The End-to-End Effects of Internet Path Selection

S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson

Presented by: Imranul Hoque

Systems Research these days

- "It's just a lot of measurement; a misinterpretation and misapplication of the scientific method."
- "Invention has been replaced by observation."
- "They may be interesting, they may even be relevant, but they aren't research."

- Rob Pike

Question to answer

- How good is Internet routing from a user's perspective, and why?
- Goodness of Internet routing:
  - Round trip time
  - Loss rate
  - Bandwidth
- Why Internet routing might not be good?

How would users react to increased RTT?

Internet routing

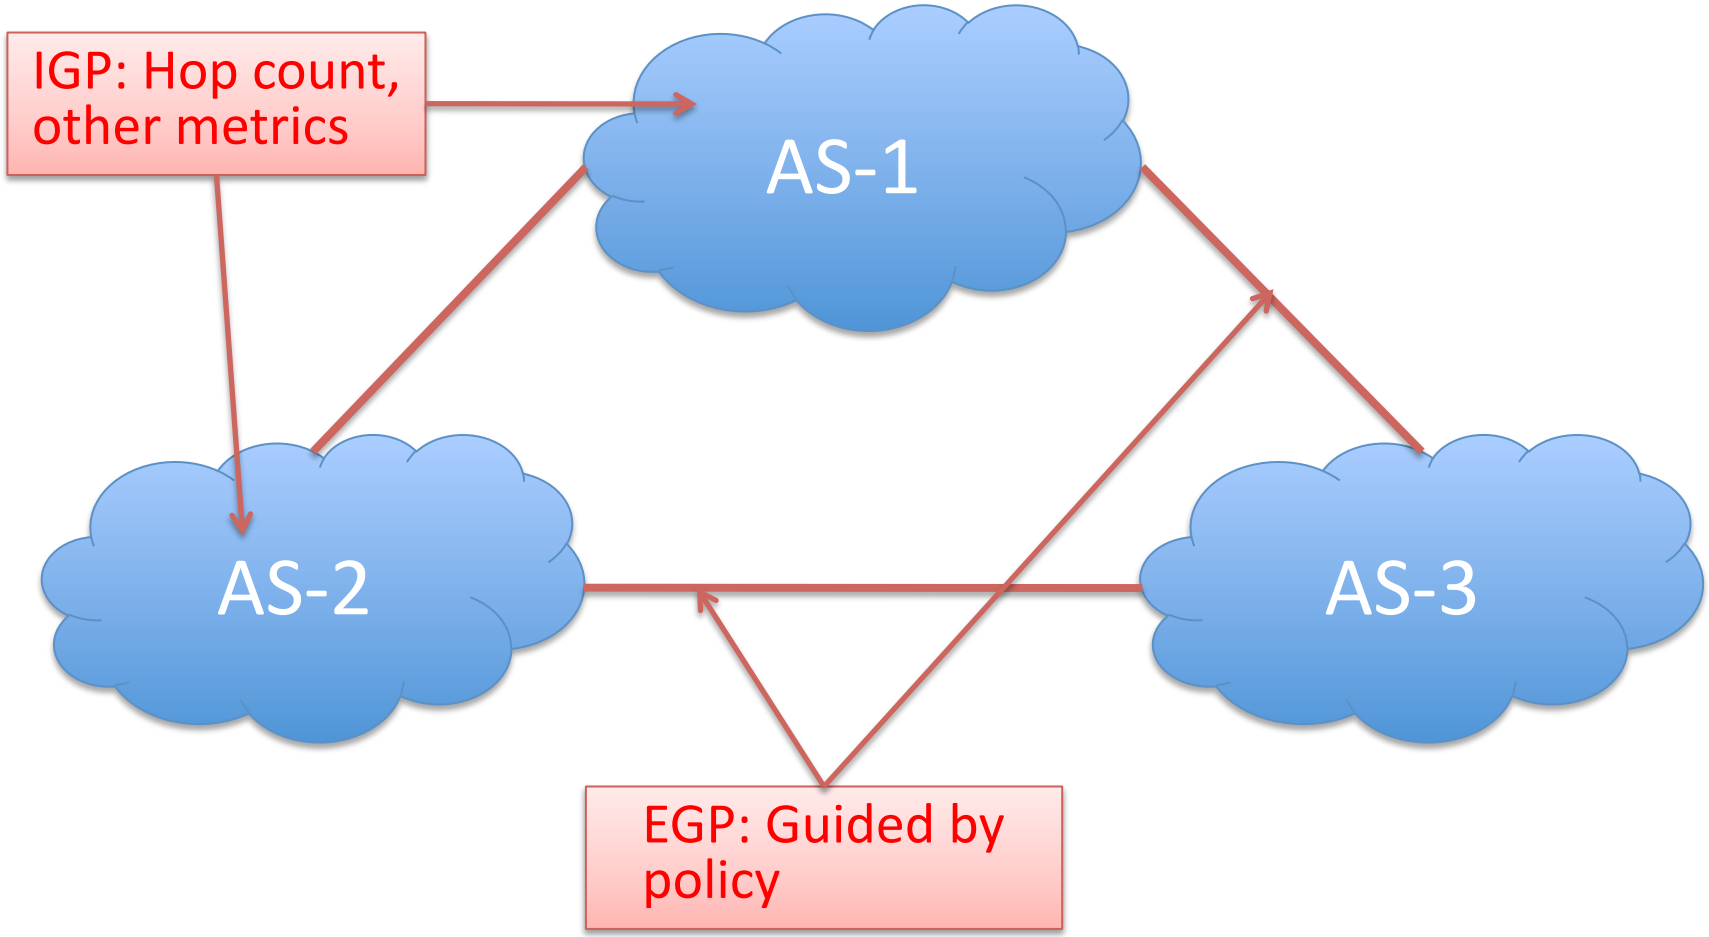

Topics

- Methodology
- Results
  - In 30-80% cases, there exists a superior alternate path
- Evaluating the results
- Using the results
  - Resilient overlay networks
  - Network coordinate systems

Methodology

- Quantifying the impact of path selection:
  - Let X = performance of default path
  - Let Y = performance of best path
  - (Y – X) = cost of using default path
- Technical problems:
  - How to find the best path?
  - How to measure the best path?

Approximating the best path

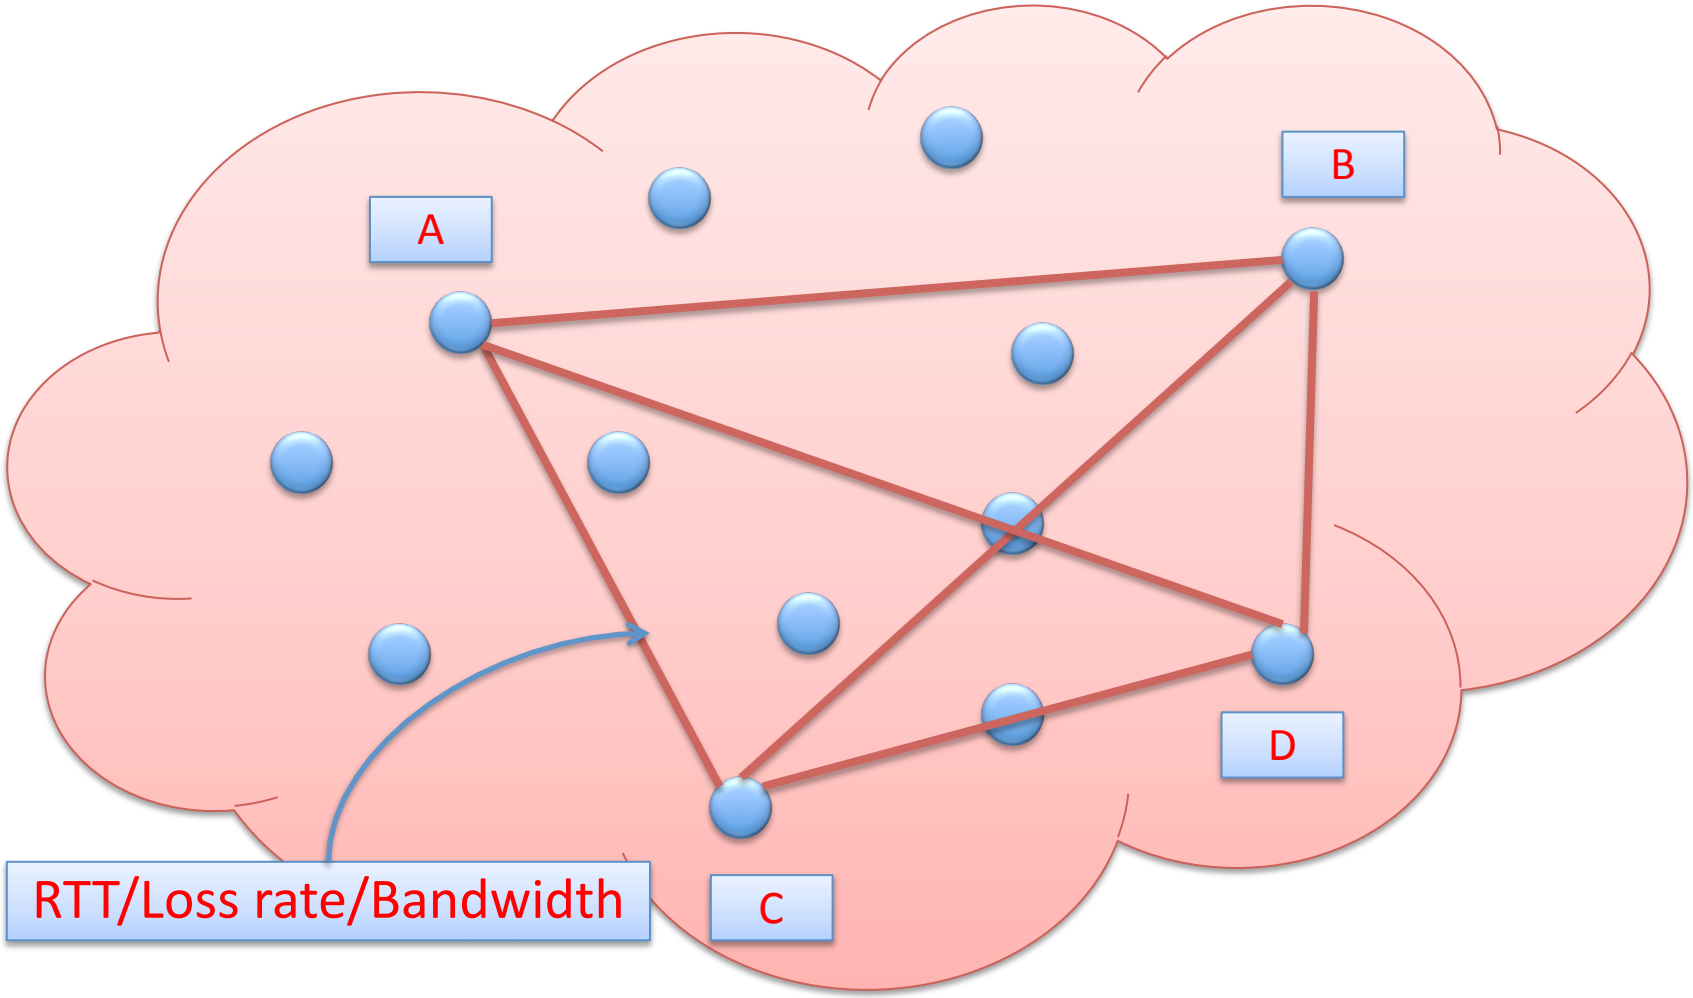

Input trace data

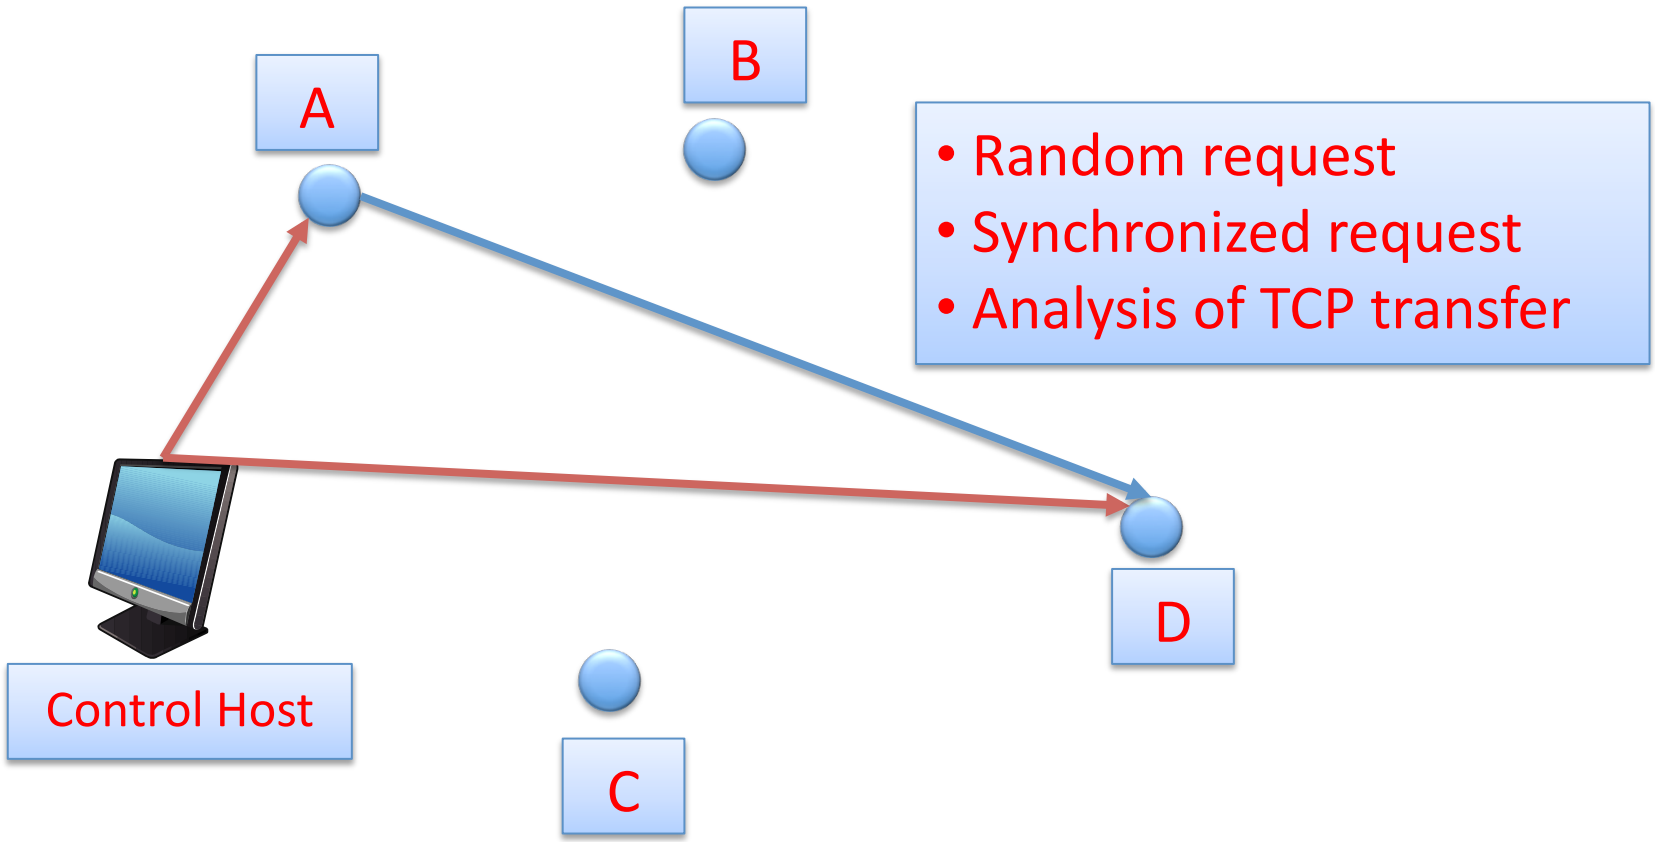

Round trip time

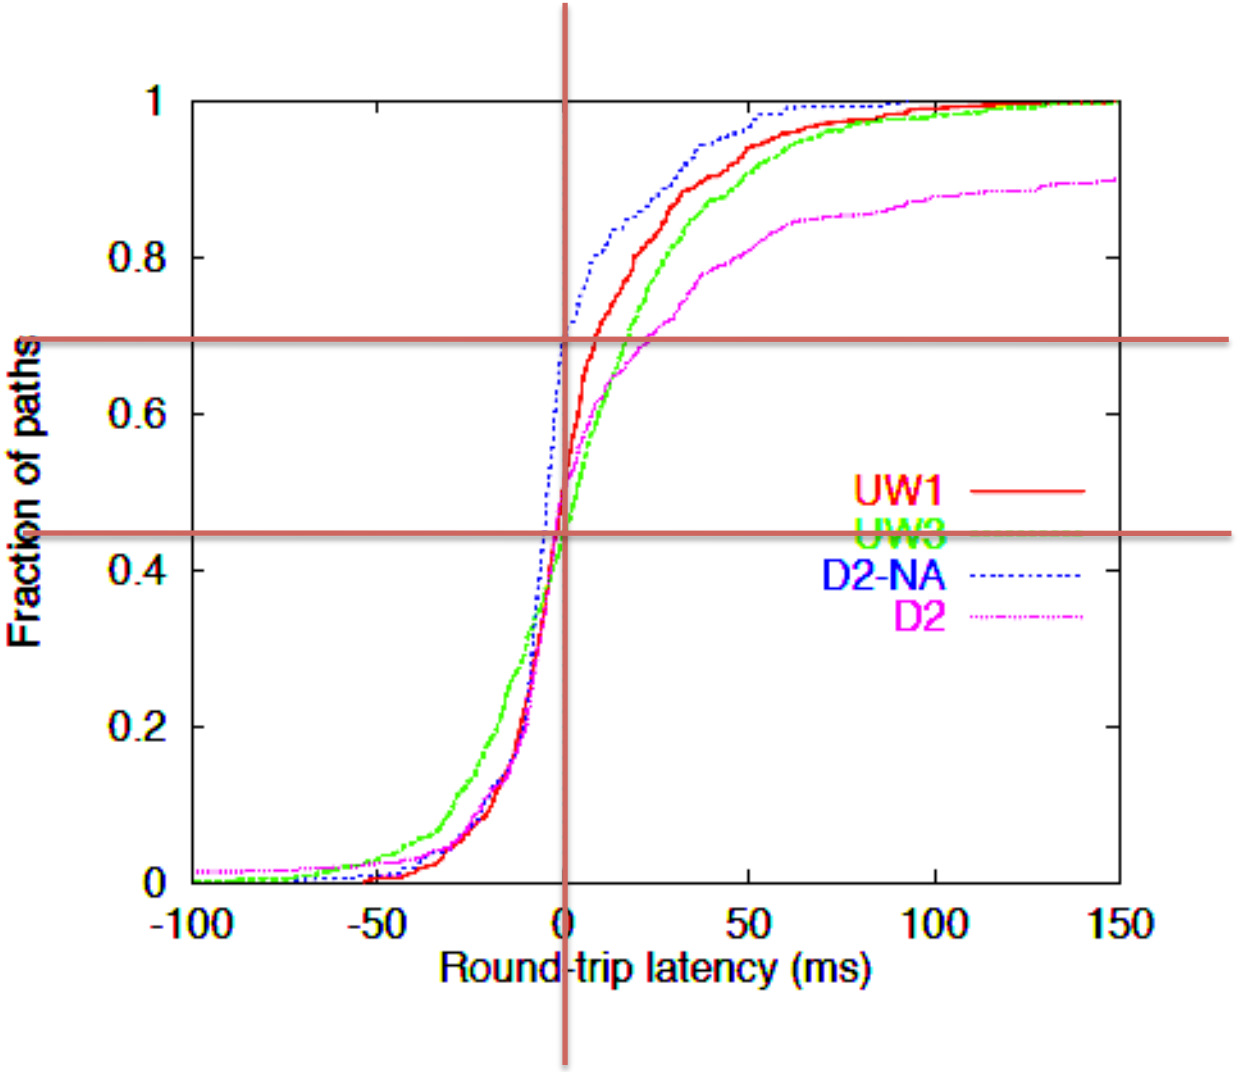

Loss rate

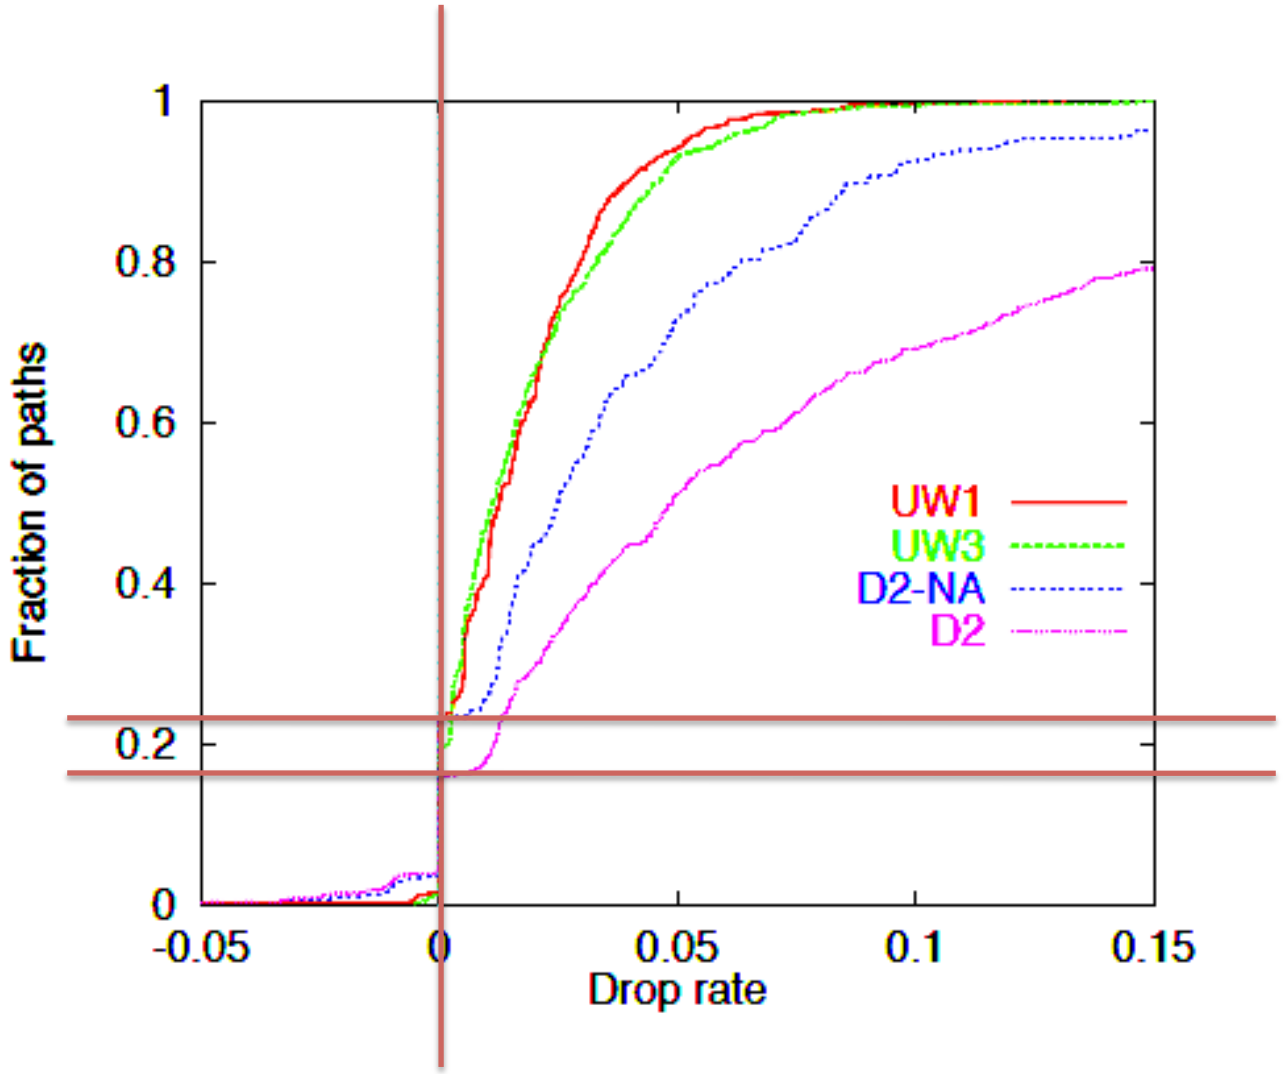
Fraction of paths
1
0.8
0.6
0.4
0.2
0
-0.05
0
0.05
0.1
0.15
Drop rate
UW1
UW3
D2-NA
D2

Bandwidth

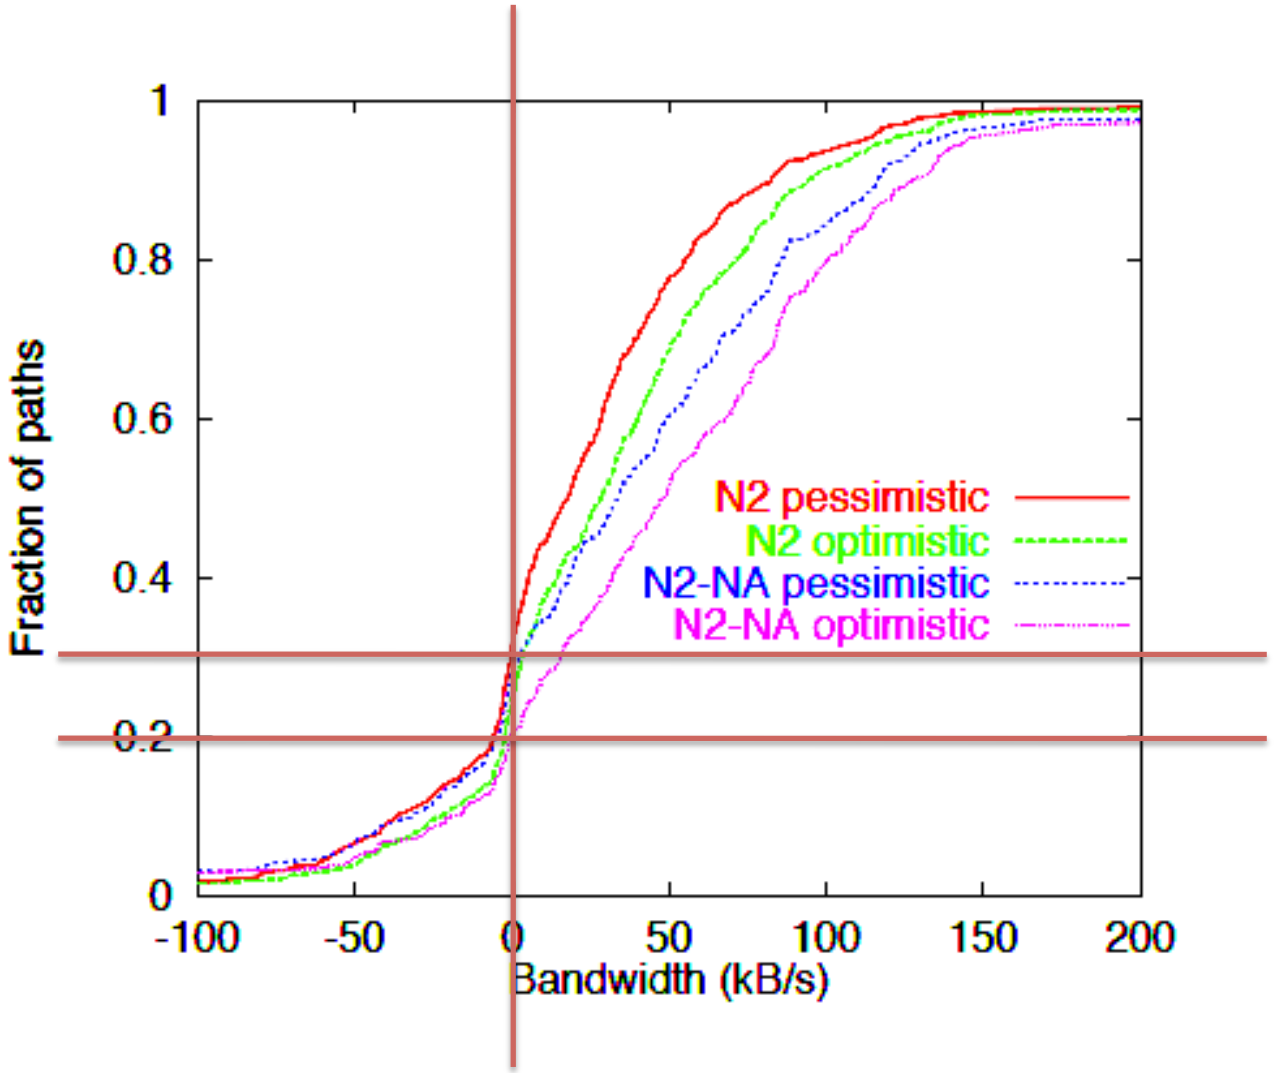
1
0.8
0.6
0.4
0.2
0
Fraction of paths
-100
-50
0
50
100
150
200
Bandwidth (kB/s)
N2 pessimistic
N2 optimistic
N2-NA pessimistic
N2-NA optimistic

Summary of results

- Default path is usually not the best
  - True for latency, loss rate, and bandwidth
  - In spite of synthetic end-host transiting
- Many alternate paths are much better
- Effect stronger during peak hours
- Better path can be:
  - Shorter, less congested, or both
  - Obtained by avoiding parts of the Internet

What makes a better path better?

- Possibilities:
  - Avoid congested queues
  - Shorter propagation delay
- Answer seems to be: both
- Visualizing propagation and congestion:
  - Estimate propagation delay (10^{th} percentile)
  - Queuing delay = RTT – propagation delay
  - Plot improvement in propagation delay vs. improvement in RTT

Propagation delay vs. congestion

- R1 = P1 + Q1 [default path]
- R2 = P2 + Q2 [best alternate path]
- We are plotting (P1 – P2) vs. (R1 – R2)
- R1 > R2:
  - P1 == P2, (Q1 – Q2) == (R1 – R2) [x-axis]
  - Q1 == Q2, (P1 – P2) == (R1 – R2) [y=x line]
  - P1 > P2, Q1 > Q2, (P1 – P2) < (R1 – R2)
  - Q1 < Q2, P1 > P2, and (P1 – P2) > (R1 – R2)
  - P1 < P2, Q1 > Q2, and (Q1 – Q2) > (R1 – R2)

Propagation delay vs. congestion

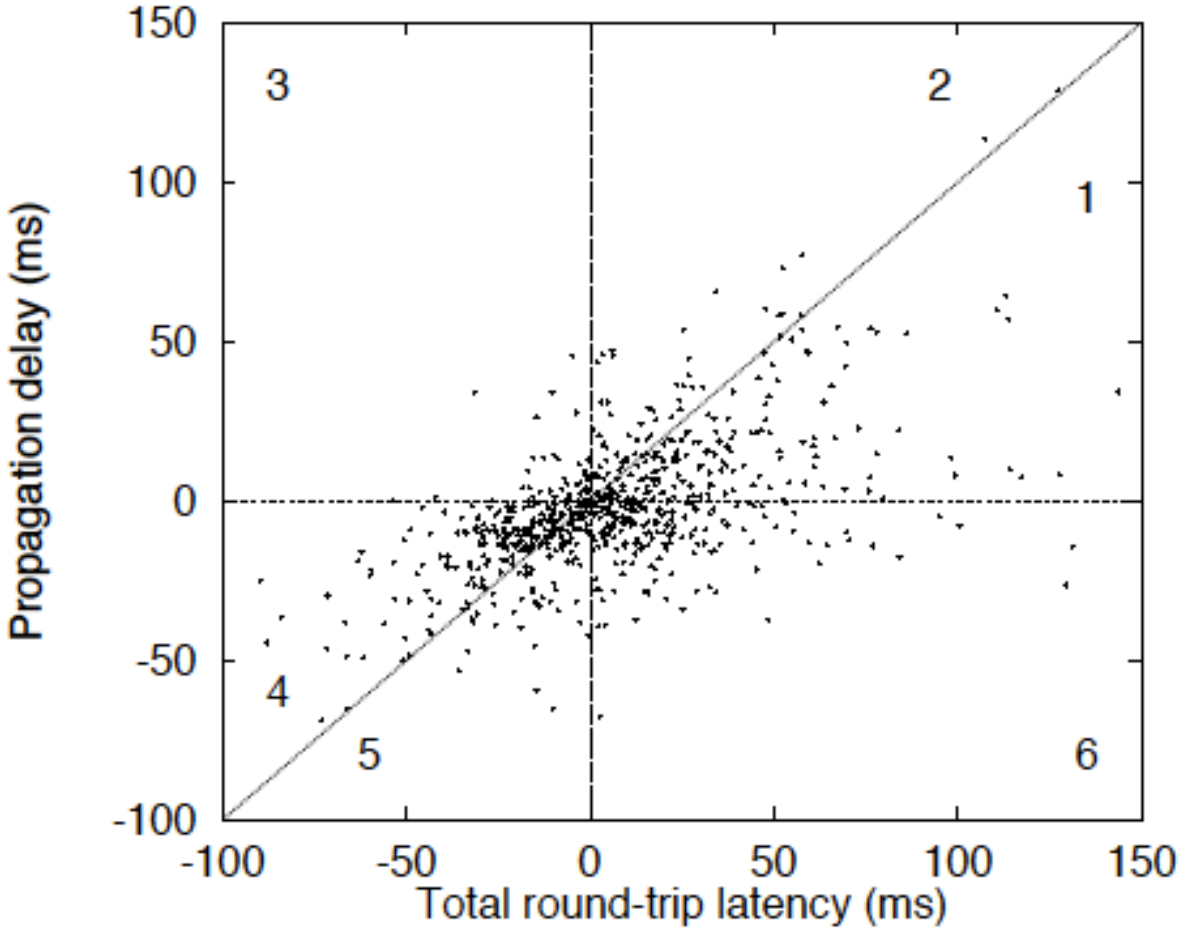

Summary

- Alternate path is better than the default path
  - So what?
- Implications of the result:
  - Resilient overlay networks (RON)
  - Network coordinate systems

Resilient overlay networks

- Routing using overlay network
- Monitor links in the overlay
  - RTT, loss rate, etc.
- Aggressively make routing decisions based on the observations
- Problems?

Network coordinate systems

- Assign coordinates to hosts to enable latency prediction
- Example: GNP, NPS, PIC, PCoord, Vivaldi, etc.

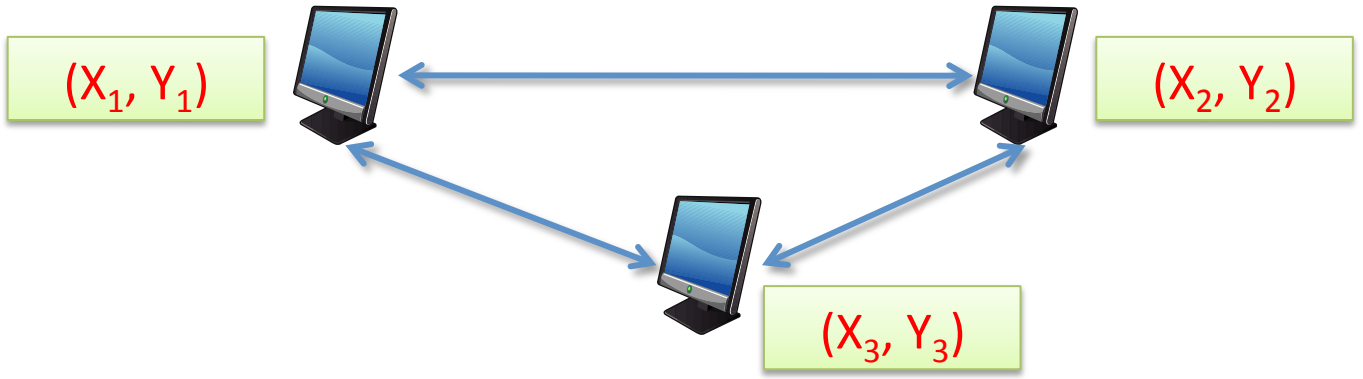

There might be a shorter path through another node

Questions?