Maximum Likelihood Network Topology Identification from Edge-based Unicast Measurements

Mark Coates
McGill University

Rui Castro, Robert Nowak
Rice University

Presented by Chunling Hu

About the Authors

- Mark Coates: Assistant Professor of ECE at McGill university. His research interest is in communications; computer networks, statistical signal and image processing
- Rui Castro: Ph.D. student of ECE at Rice University. Research interest is network tomography
- Robert Nowak: Associate Professor of ECE at Rice University. Reasearch interest: digital signal processing

Motivation and Main Point

- Motivation
  - Network tomography requires knowledge of network topology
  - Network topology identification is a first critical step in the tomography process
- Main Point
  - Discovering network topology solely from host-based, unicast measurements, without internal cooperation

Introduction

- Performance optimization of high-end applications
- Spatially localized information about network performance
  - Two gathering approaches:
    - Internal: impractical(CPU load, scalability, administration…)
    - External: network tomography
- Cooperative conditions: increasingly uncommon
- Assumption: the routers from the sender to the receiver are fixed during the measurement period, that is, a tree-structured graph

Contributions

- A novel measurement scheme based on special-purpose unicast “”sandwich” probes
  - Only delay differences are measured, clock synchronization is not required
- A new, penalized likelihood framework for topology identification
  - A special Markov Chain Monte Carlo (MCMC) procedure that efficiently searches the space of topologies

Sandwich Probe Measurements

- Sandwich: two small packets destined for one receiver separated by a larger packet destined for another receiver

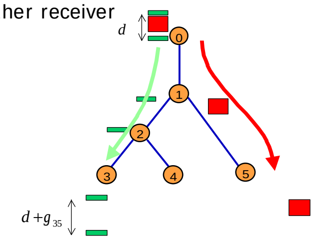

Sandwich Probe Measurements

- Three steps
  - End-to-end measurements are made
  - A set of metrics are estimated based on the measurements
  - Network topology is estimated by an inference algorithm based on the metric

Step 1: Measuring (Pairwise delay measurements)

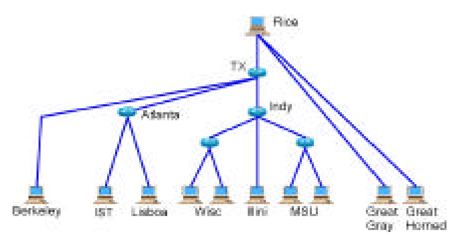

Step 1: Measuring (Continue)

- Each time a pair of receivers are selected
- Unicast is used to send packets to receivers
- Two small packets are sent to one of the two receivers
- A larger packet seperates the two samll ones and is sent to the other receiver
- The difference between the starting times of the two small packets should be large enough to make sure that the second one arrives the receiver after the first one
- Cross-traffic has a zero-mean efffect on the measurements (d is large enough)

Step 1: Measuring (Continued)

- γ_{35} is resulted from the queuing delay on the shared path

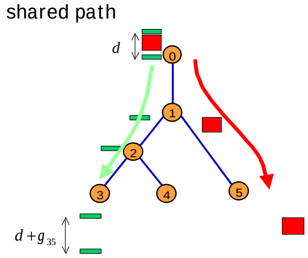

Step 1: Measuring (Continued)

- More shared queues→ larger γ : $\gamma_{34} > \gamma_{35}$

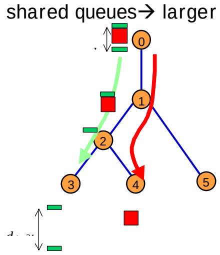

Step 2: Metric Estimation

- More measurements, more reliable the logical topolody identification is.
- The choice of metric affects how fast the percentage of successful identification improves as the number of measurements increases
- Metrics should make every measurement as informative as possible
- Mean Delay Differencies are used as metrics
  - Measured locally
  - No need for clock synchronization

Step 2: Metric Estimation(Continued)

- The difference between the arrival times of the two small packets at the receiver is related to the bandwidth on the portion of the path shared with the other receiver
- A metric estimation is generated for each pair of receivers.

Step 2: Metric Estimation(Continued)

- Formalization of end-to-end metric construction
  - N receivers → N(N-1) different types of measurements
  - K measurements, independent and identically distributed
  - (k) – difference between arrival times of the 2 small packets in the kth measurement
  - Get the sample mean and sample variance of the measurement for each pair (i,j): $x_{i,j}$ and $_{i,j}^2$

(Sample mean of sample $\boldsymbol{X} = (X_1, X_2, \ldots)$ is
$M_n(\boldsymbol{X}) = (X_1 + X_2 + \cdots + X_n) / n$ (arithmetic mean)
Sample variance is $(1/n)$Ó$_{i=1..n}$ $(X_i$ ì$)^2$
$E(M_n) =$ ì
)

Step 3: Topolody Estimation

- Assumption: tree-structured graph
- Logical links
- Maximum likelihood criterion: find the true topolody tree T* out of the possible trees (forest) F based on x
- Probability model for delay difference
  - Central Limit Theorem → $x_{i,j} \sim N(\ _{i,j}, \ _{i,j}/n_{i,j})$
  - $_{i,j}$ is the the theoretical value of $x_{i,j}$
  - That is, sample mean be approximately normally distributed with mean $_{i,j}$ and variance $_{i,j}/n_{i,j}$
  - The larger $n_{i,j}$ is, the better the approximation is.

Step 3: Topolody Estimation(Continuted)

- Probability desity of x is p(x|T, (T)), means (T) is computed from the measurements x
- Maximum Likelihood Estimator (MLE) estimates the value of (T) that maxmizes p(x|T, (T)), that is, $\hat{\mu}(T)$
- Log likelihood of T is

  $$L(x|T) = \log p(x|T, \hat{\mu}(T)).$$

- Maximum Likelihood Tree (MLT) T*

  T* = argmax $_{T^a F}$ $\log p(x|T, \hat{\mu}(T))$

Step 3: Topolody Estimation(Continuted)

- Over fitting problem: the more degrees of freedom in a model, the more closely the model can fit the data
- Penalized likelihood creterion:

  $L_\lambda(x|T)$ $\hat{T}_\lambda = \max_{T \in F} L_\lambda(x|T) - \lambda n(T)$

  - Tradeoff between fitting the data and controlling the number of links in the tree
- Maximum Penalized Likelihood Tree(MPLT) is

  $$\hat{T}_\lambda = \max_{T \in F} L_\lambda(x|T)$$

Finding the Tallest Tree in the Forest

- When N is large, it is infeasible to exhaustively compute the penalized likelihood value of each tree in F.
- A better way is concentrating on a small set of likely trees
- Given $p(T) \propto \exp(-\lambda n(T))$

  $\exp(L_\lambda(x|T)) = e^{-\lambda n(T)} p(x|T,\mu) \propto p(T,\mu|x)$
- Posterior density $p(T,\mu|x)$ $= p(T)$ x $p(x|T,\mu)$ can be used as a guide for searching F.
- Posterior density is peaked near highly likely trees, so stochastic search focuses the exploration

Stochastic Search Methodology

- Reversible Jump Markov Chain Monte Carlo
  - Target distribution: $p(T,\mu|x)$
  - Basic idea: simulate an ergodic markov chain whose samples are asymptotically distributed according to the target distribution
  - Transition kernel: transition probability from one state to another
  - Moves: birth step, death step and step

Birth Step

- A new node l* is added→ extra parameter $_l$*
- The dimension of the model is increased
- Transformation (non-deterministic)
  - $_l$* = r x min($_c$(l,1), $_c$(l,2))
  - '$_c$(l,1) = $_c$(l,1) – $_l$*
  - '$_c$(l,2) = $_c$(l,2) - $_l$*

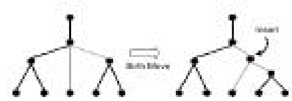

Death Step

- A node l* is deleted
- The dimension of the model is reduced by 1
- Transformation (deterministic)
  - ${}_c(l,1) \;=\; {}'_c(l,1) + {}_l{}^*$
  - ${}_c(l,2) \;=\; {}'_c(l,2) + {}_l{}^*$

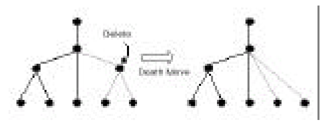

–step

- Choose a link l and change the value of ${}_l$
- New value of ${}_l$ is drawn from the conditional posterior distribution

The Algorithm

- Choose a starting state s_0
- Propose a move to another state s_1
  - Probability = [illegible]
- Repeat these two steps and evaluate the log-likelihood of each encountered tree
- Why restart?

Penalty parameter

- Penalty = $1/2\log_2 N$
- N: number of receivers

Simulation Experiments

- Compare the performance of DBT(Deterministic Binary Tree) and MPLT
- Penalty = 0 (both will produce binary trees)
- 50 probes for each pair in one experiment, 1000 independent experiments
- When the variabilities of the delay difference measurements differ on different links, MPLT performs better than DBT
- Maximum Likelihood creterion can provide significantly better identification results than DBT

ns Experiment

- Topology used for the experiment

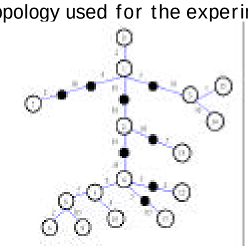

Experiment Results

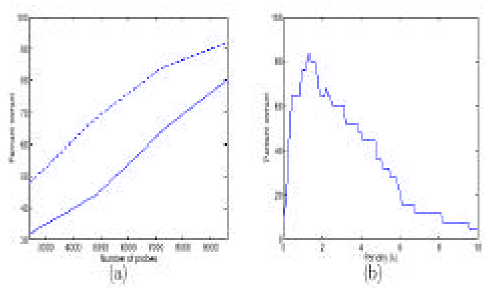

Internet Experiment

- Source host: data collection and inference
- Receivers: a low overhead receiver task
- 8 minutes/experiment, 6 independent experiments
- 1 sandwidth probe / 50ms
- Penalty = 1.7
- topology

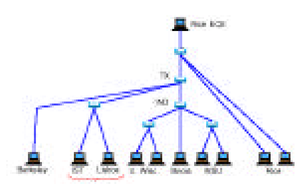

Experiment Result

- Estemated topology

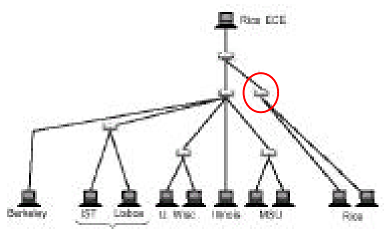

Conclusion and Future work

- Conclusions:
  - Delay-based measurement without the need for synchronization
  - MCMC algorithm to explore forest and identify maximum (penalized) likelihood tree
  - Foundation for multi-sender topology identification
  - Localization of layer-two elements
- Future work
  - Adaptive methods for selecting penalty parameter
  - Adaptivity in the probing scheme

End

Thank you!