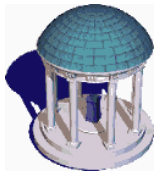

Tomography with Available Bandwidth

Alok Shriram **Jasleen Kaur**

Department of Computer Science
University of North Carolina at Chapel Hill

http://www.cs.unc.edu/~jasleen/research/

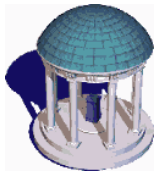

Network tomography

- Idea: Use end-to-end probes to estimate state of internal links
  - Send simultaneous probes to destinations that share portions of their path
  - Study correlations in end-to-end metrics

$$loss_{XY} = 1 - (1 - loss_{XR})(1 - loss_{RY})$$
$$loss_{XZ} = 1 - (1 - loss_{XR})(1 - loss_{RZ})$$

If $loss_{XY} = loss_{XZ} = L$,
then $loss_{RY} = loss_{RZ} = 0$,
and $loss_{XR} = L$

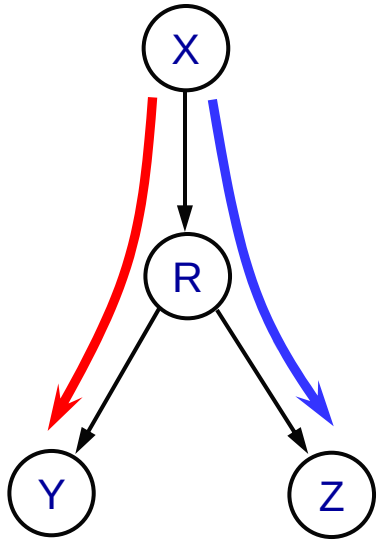

- Past tomographic work done mostly with delay and loss

Focus: Available Bandwidth

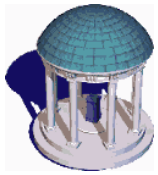

Extending tomography to A.B.

$AB_{XY} = \min(AB_{XR}, AB_{RY})$

$AB_{XZ} = \min(AB_{XR}, AB_{RZ})$

$\Rightarrow \quad AB_{XR} \geq \max(AB_{XY}, AB_{XZ})$

$AB_{RY} \geq AB_{XY}$

$AB_{RZ} \geq AB_{XZ}$

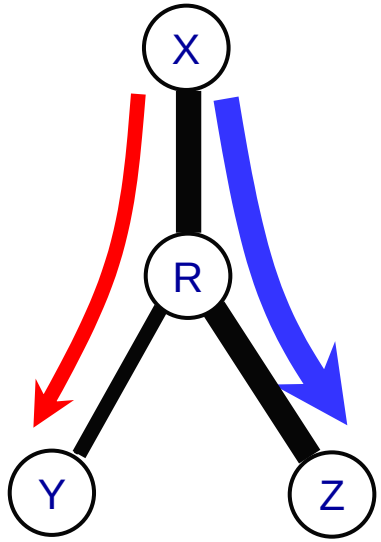

Tomography with several sources and destinations may help identify <u>multiple</u> bottlenecks on end-to-end paths

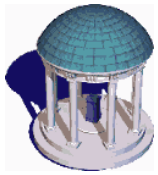

Identifying bottleneck links

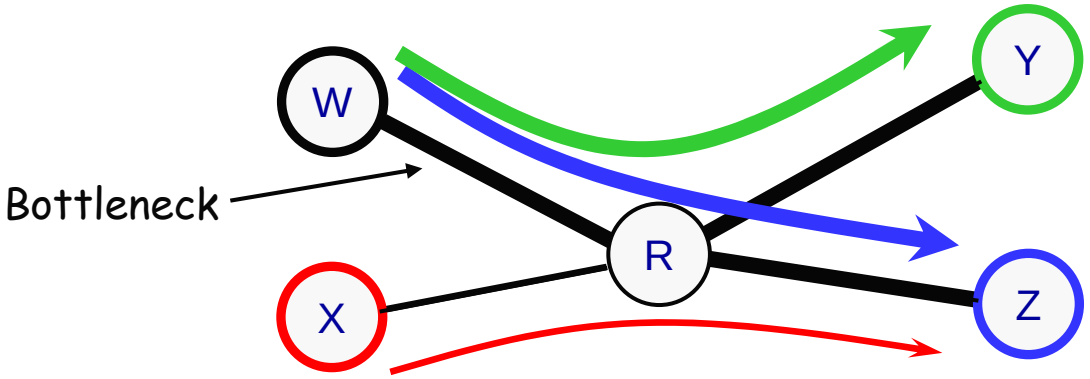

Bottleneck identification rules:

- Rule 1: For each path,
  - Links with the least A.B. are potential bottlenecks

    Could lead to false positives

- Rule 2: For every pair of 2 paths with **equal** end-to-end A.B.
  - Non-shared links are non-bottlenecks

    Could lead to false negatives

Challenge 1: probing tool inconsistency

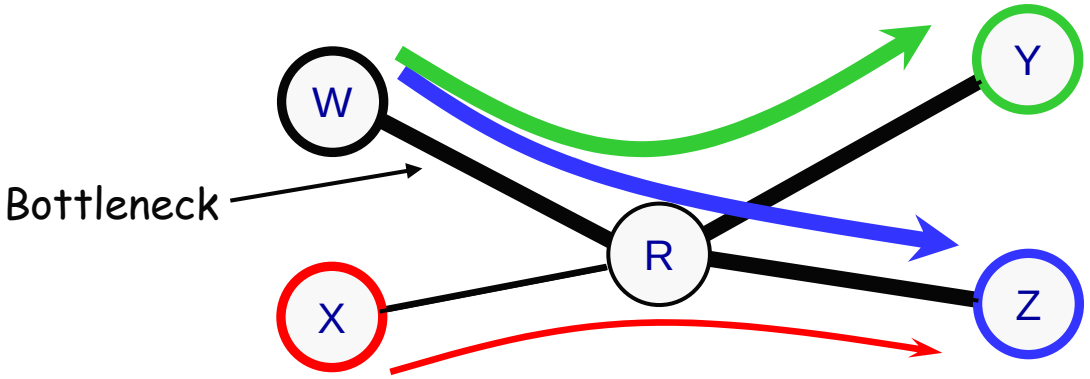

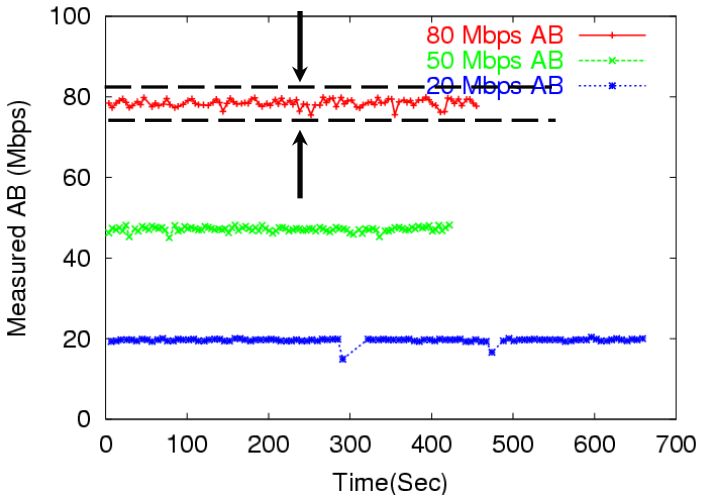

Inconsistency of the probing tool

Tool inconsistency limits the ability to distinguish between bottleneck links

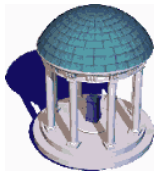

Challenge 2: probe scheduling

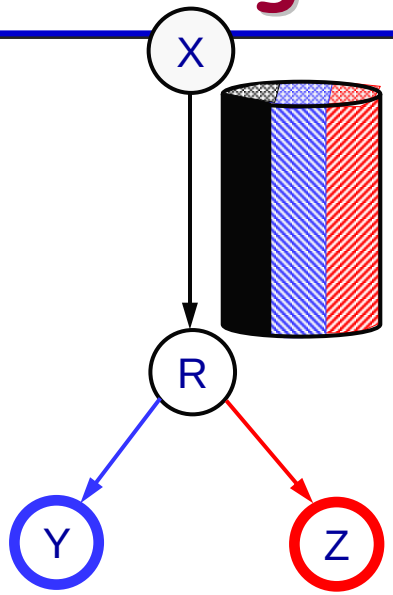

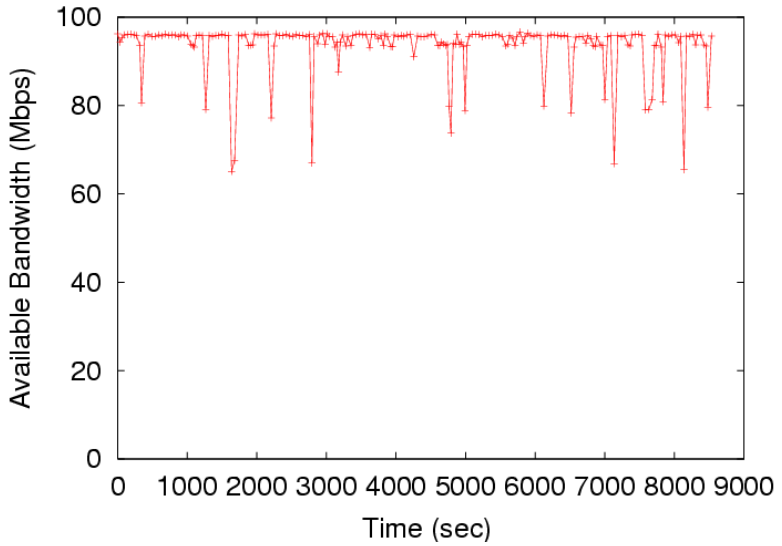

- Requirements:
  - Paths that share links should not be probed concurrently
  - Paths that share links should be probed concurrently
- Solution:
  - Schedule link-sharing probes in separate steps
  - Minimize the total number of steps used

This scheduling problem is NP-hard!

Challenge 3: limit on topology

Available bandwidth dynamics

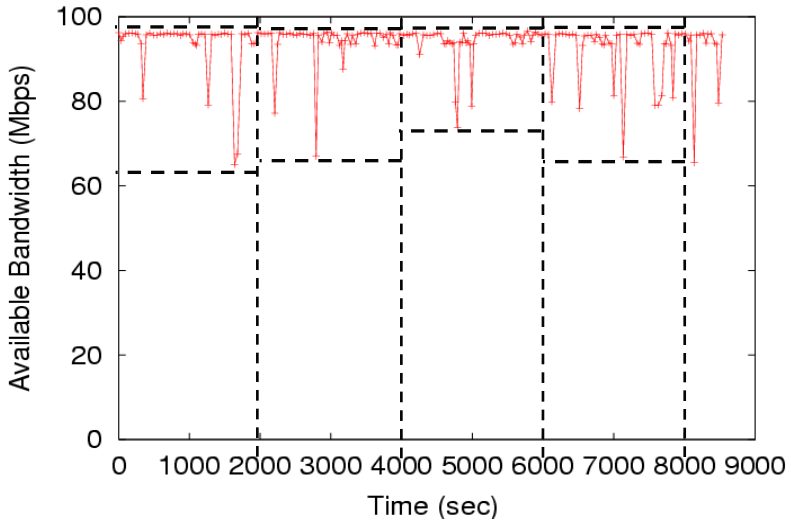

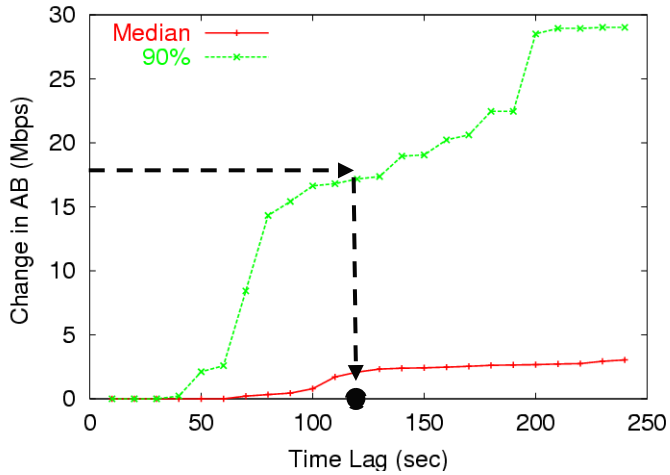

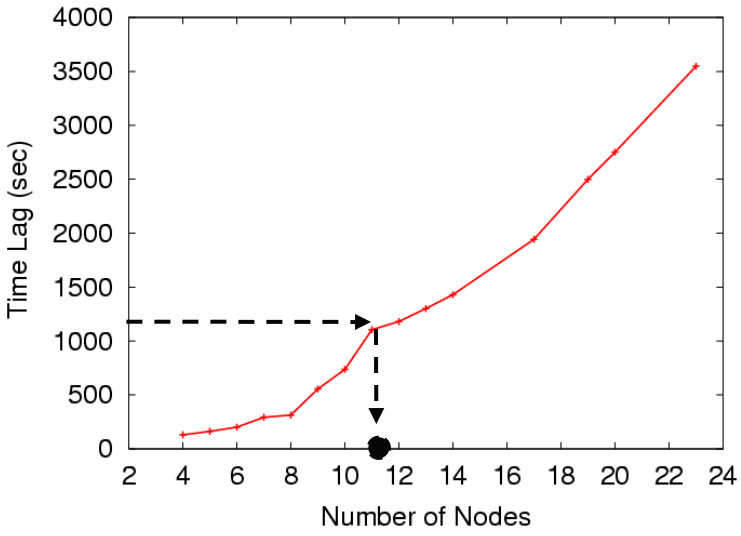

Tool run-time limits the number of participating end-nodes

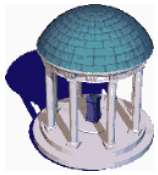

Hurdles in identifying bottleneck links

- Tool inaccuracy
  - Limits the accuracy of detecting bottlenecks
- Tool run-time and dynamics of available bandwidth
  - Limit the number of participating end-nodes
  - Limit the reduction in false positives

PlanetLab measurements with Pathload: ~ 4 end-nodes

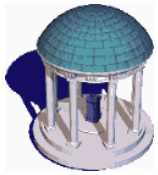

PlanetLab tomography results

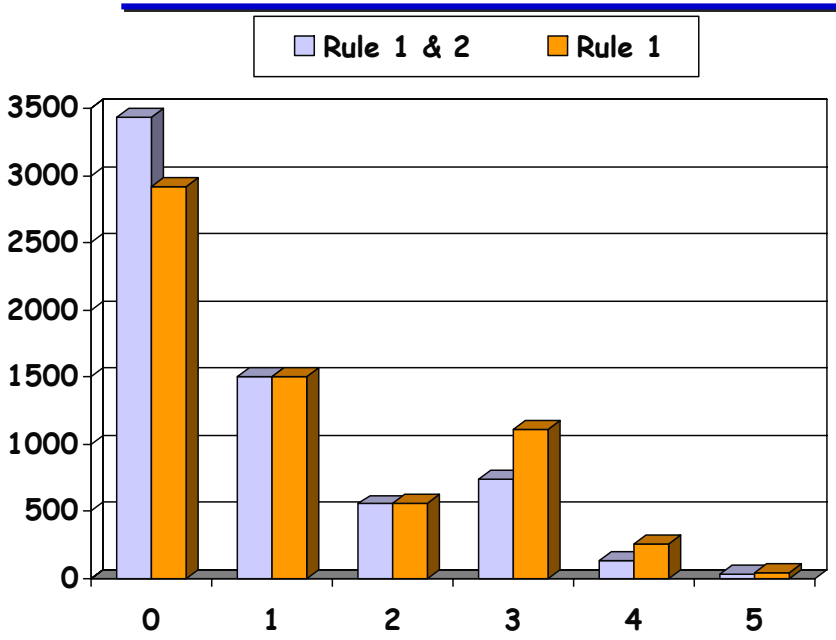

Number of bottlenecks per path

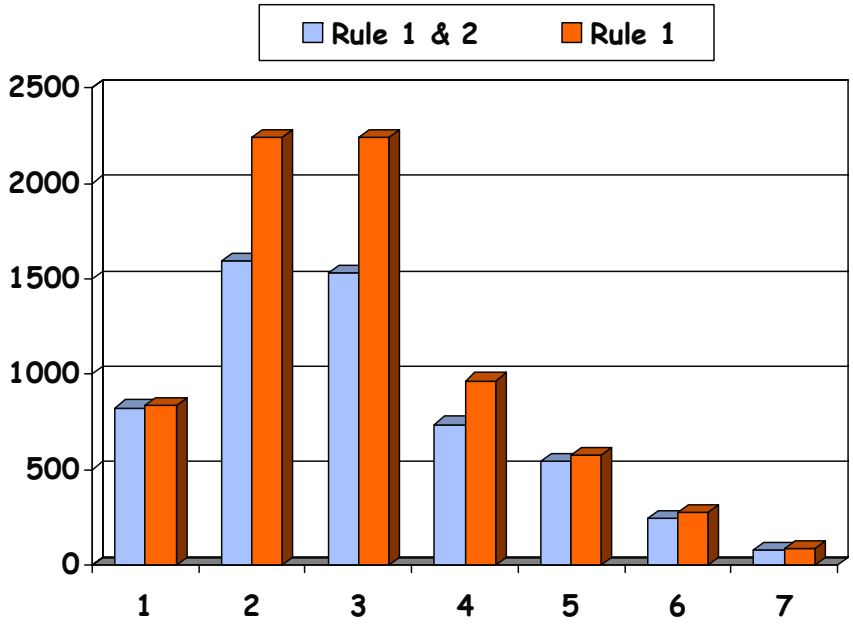

Distance of bottlenecks from source

- Results from 4 sets of 4-node PlanetLab topologies used
  - At least 1 bottleneck discovered on half the paths
  - No more than 3 bottlenecks listed for most paths
  - Most bottlenecks lie at 2-3 hops from the source

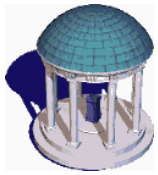

Wish-list for a probing tool

- High accuracy and consistency
  - Within 1 Mbps?
- High speed
  - Within 1 sec?
- Non-interference
  - With cross-traffic
  - With concurrent probing tools?