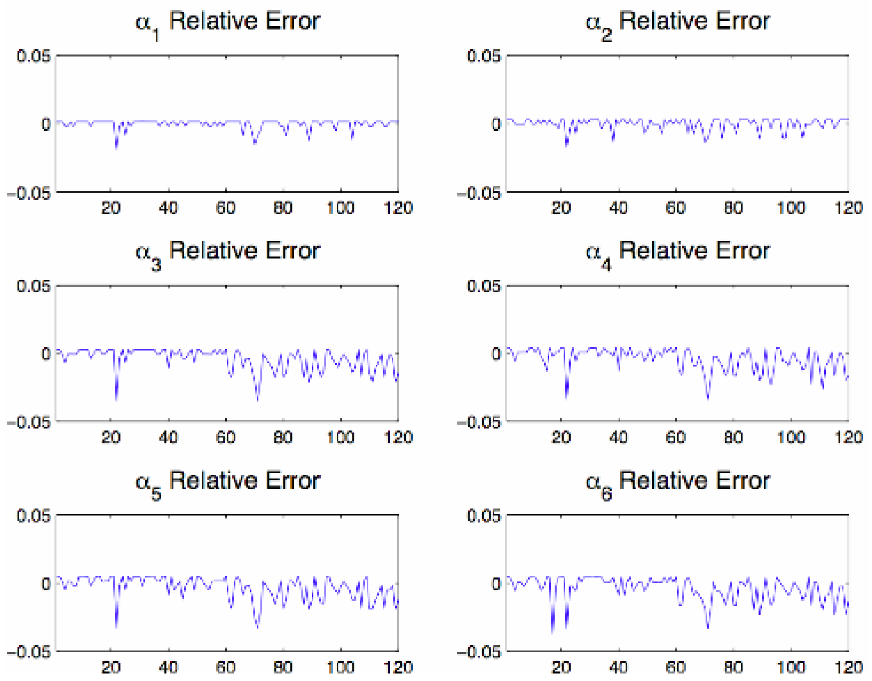
α_1 Relative Error
α_2 Relative Error
α_3 Relative Error
α_4 Relative Error
α_5 Relative Error
α_6 Relative Error

Results

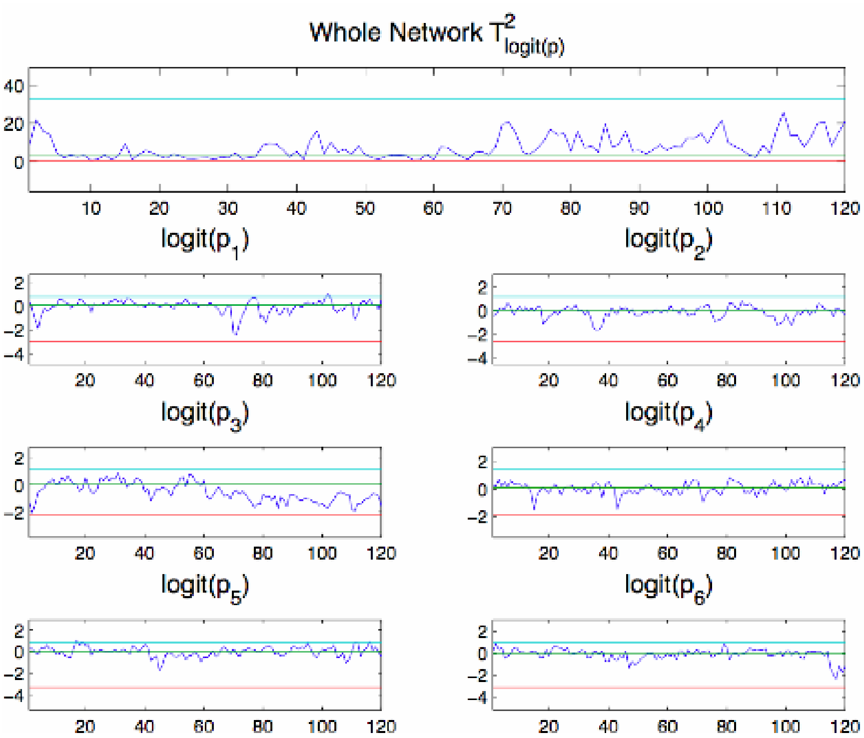
Whole Network $T^2_{logit(p)}$
logit(p_1)
logit(p_2)
logit(p_3)
logit(p_4)
logit(p_5)
logit(p_6)

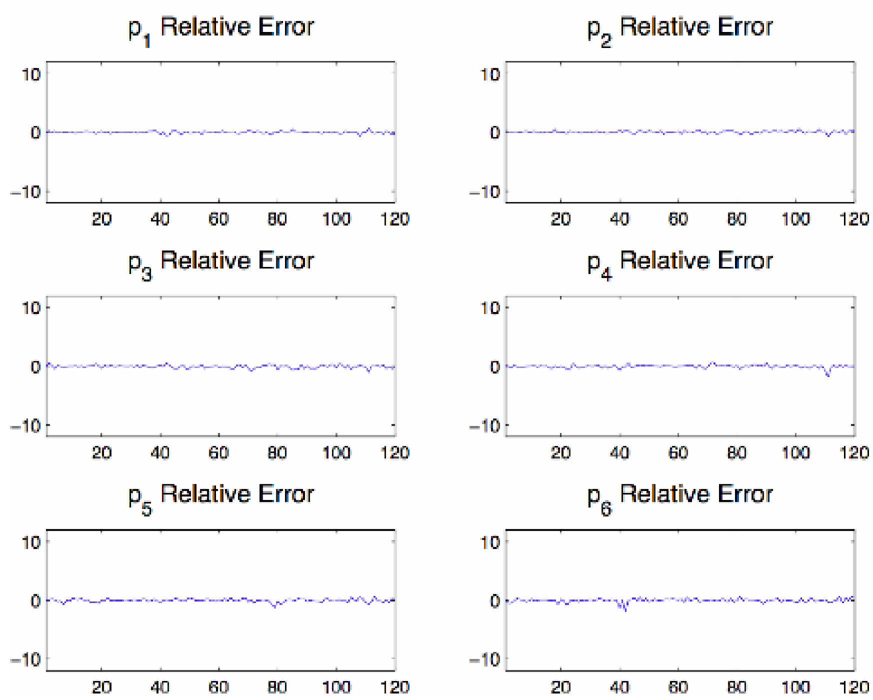
p_1 Relative Error
p_2 Relative Error
p_3 Relative Error
p_4 Relative Error
p_5 Relative Error
p_6 Relative Error

Results

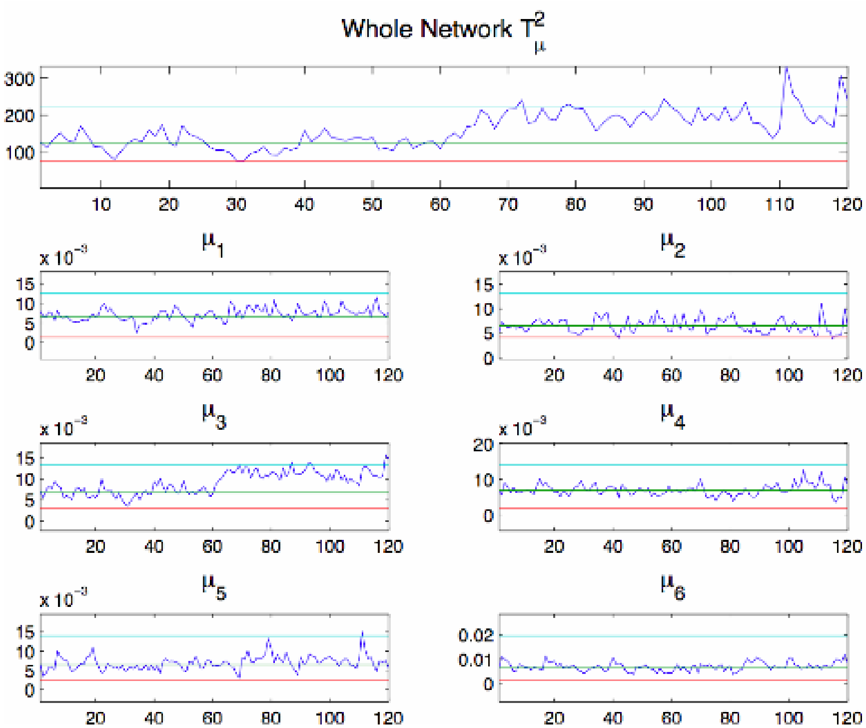
Whole Network T^2_μ
300
200
100
10 20 30 40 50 60 70 80 90 100 110 120
μ_1
x 10^{-3}
15 10 5 0
μ_2
x 10^{-3}
15 10 5 0
μ_3
x 10^{-3}
15 10 5 0
μ_4
x 10^{-3}
20 10 0
μ_5
x 10^{-3}
15 10 5 0
μ_6
0.02 0.01 0
20 40 60 80 100 120

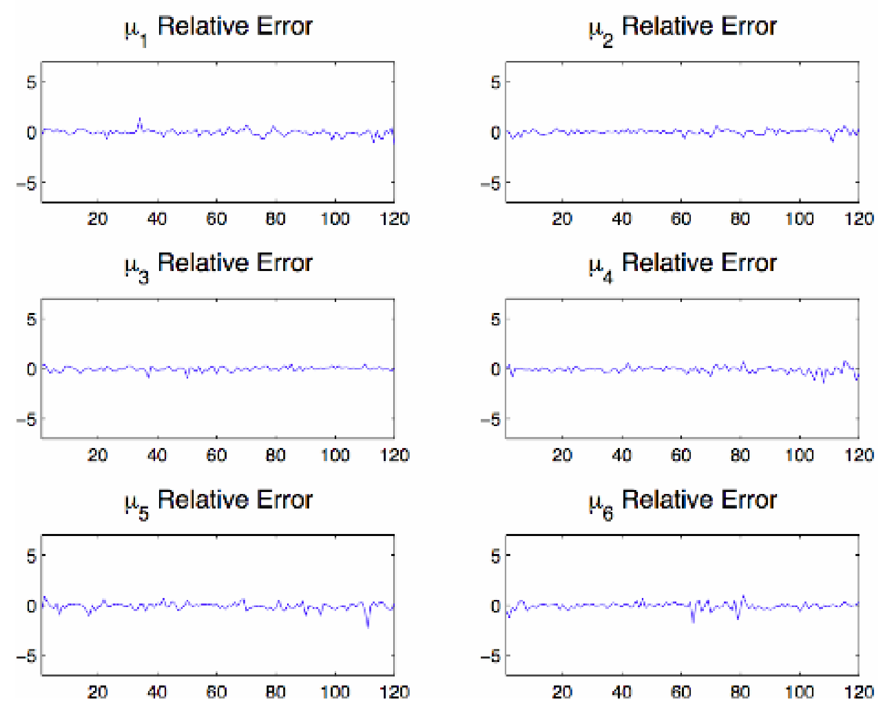
μ_1 Relative Error
μ_2 Relative Error
μ_3 Relative Error
μ_4 Relative Error
μ_5 Relative Error
μ_6 Relative Error
5 0 −5
20 40 60 80 100 120

Mixtures of Phase Distributions

- Let F_k be a mixture of gammas.
- Simple structure that can capture the known heavy-tailed nature of network traffic.
- Can approximate any distributions over the positive real line.
- Identifiable under the general framework or by higher order moments.

Estimation

- The likelihood for exponential mixtures of known order on the simplest tree is already quite crappy:

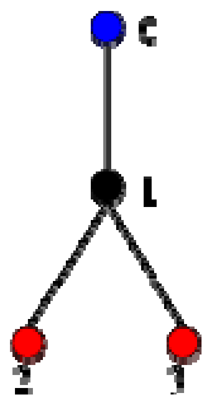

$$
\begin{aligned}
l(\vec{\alpha}, \vec{\lambda}; \mathbf{Y}) \quad = \quad & \textstyle\prod_{j=1}^{k_1} \prod_{m=1}^{k_2} \prod_{q=1}^{k_3} \alpha_{1,j}\alpha_{2,m}\alpha_{3,q} \\
(\quad & N \log(\lambda_{1,j}) + N \log(\lambda_{2,m}) + N \log(\lambda_{3,q}) - N \log(\lambda_{1,j} - \lambda_{2,m} - \lambda_{3,q}) \\
- \quad & \lambda_{2,m} \textstyle\sum_{i=1}^{N} y_{i,2} - \lambda_{3,q} \sum_{i=1}^{N} y_{i,3} \\
+ \quad & \textstyle\sum_{i=1}^{N} \log[1 - \exp\{-(\lambda_{1,j} - \lambda_{2,m} - \lambda_{3,q}) \min(y_{i,2}, y_{i,3})\}] \ )
\end{aligned}
$$

Estimation

- Extensions to larger trees, unknown order, and mixtures of gamma make it difficult to write down a closed form.
- I will consider schemes that only require evaluation of the likelihood.

Bayesian Estimation via Metropolis

- Lots of work on unknown mixing order that can be extended.
- Good complexity (per iteration at least).
- Prior information from previous estimates can feed into a monitoring scenario.
- The people where I work seem to like this Metropolis guy.

Example

- 1000 points from a four component exponential mixture model on each link: $\lambda=[1,.1,.01,.001]$, $\eta=[.4,.3,.2,.1]$.

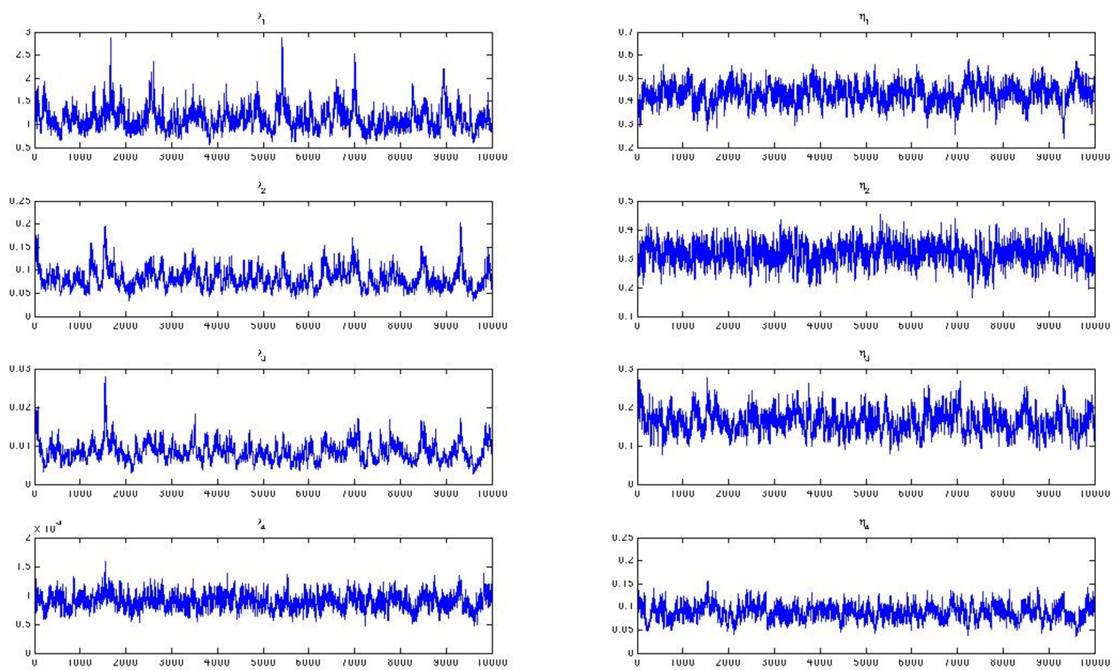

Gaussian Process Approximation

- Treat the likelihood like the output of a complex simulator.
  - Lots of numerical integration, etc.
- Evaluate the likelihood at a set of parameter inputs.
- Fit a Gaussian process to approximate the likelihood.
- Output can be used to maximize or find expected values.

Spatio-Temporal Loss Modeling

- Loss modeling: did the packet reach the receiver (α_k).
- Latent variable modeling:

$$\begin{aligned} Z_k &\sim \mathrm{N}(\mu_k, 1) \\ Z_k > 0 &\Rightarrow X_k = \infty \end{aligned}$$

- Build dependency on the latent variable structure.

Multivariate Probit Modeling

- Consider a vector of latent variables that determine loss for the entire tree:

$$\begin{aligned} \vec{Z} &\sim N_q(\vec{\mu}, \Sigma) \\ \vec{Z} &\Rightarrow \vec{X} \end{aligned}$$

- When *X* is known, there are several techniques.
  - I recommend Lawrence et al. (2007) which uses parameter expansion and Gibbs algorithm.

MVP Loss

- Extend to observed end-to-end values *Y* with an imputation step at each iteration to get *X*.
- Extend to a structured covariance matrix to model temporal dependency.

Summary

- Accurate and fast estimation under simplifying assumptions.
- Poised to have models with flexible structure that incorporate complex dependency.