Internet-wide systems need Internet-wide measurement platforms

Fabián E. Bustamante

EECS, Northwestern U.

Building Internet-wide systems

- Research on networks and distributed systems
  - Public network testbed (e.g. Emulab, ModelNet, PlanetLab, …)
  - Public views of network topologies (e.g. RouteViews)
- Assume (hope?) testbed results extend to the wider network

Using public data

- *Our study is comprehensive, we cover all prefixes made available by RouteViews*
- Actually, a very incomplete picture [CONEXT09]
  - Use traceroutes between P2P users
  - Apply magic filter…
  - Added 24,000 new AS links, inferred relationships
    - ~41% more peering links
    - ~13% more customer-provider links

Mapping flows with public data

- Understanding locality and cost of P2P traffic
  - 3 month data from BitTorrent – 2.8M AS pairs from all connections
  - Look for paths that contain ASes in the public view
    - Can't map >50% for 68% of hosts!
    - Need traceroutes from the edge to complete the picture

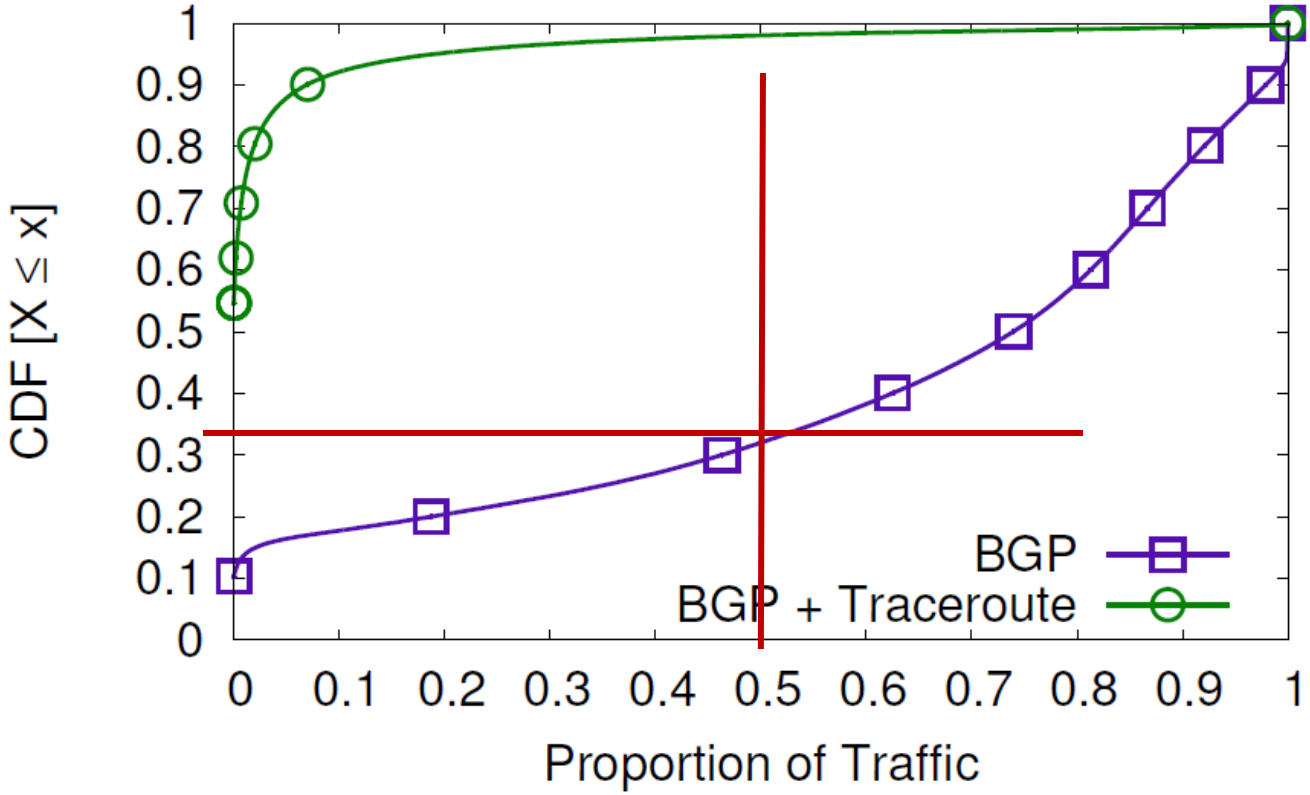

Using a public platform

- *Our system works on PlanetLab, so*
  - Results are representative of a large deployment
  - Translate well up to some constant deviation
- Reality
  - Ledlie et al., Network Coordinates in the Wild
  - From PlanetLab to hosts at the network edge
    - Vivaldi out of tune

Latencies from where you measure

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

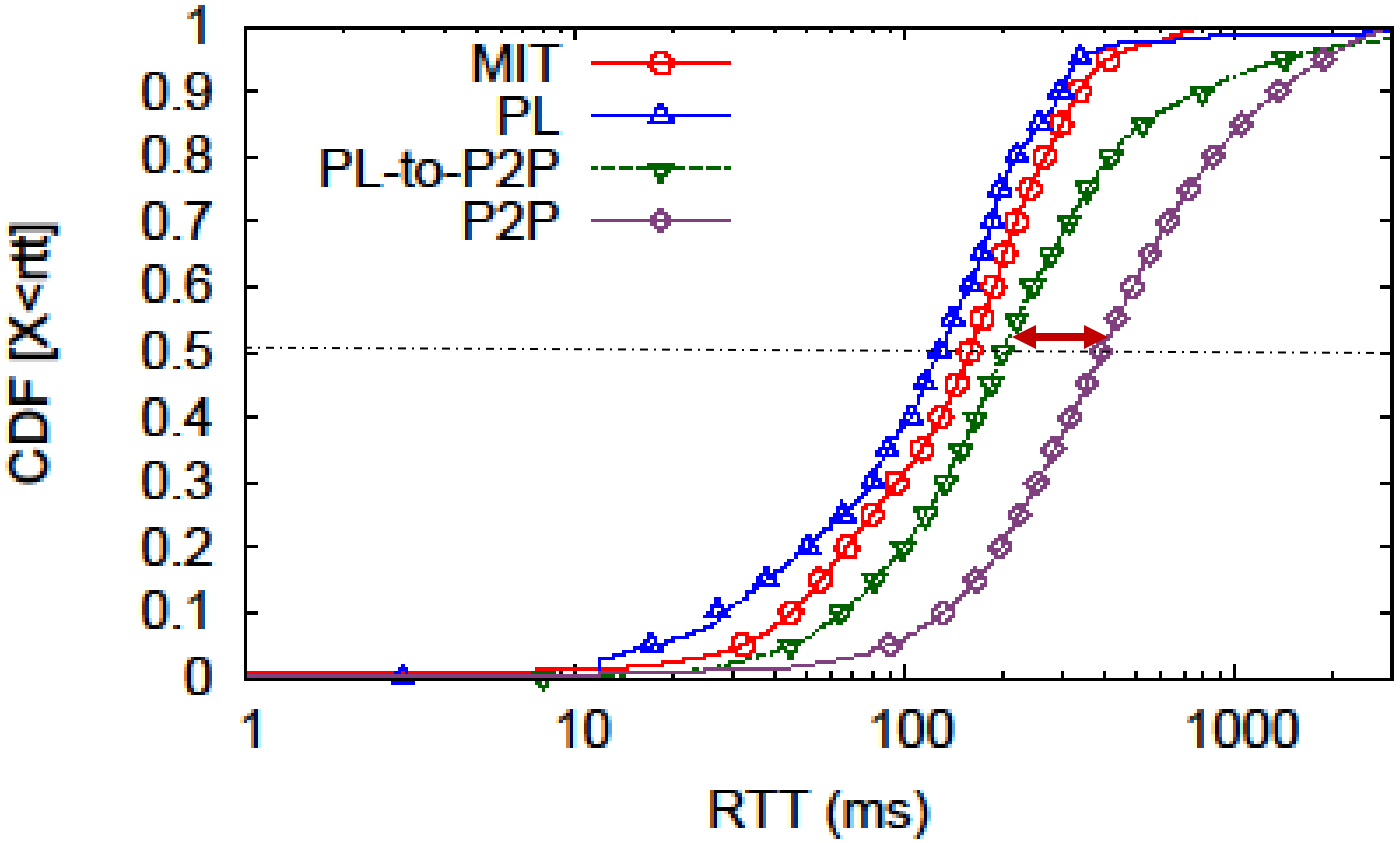

Impact on net positioning

- Errors in predicted latencies (median errors)

GNP – 59.8ms

Vivaldi –
150ms (v1),
165ms (v2), but
smaller relative error

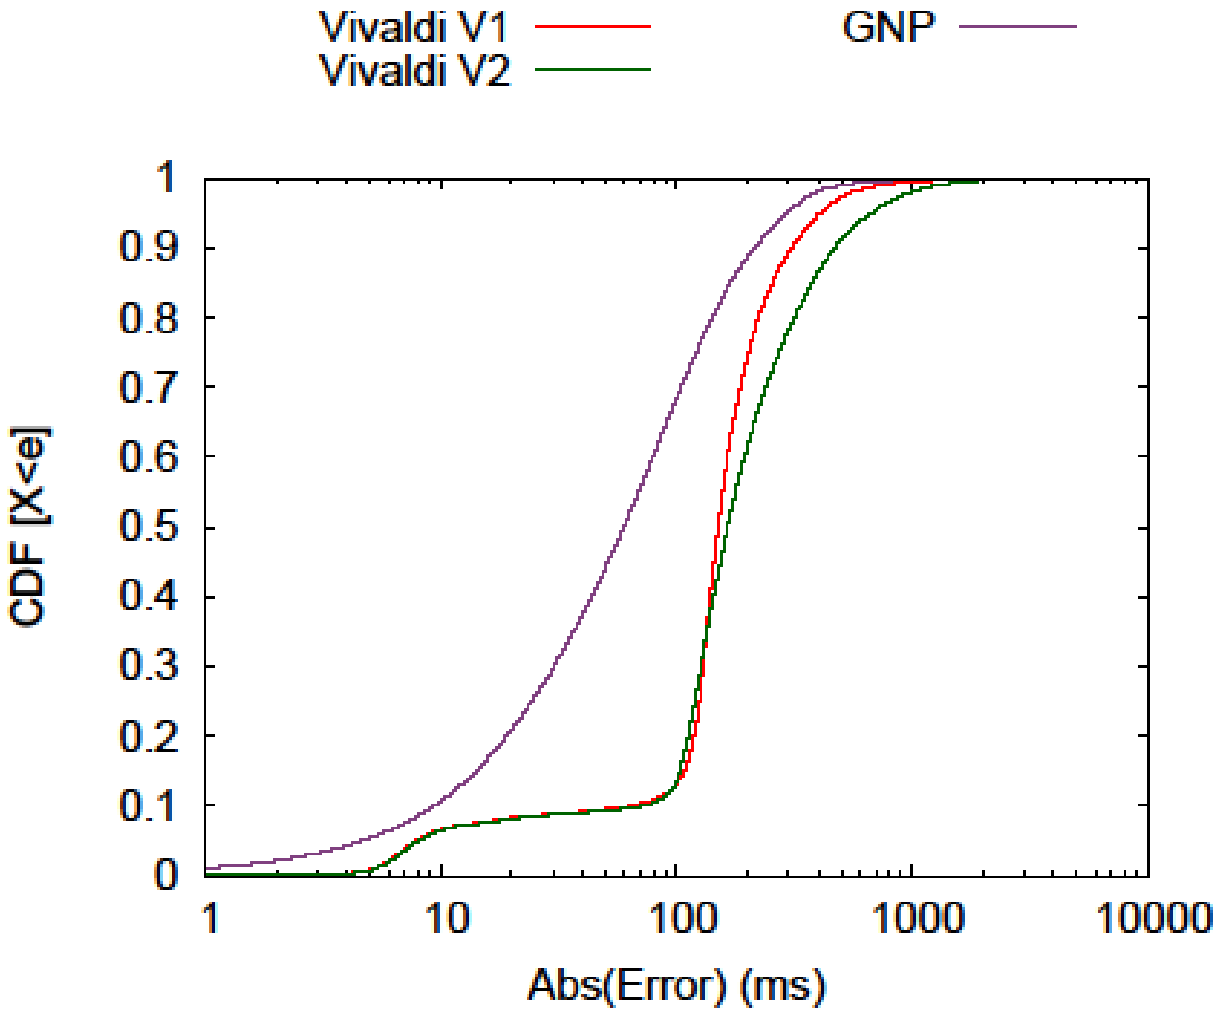

Using a public platform and public data

- Perfect combination for misleading conclusions
  - Incomplete view
  - Unrepresentative results
- Example
  - On the nature of P2P traffic
    - Growing or shrinking?
    - Local or nonlocal?
    - Costly or mostly harmless?
  - Answer depends on perspective
  - The **right** answer requires both a (mostly) complete view and representative measurements

BitTorrent volume as a proxy for popularity

- Downloading going up still (23% increase)
- Uploading has remained steady

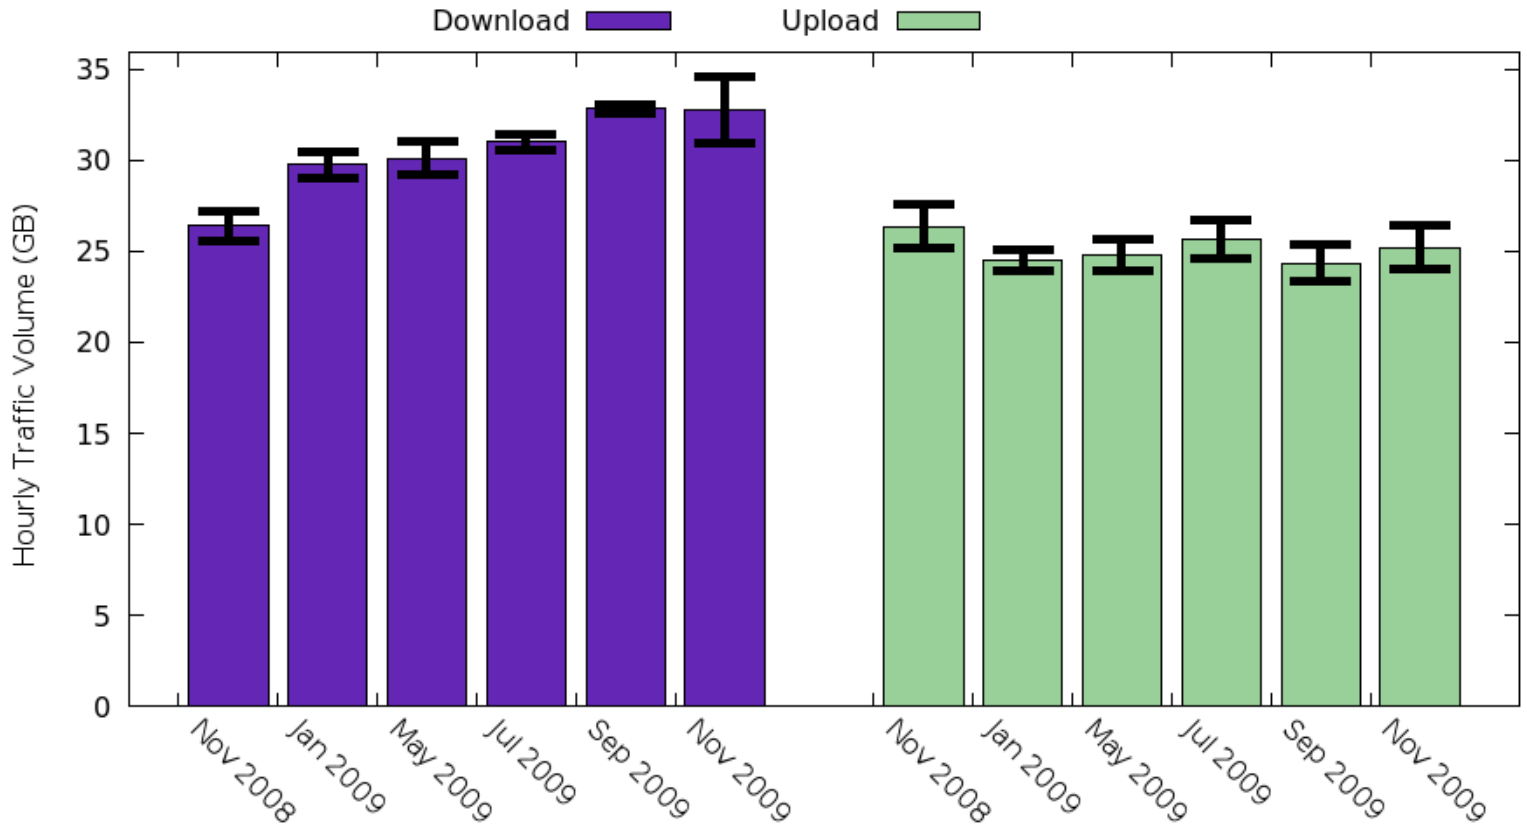

BitTorrent's online peers over time

- How's BitTorrent doing in terms of users?
  - 20% drop on number of connections per vantage point

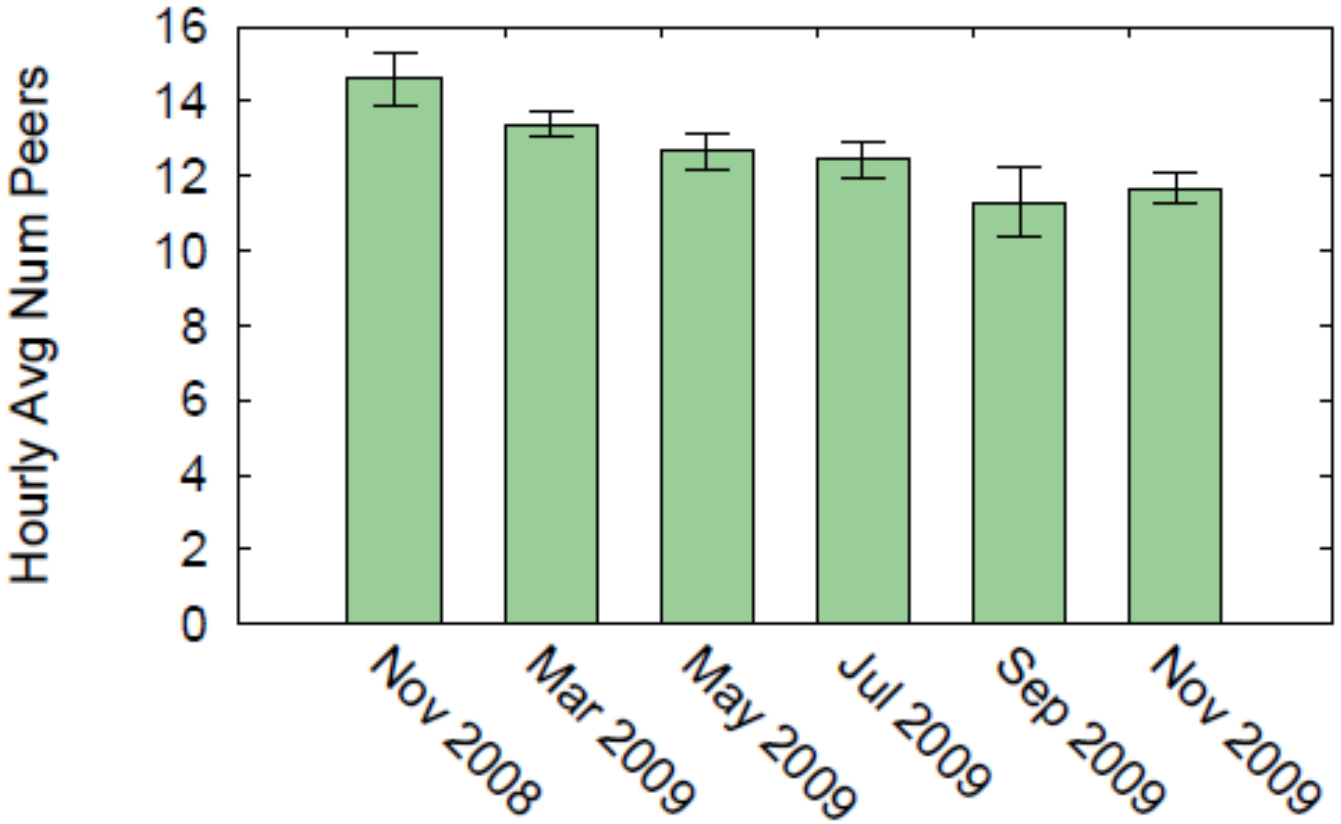

Online peers over time per continent

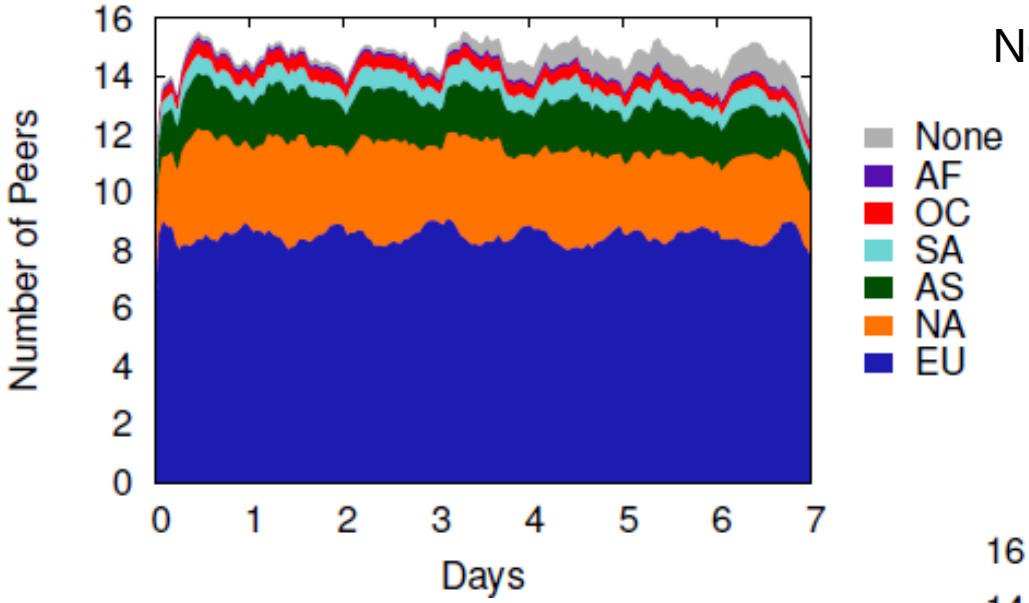

November 2008

Drop in Europe and US with more clear daily patterns in Europe (and shorter session times)

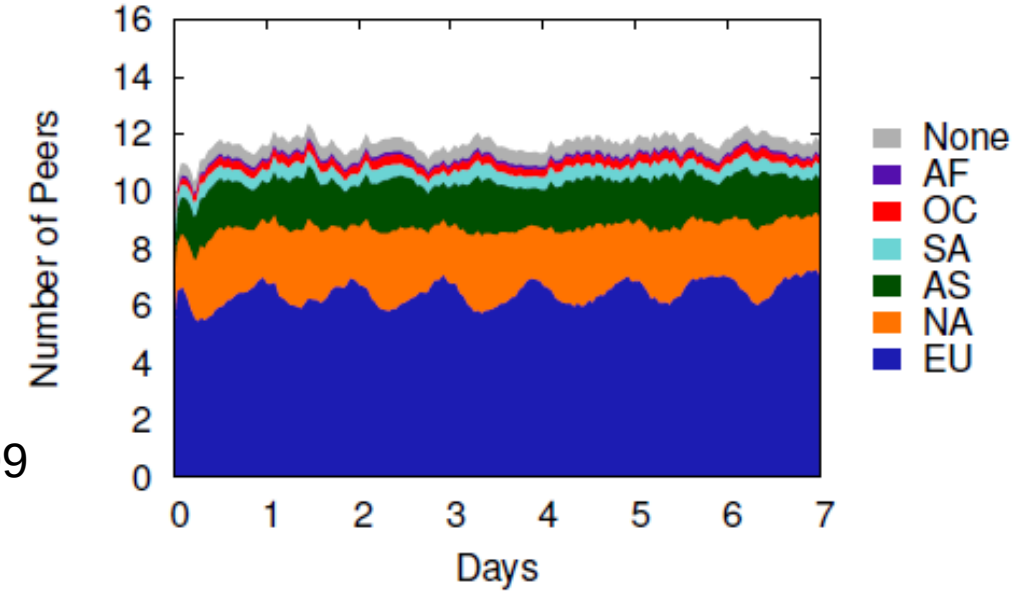

November 2009

Locality – North America and Europe

- Everybody likes going to Europe (and staying there)

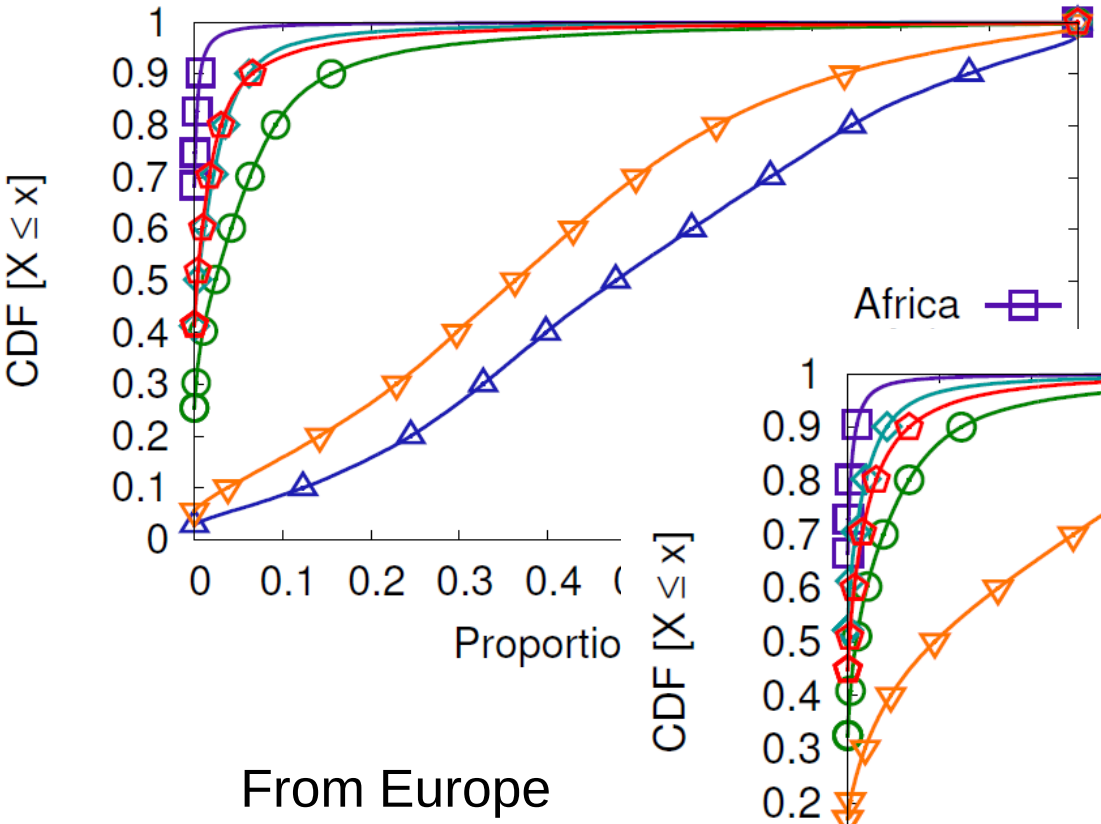

From Europe

From North America

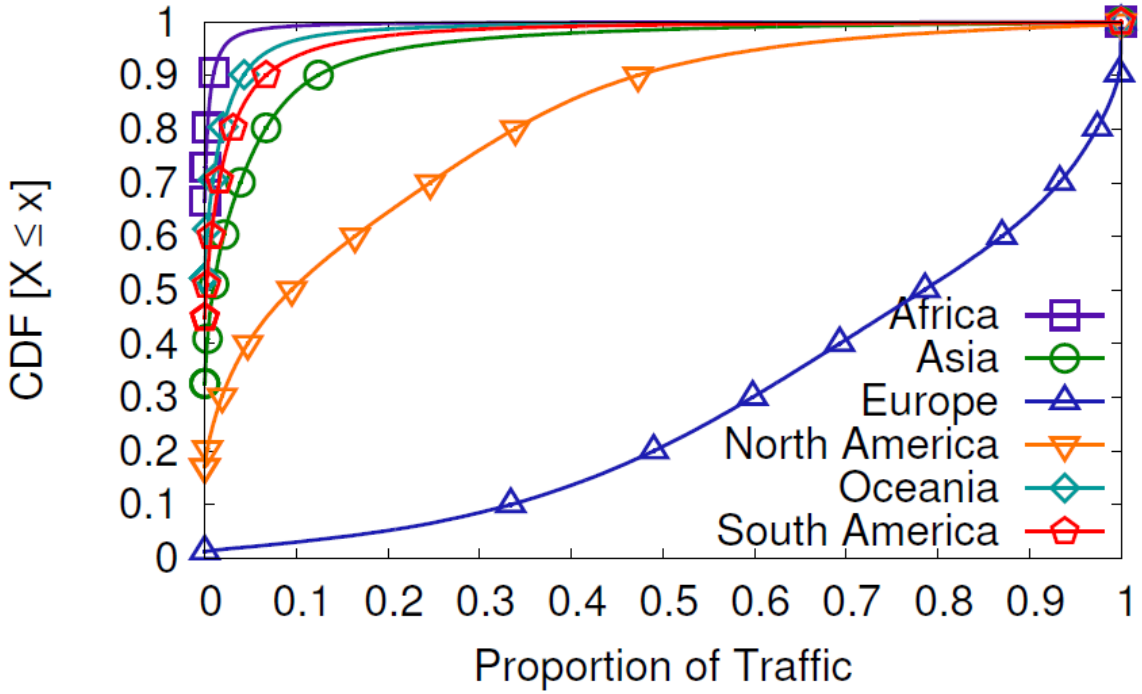

Who pays and how much does it cost?

- Customer:provider traffic per tier (simplistic cost model)
  - Value >1 – AS receives more customer traffic than it sends to providers (revenue?)
  - 39%of tier-2 Ases and 12% of tier-3 may be profiting

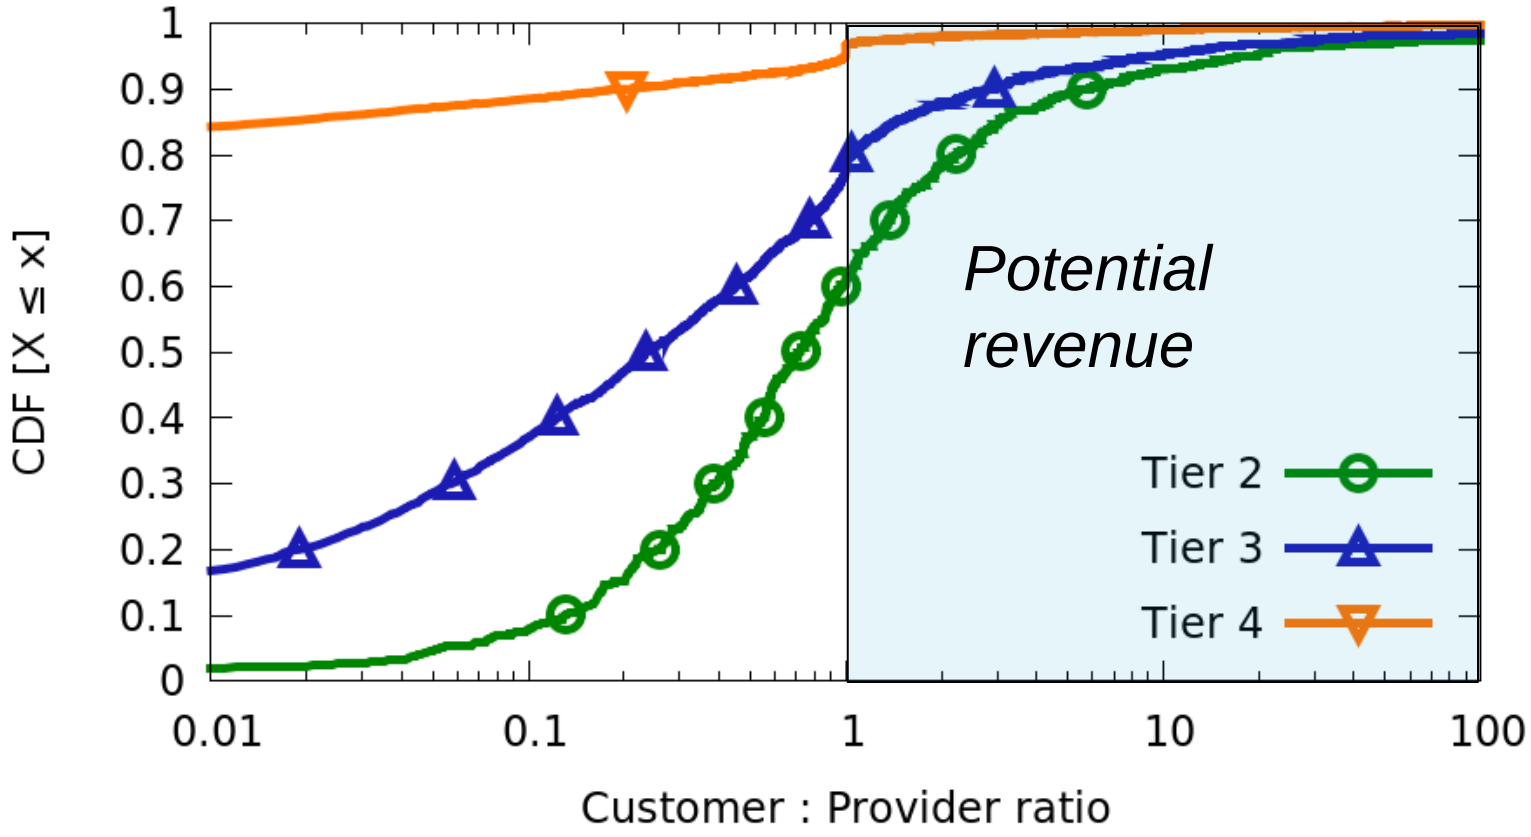

Who pays and how much does it cost?

- Revenue/cost from P2P traffic (include peering and sibling traffic in total traffic)
  - Customer – provider bytes / all bytes flowing through

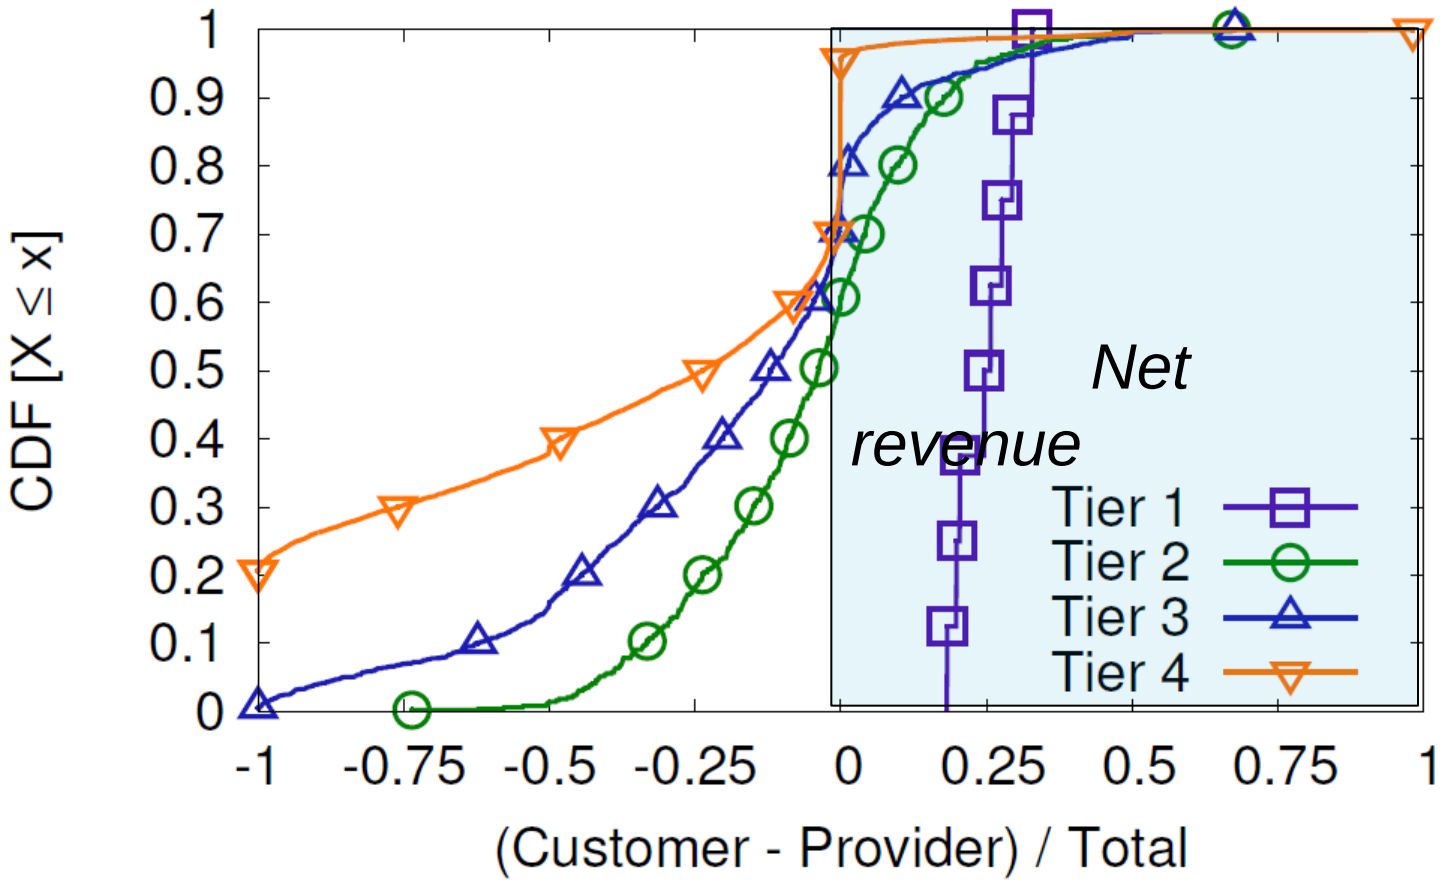

Conclusions

- Today's public network testbeds
  - Emulation
    - Control & repeatability
    - If you are ok with artificial network conditions
  - Overlay testbed
    - More realistic network conditions
    - At the price of control and repeatability
- Going for the edge
  - Today – opportunistic measurement
  - Room for a third stage?
  - Natural experimentation for Internet systems