Network Positioning from the Edge

David R. Choffnes, **Mario A. Sánchez,**
Fabián E. Bustamante
EECS, Northwestern U.

IEEE INFOCOM 2010

Network Positioning Systems (NPS)

- *How to pick among alternative hosts?*
- A common problem in distributed systems
  - Locate closest game server
  - Ensure inter-node latency bounds for a cluster
- *Scalable way of determining relative location of hosts*
- Different Approaches
  - Landmark-based network coordinates (e.g. GNP)
  - Landmark-free network coordinates (e.g. Vivaldi)
  - Direct measurement (e.g. Meridian)
  - Measurement reuse (CRP)

Large-scale P2P – A good client

- Benefits for large-scale P2P systems
  - High performance
  - Reduced cross-ISP traffic
  - Improved robustness to failures
- How do they fare when deployed?
  - At the scale of large P2P systems?
  - At the edge of the network?
- An open question
  - NPS evaluated in simulation and research testbeds
  - Lack of representative traces of P2P environments
  - No platform suited for experimentation at appropriate level

The world looks different from the edge...

- Median latency between P2P nodes
  - **2x what reported by Ledlie et al.**

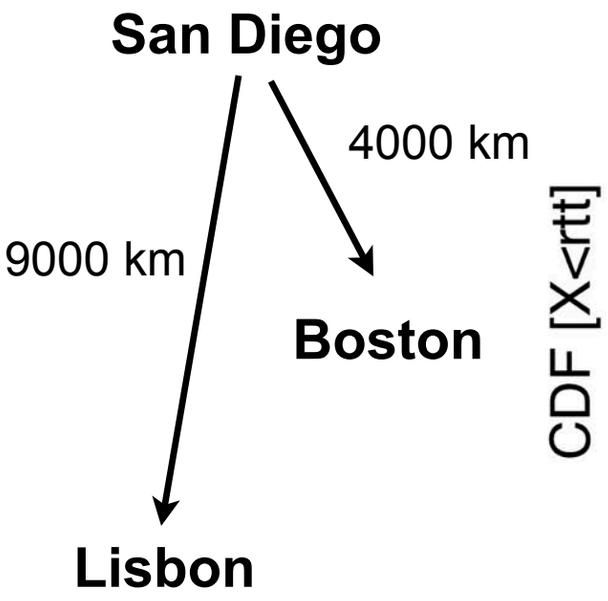

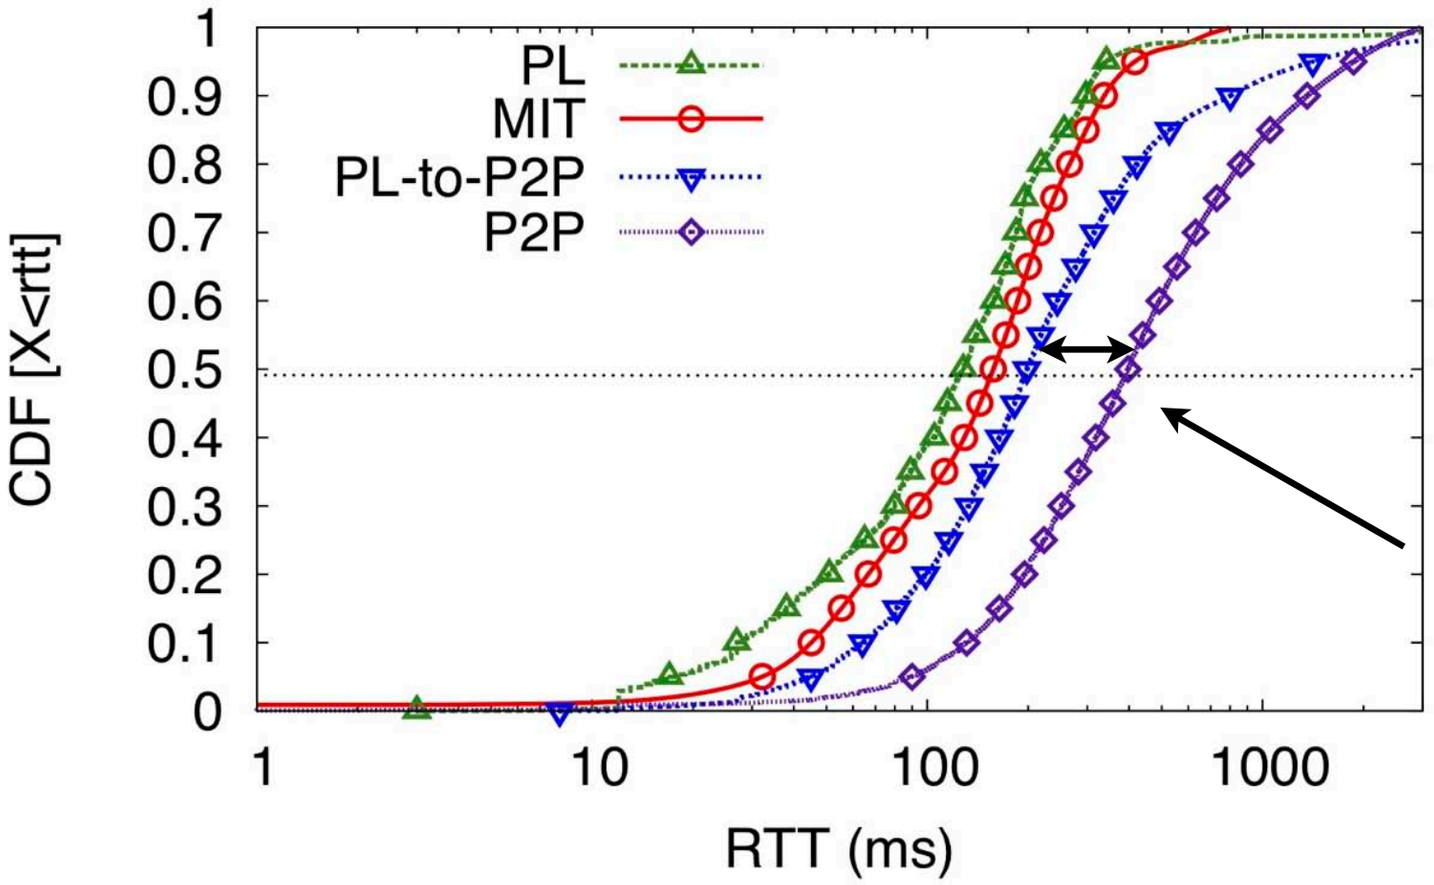

Going to the edge

- To know if NPS work for P2P you need to go to the edge
- Compare representative network positioning systems
  - GNP, Vivaldi, Meridian and CRP
- Use dataset gathered from BitTorrent users
  - Directly sample Vivaldi and CRP positions
  - Collect latency and traceroute between connected peers
    - 15 days, ~20k peers, ~8k routable prefixes - **2 billion latency samples**
- Latency Matrix
  - Matrix of source and destination **routable BGP prefixes**
  - Use minimum observed RTT for each matrix element
  - Obtain a 95%-full matrix: 479x479 matrix

Measuring the accuracy of net pos systems

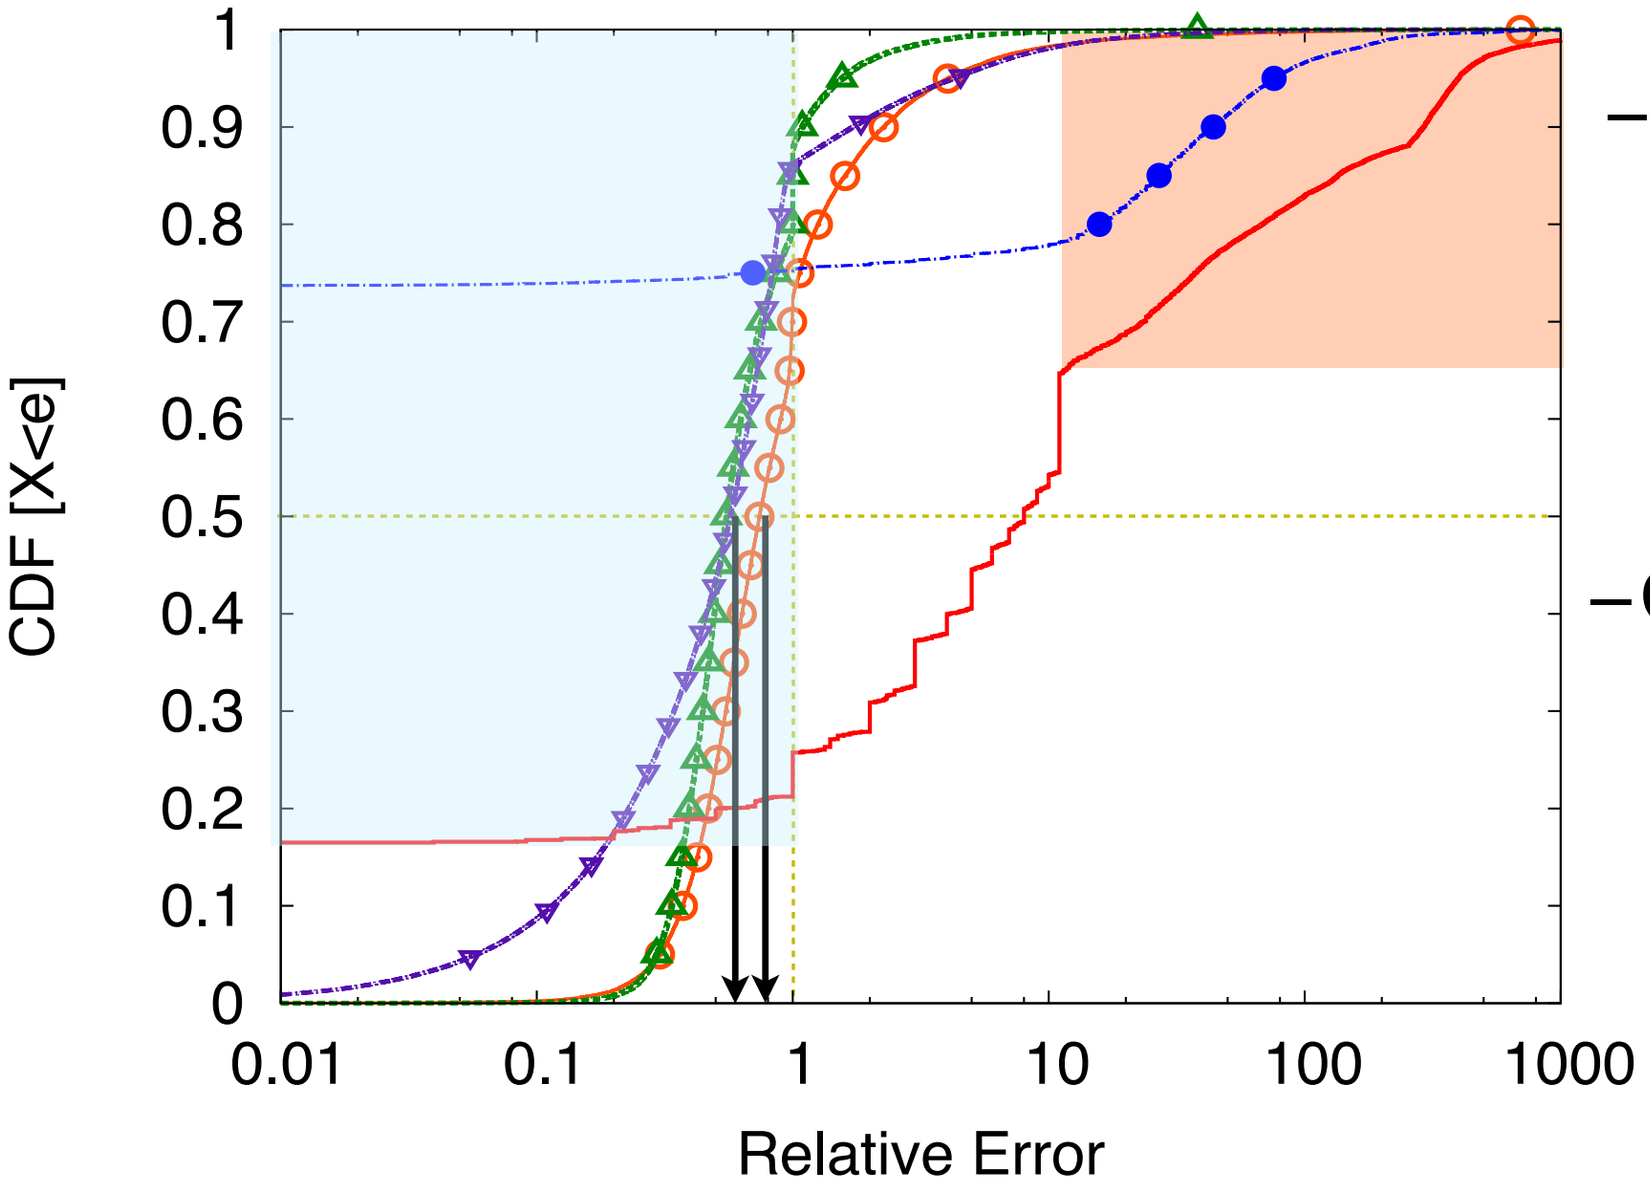

- Vivaldi
  - 74% (v1)
  - 55% (v2)
  - >> 26%
- CRP, Meridian
  - Still high degree of error

Vivaldi V1
Vivaldi V2
GNP
CRP
Meridian

* **GNP:** 3 runs, 15 landmarks , 464 targets, 8-dimensional coordinate space

** **Meridian:** 379 random Meridian nodes, 100 target nodes, 16 nodes per ring and 9 rings per node

Average RALP (Rel. App-Latency Penalty)

- RALP: (selected - optimal)/optimal

  * Much greater than 0.4 RALP reported by Ledlie et al.

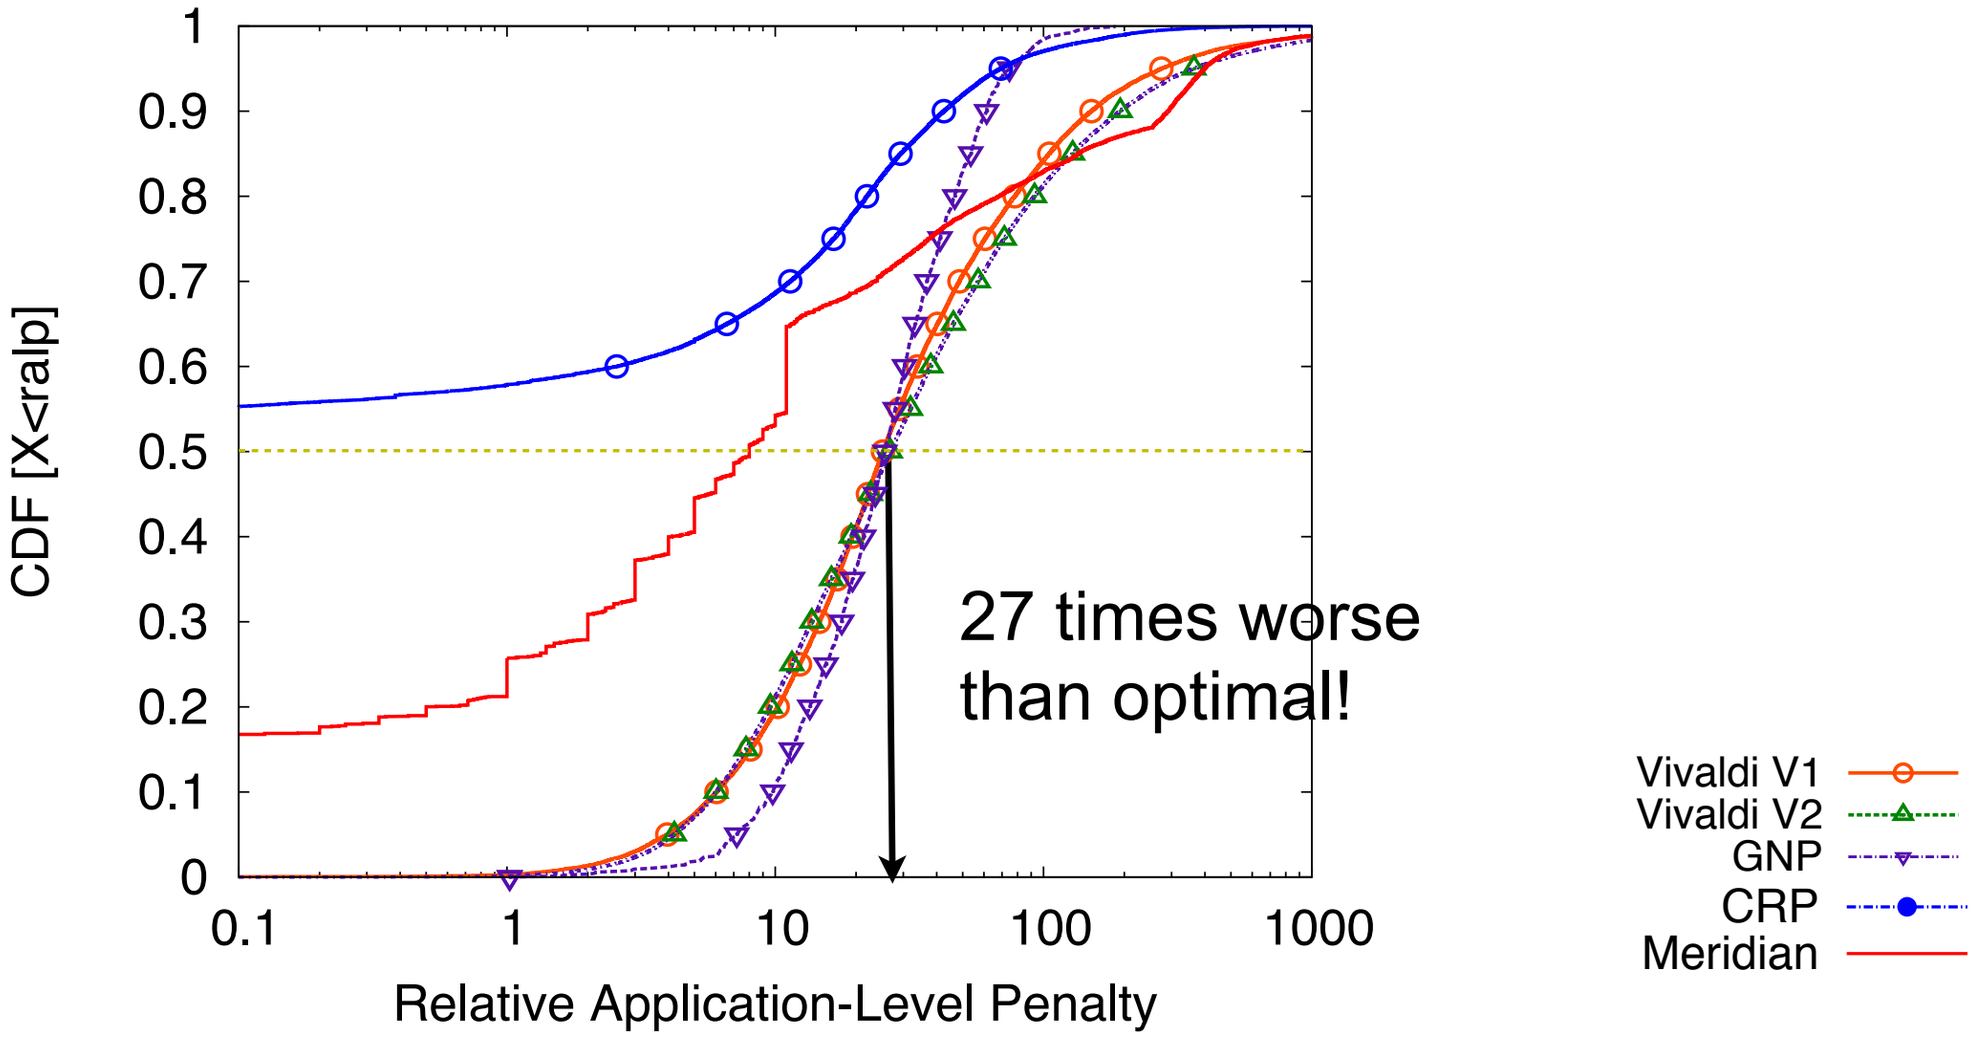

Sources of error

- High Variance
- **First- and Last- Mile Issues**
- **Triangle Inequality Violations**

First- and last-mile issues

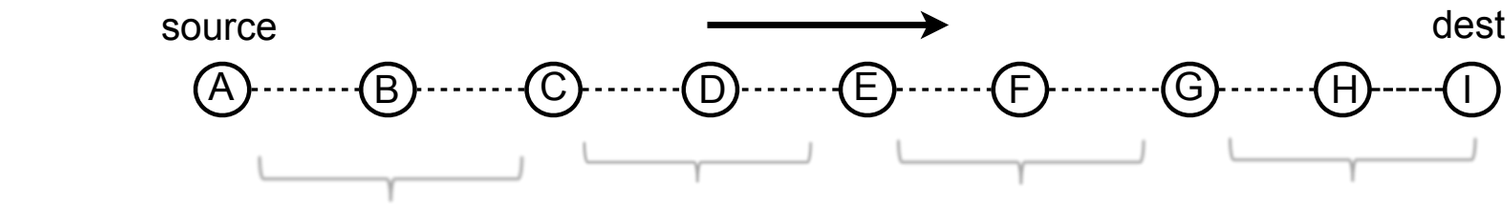

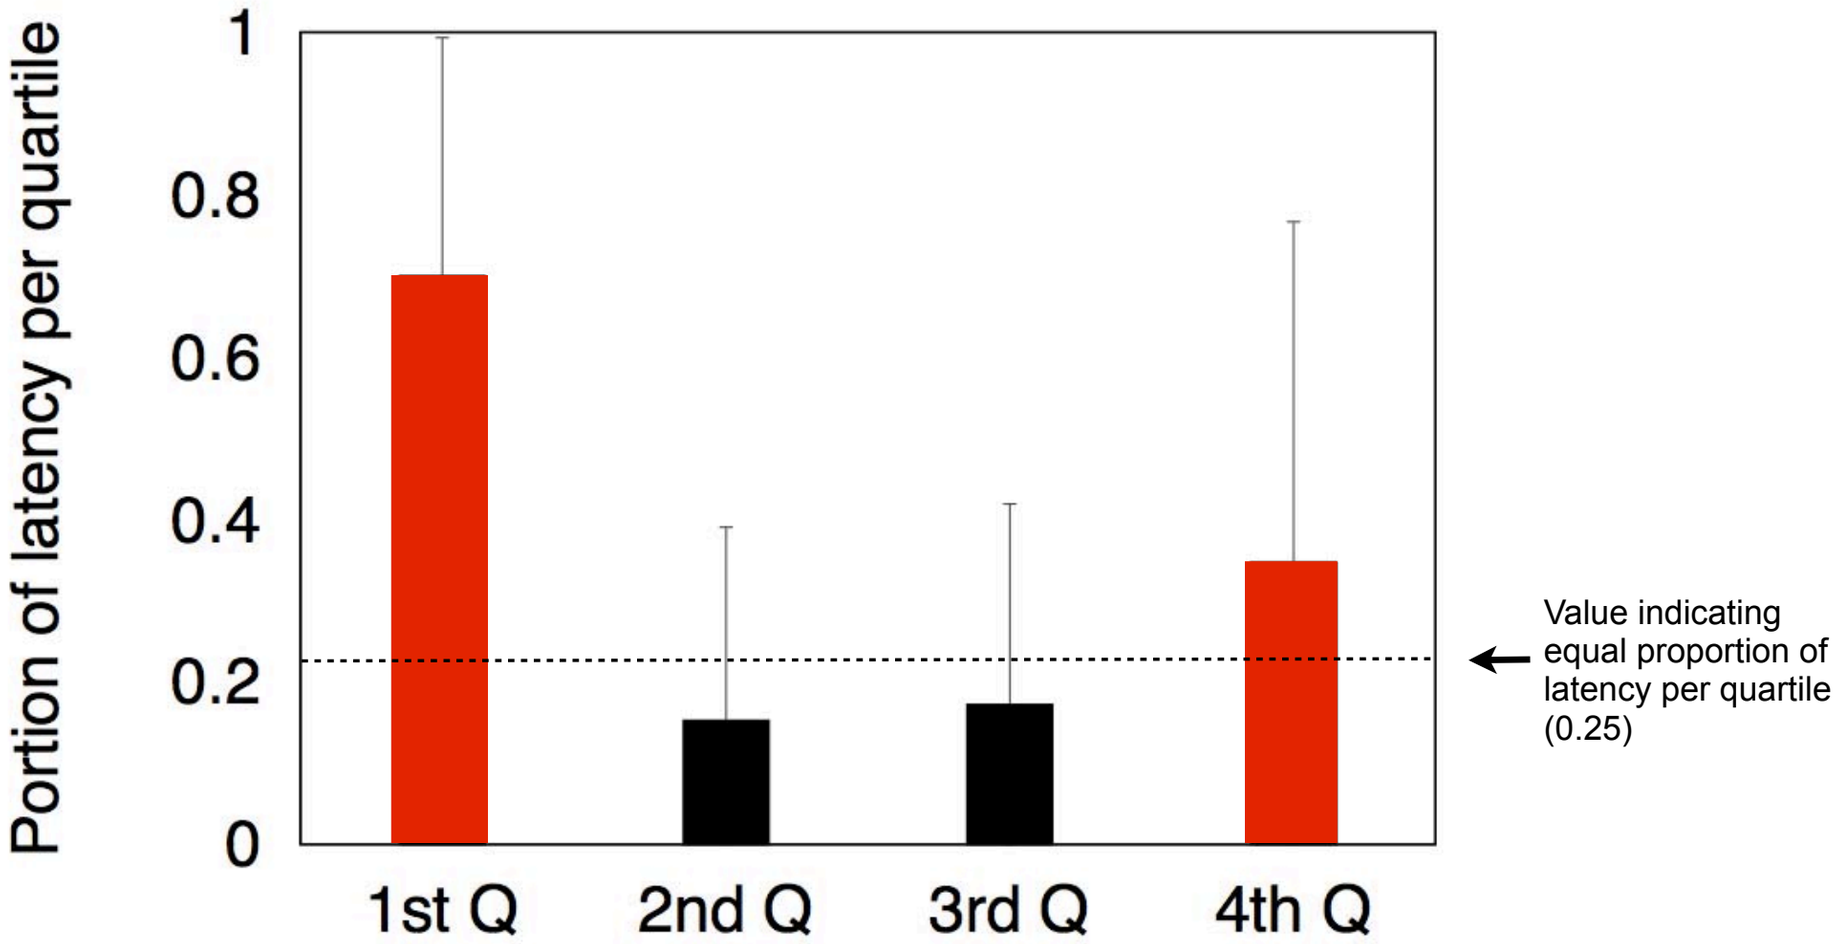

Triangle Inequality Violation Severity

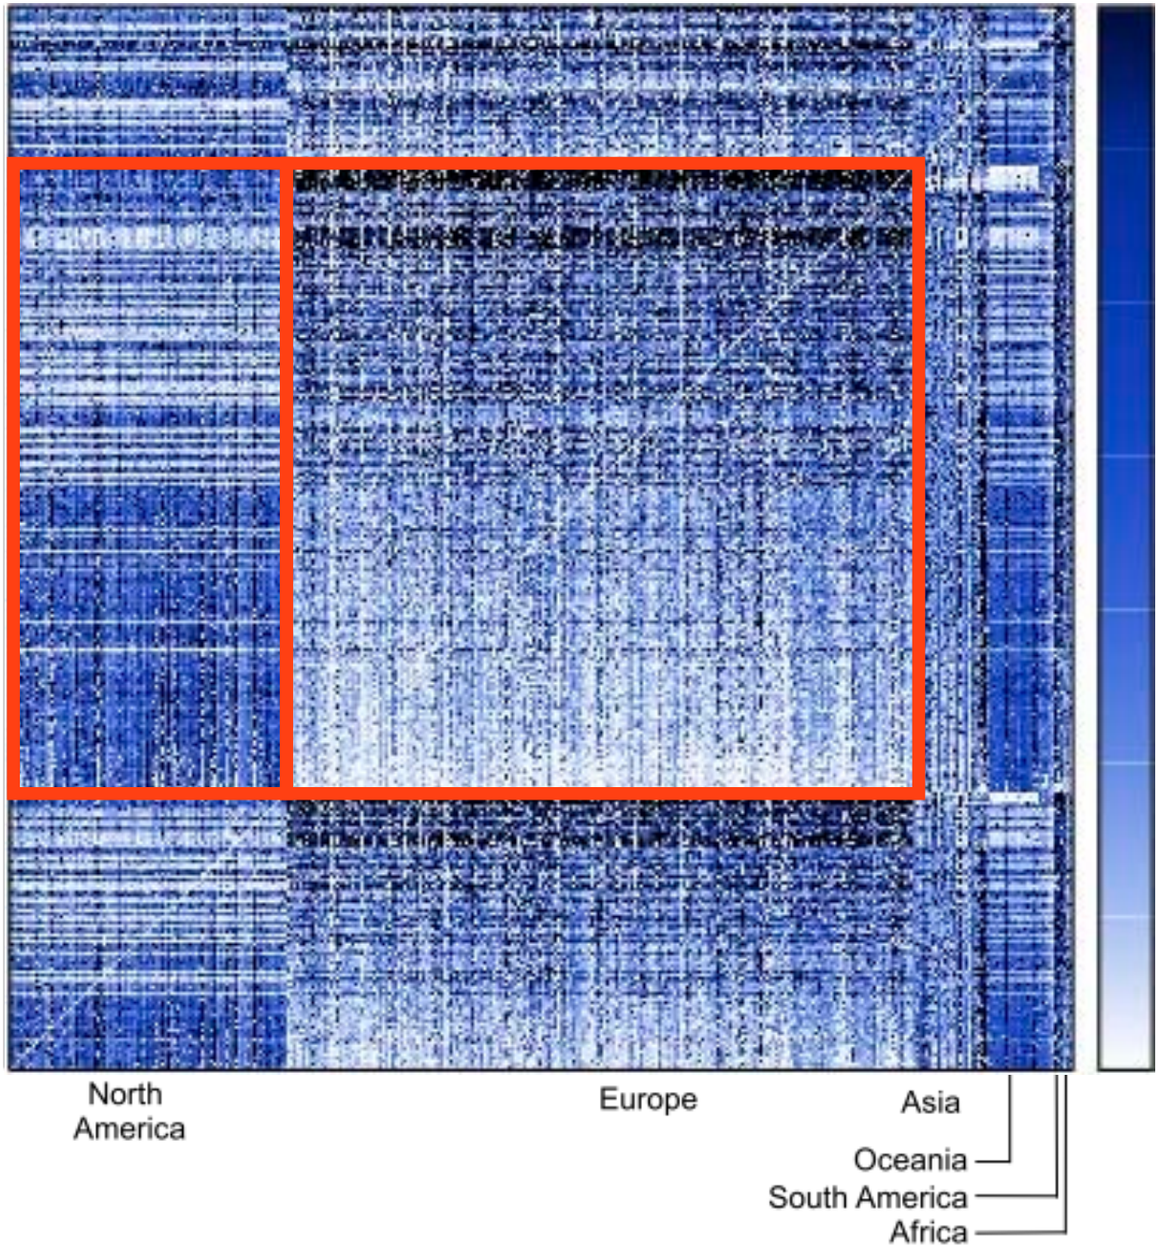

- Lighter colored → less severe TIVs
- 13% of triangles have TIVs (99.5% of source/destination pairs)
- No clear pattern

C

AB > ACB

A B

Conclusions

- NPS performance at the edge is important but generally unknown
  - Requires evaluating them where they are used
- Most NPS deployed to edge perform much worse than predicted
  - Not particularly new problem: Last-mile issues, TIVs, etc.
  - **Severity** of issues much worse

Time to make the subject hot again?

Thanks...

Questions?