Network Tomography Based on end-to-end Measurements

Francesco Lo Presti

Dipartimento di Informatica - Università dell'Aquila

Talk Outline

- Motivation
- Network Tomography
- Challenges
- Summary

Tomography

- Noninvasive imaging technique
- X-ray beam directed through one part of the body onto a film/sensors
- Computer reconstructs cross sectional views of the body from data received from the sensors

Network Tomography

- Characterize internal network characteristics from end-to-end packet behavior

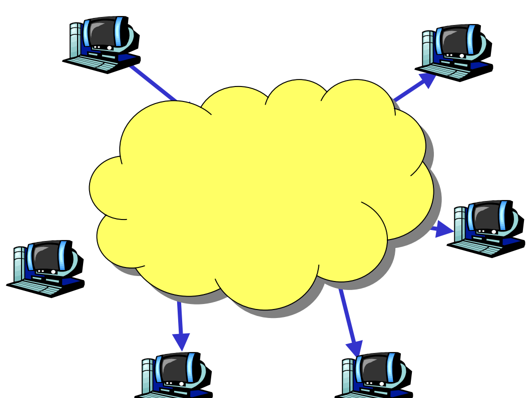

Network Tomography

- Characterize internal network characteristics from end-to-end packet behavior

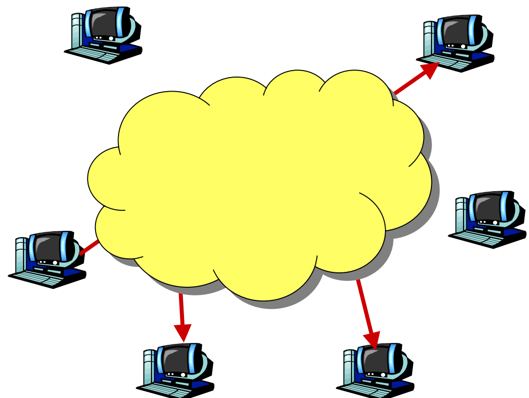

Network Tomography

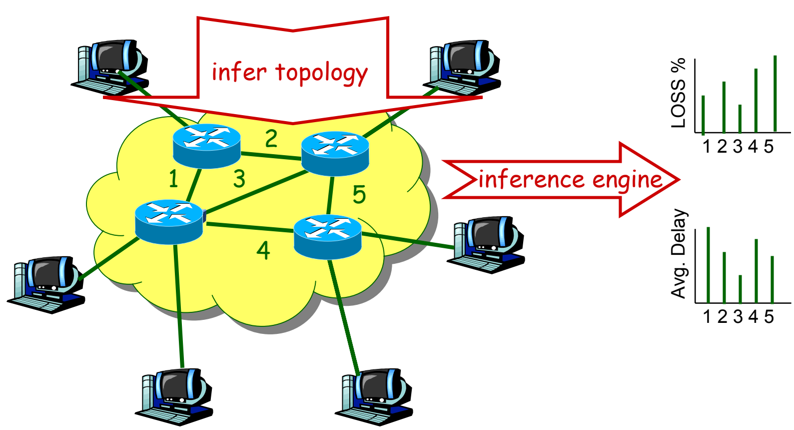

Network Tomography

- Why end-to-end
  - decentralized stateless network
  - ISPs do not share internal measurements
- Applications
  - network performance monitoring
    - network management
    - research
    - performance sensitive applications
- Key Ingredients
  - measurement infrastructure
  - measurement/inference techniques

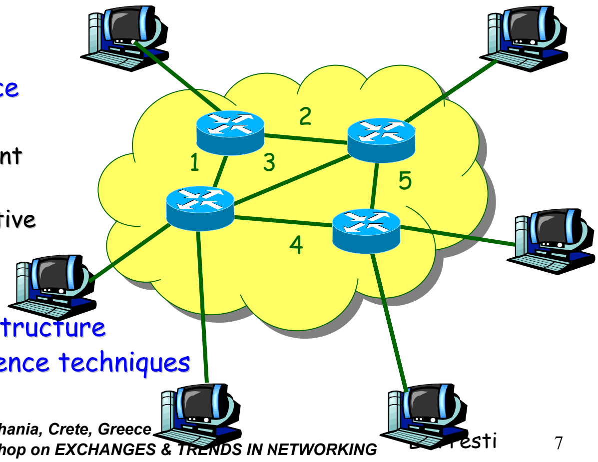

Key Ingredients

- Network Model
  - topology, loss, delay model, etc.
  - parameterized by a suitable vector α
- Goal: Infer α
- Probing Technique
  - active, passive probing
  - unicast, multicast
    - packet pairs, stripes, cartouches, etc.
- Measurements
  - receivers observations X
    - distribution is function of α
- Inference Technique
  - MLE, Bayesian, Method-of-Moments, etc.

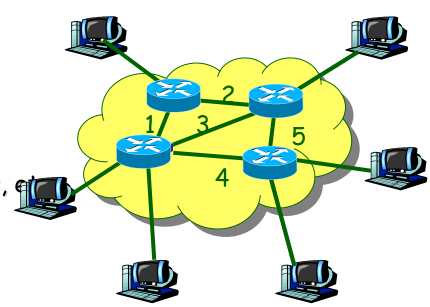

Identifiability

- captures the property that a particular internal network characteristic can be uniquely identified from a given type of measurements (as the # of probes goes to ∞)

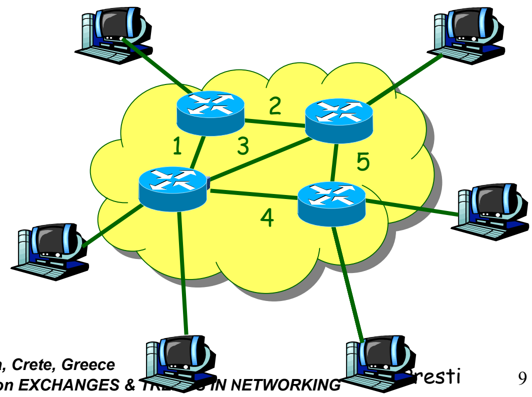

Probing, Measurements and Inference

- Unicast probes
  - infer path characteristics
    - E.g., available bandwidth

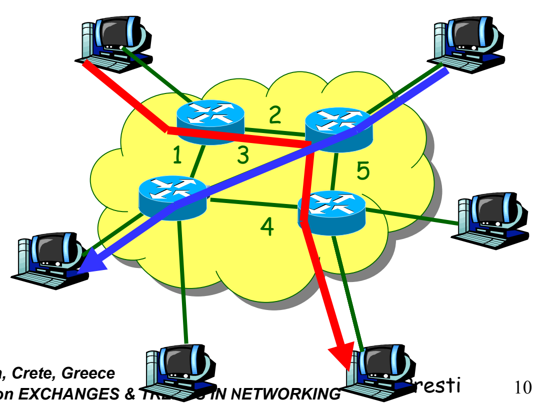

Probing, Measurements and Inference

- Multicast probes
  - ⇒ exploit correlation observed at receivers
  - infer tree segments loss/delay characteristics
    - Segment: Path between source/receivers/branch points
  - logical tree topology
- Assumptions
  - independent loss/delay
    - spatial correlation
      - » bias
    - temporal correlation:
      - » slower convergence
  - stationary behavior

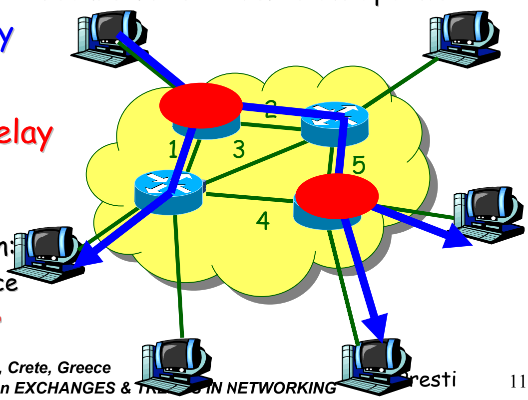

Probing, Measurements and Inference

- Multicast probes & Multiple Trees
  - increase network cover
  - increase # of links which can be identified

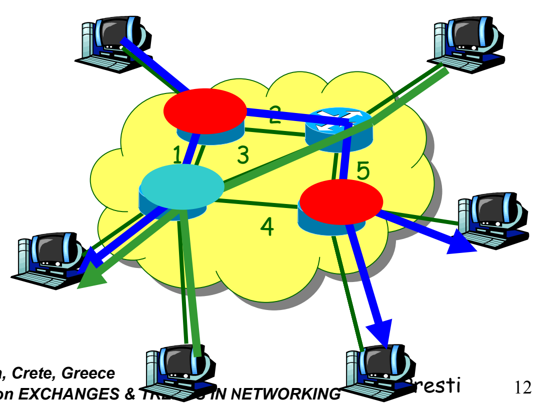

Probing, Measurements and Inference

- Unicast pairs (back-to-back packets to different receivers
  - assimilated to the corresponding 2-recvs tree case
    - but imperfect correlation leads to bias
  - better use longer sequences (Stripes)
- To cover the net use
  - different receivers
  - different senders

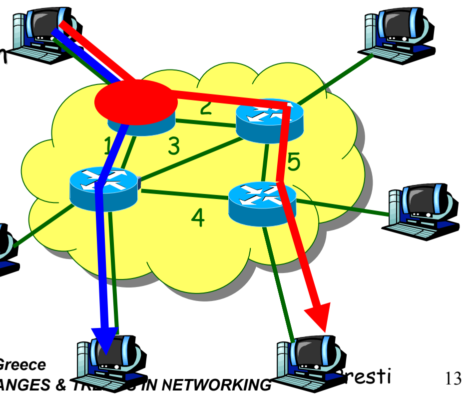

Inference Techniques

- Maximum Likelihood
  - link k behavior characterized by α_k
    - α_k transmission probability
    - α_k delay p.m.f.
  - $L(X;\alpha)$= likelihood of the recvs observations X
  - Estimate $\hat{\alpha} = \text{argmax}_{\alpha} L(X;\alpha)$
    - loss: closed form expressions
    - delay: EM Algorithm
    - topology: exhaustive
- MLE Properties
  - asymptotic consistency
  - asymptotic normality
  - asymptotically efficient
- Bayesian
  - mcast tree topology
- Greedy Heuristics
  - mcast tree topology
- Method of Moments
  - delay cumultants
- ...

Example: Mcast Loss Prob. Estimation

- Estimate segment transmission probability $\alpha_1\ \alpha_2\ \alpha_3$

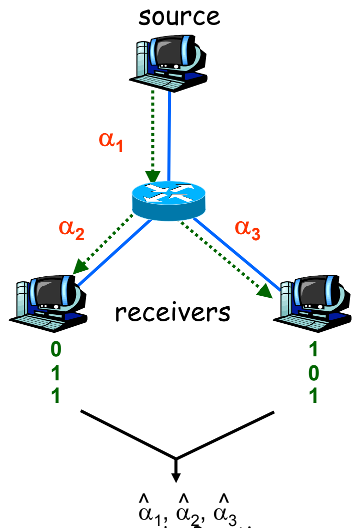

Example: Mcast Loss Prob. Estimation

- Estimate segment transmission probability $\alpha_1\ \alpha_2\ \alpha_3$
- Express $\alpha_1\ \alpha_2\ \alpha_3$ as function of observable events
- Pr[1 1] = $\alpha_1\ \alpha_2\ \alpha_3$
- Pr[1 *] = $\alpha_1\ \alpha_2$
- Pr[* 1] = $\alpha_1\ \alpha_3$
- $\alpha_1 = \dfrac{\text{Pr[1*] Pr[*1]}}{\text{Pr[11]}}$
- ...
- estimate the different Pr[...] by their empirical means
- Substitute above to get the estimate $\hat{\alpha}_1$

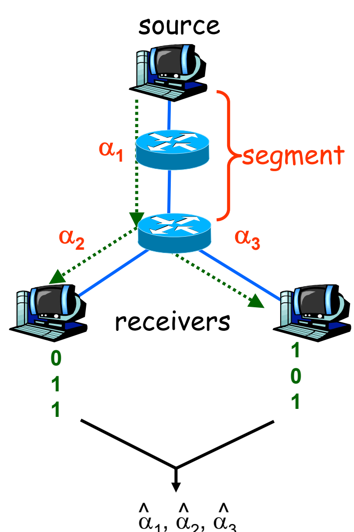

Expectation Maximization (EM) Algorithm

- Computation of $\mathrm{argmax}_{\alpha}\ L(X;\alpha)$ is generally not trivial
  - delay inference, mcast inference with multiple trees, unicast, missing data, etc.
- EM Approach: Augment the observed data with "latent data" so that the resulting MLE is simpler

EM: Delay Inference

- Measurement Experiment:

  send probes and collect delay data (observations) at receivers $X=((X^{(i)}(k))_{k\in R})_{i=1,\ldots,n}$
- Likelihood function:

  $L(X;\alpha)$=Probability of observing X given under the probability law induced by α]=$Pr_{\alpha}[X]$;
  - expressions for $Pr_{\alpha}[X]$ are complex
    - Convolutions
  - finding the maximizer not easy

EM: Delay Inference (cont'd)

- Expectation-Maximization (EM) algorithm: Iterative approach to MLE
- Idea: Assume to know more…so that the MLE becomes simpler
- Assume to observe the entire process
  - for each packet the delay along each link
  - n(k,d)=# of probes which experienced delay=d along link k
- MLE of the new (complete data) problem reduces to MLE of multinomial random variables

$$\widehat{\alpha}_k(d)=n(k,d)/n$$

EM: Delay Inference – E step

1. Choose an initial estimate $\alpha^{(0)}$
2. Expectation Step:

   Given the current estimate $\alpha^{(i)}$ estimate the unknown counts n(k,d) by their conditional expectation given the observation y under the probability law induced by the current estimate $\alpha^{(i)}$

$$\widehat{n}(k,d)=E_{\alpha(i)}[n(k,d) \mid X]$$

   - computationally intensive
   - approaches
     - probability propagation alg.
     - FFT
     - Monte Carlo

EM: Delay Inference – M step

3. Maximization Step:

   Compute the new estimate $\alpha^{(i+1)}$ using the estimated counts n(k,d) in place of the unknown n(k,d)

   - In other words

$$\alpha^{(i+1)}(k,d)=\widehat{n}(k,d)/n$$

4. Iterate Step 2 and 3 until some termination criteria is satisfied. Set $\widehat{\alpha}=\alpha^{(i)}$

- Sequence $L(X;\alpha^{(i)})$ is increasing
- Under mild hypothesis – here satisfied - $\alpha^{(i)}$ converge to a local maximizer of the likelihood of the original (incomplete data) problem

Network Tomography: the Challenge

- Can Network Tomography support online network operation and applications?
  - network monitoring, admission control, SLA monitoring, CDN
- We need to:
  - extend end-to-end Network Tomography
  - tailor tomography to potential applications
  - scalability, scalability, scalability!!!

Network Tomography: Challenges

- Richer yet parsimonious network models
  - overcome simplifying models assumptions
    - correlation, non stationary, second order statistics, etc
- Probing Theory/Framework
  - explore probing techniques space
    - grouping destination, size, timing, TTL, scoping
- Innovative Approaches to Inference
  - statistics theory/alg. tailored to networking

Challenges: Distributed Measurement Infrastructure

- Issues:
  - scale: cooperation, data exchange
  - delay measurements accuracy,
  - we can only infer attributes of a network we can "surround"
- Explore Single Node Measurements
  - relies on TTL expiry techniques
  - no cooperation
  - no data exchange

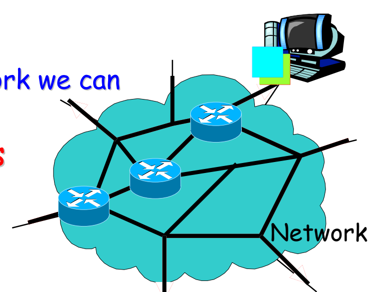

- Allows performance inference along both direction of traversed links
- Intuition: We still rely on a logical tree topology where sender and receivers coincide!!

Application: SLA Monitoring

- Application: SLA Monitoring

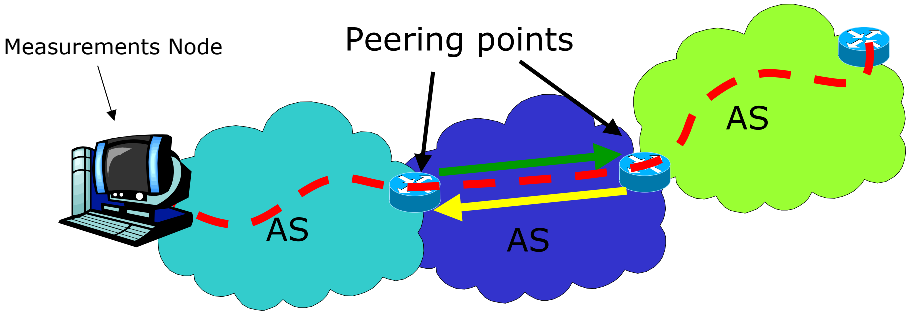

Summary

- Network Tomography based on end-to-end measurements
- Challenges
  - Modeling
  - Probing techniques
  - Inference
  - Measurements