Some Findings on the Network Performance of Broadband Hosts

Karthik Lakshminarayanan, UC Berkeley
Venkat Padmanabhan, Microsoft Research

Primary focus of our work

- Increasing number of broadband hosts in “real world”
- Understand broadband connectivity:
  - Raw characteristics
  - Applicability of traditional measurement techniques
- Impact on P2P systems:
  - Find ‘good’ peers in a P2P system
  - Impact on applications, in particular overlay multicast

Related work

- Direct measurements: (such as NPD, Nimi)
  - Restricted to “well-connected” hosts
- Indirect inferencing: (UW study, CMU study)
  - Use peers of file-sharing system as vantage points
    - Most vantage points were “well-connected” machines
  - Inability to measure directly between broadband hosts
- Our work:
  - Run measurement agents on broadband hosts
  - Perform direct measurements
  - Ability to study properties at a micro-scale

Constraints

- Difficulty in recruiting volunteers:
  - Privacy concerns:
    - Could not measure/use existing traffic
  - Restriction on bandwidth consumption:
    - Imposed a limit of 10 kbps (averaged over few minutes)
  - Cannot obtain login access to machines
    - Run as a Windows service (self-starting daemon)
- NATs:
  - Used techniques similar to IETF STUN proposal for UDP packets traversing NATs

Design of PeerMetric

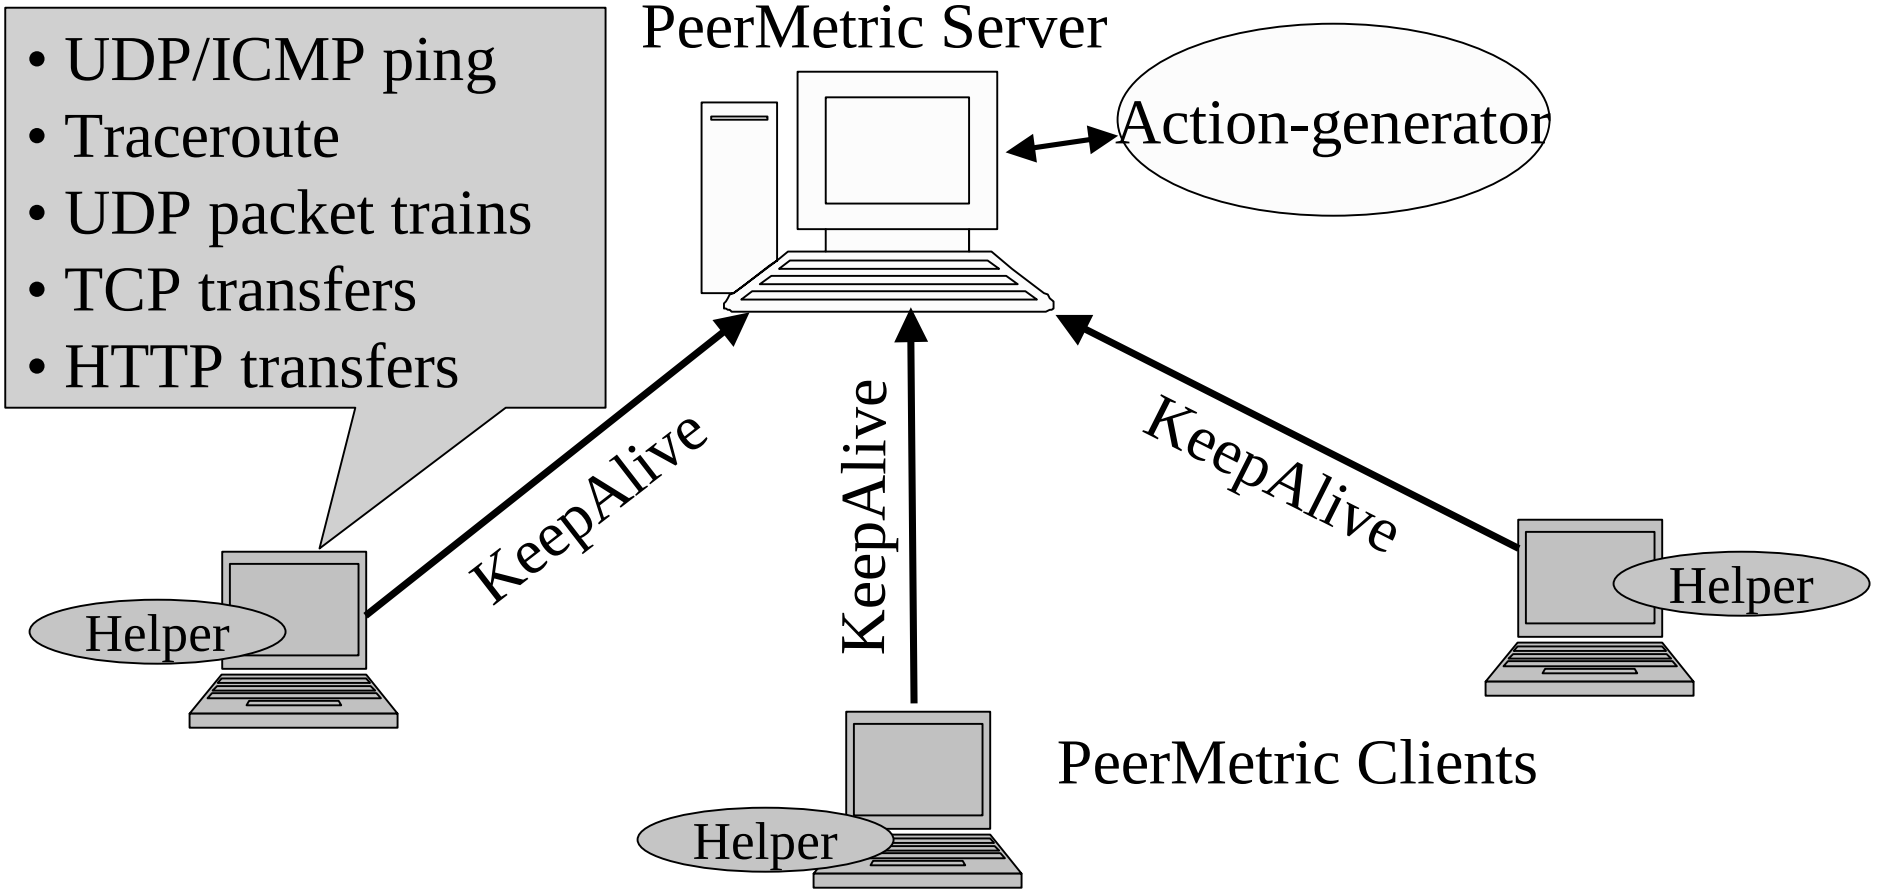

Intelligence about tests to perform is placed in the *Action-generator* which runs at the server

Each PeerMetric client performs the basic P2P tests

Initial deployment of PeerMetric

Connection Type

Cable modem: 13

DSL: 12

Main ISPs

AT&T Broadband: 9

Verizon DSL: 8

Initial PeerMetric deployment had 25 hosts

Summary of results

- Confirmation of known results:
  - Asymmetry in bandwidth
    - Median Upstream =212 kbps, Downstream = 900kbps
  - Latency between hosts is high
    - Median of 40ms between hosts in the same city compared to 3-4 ms between well-connected hosts
- Interesting results:
  1. Broadband link “management” affects measurements
  2. Delay-vector technique picks proximate peers well
  3. P2P latency is a poor predictor of P2P throughput
  4. Locality-based heuristics for tree construction perform poorly

#1: Impact of broadband link management

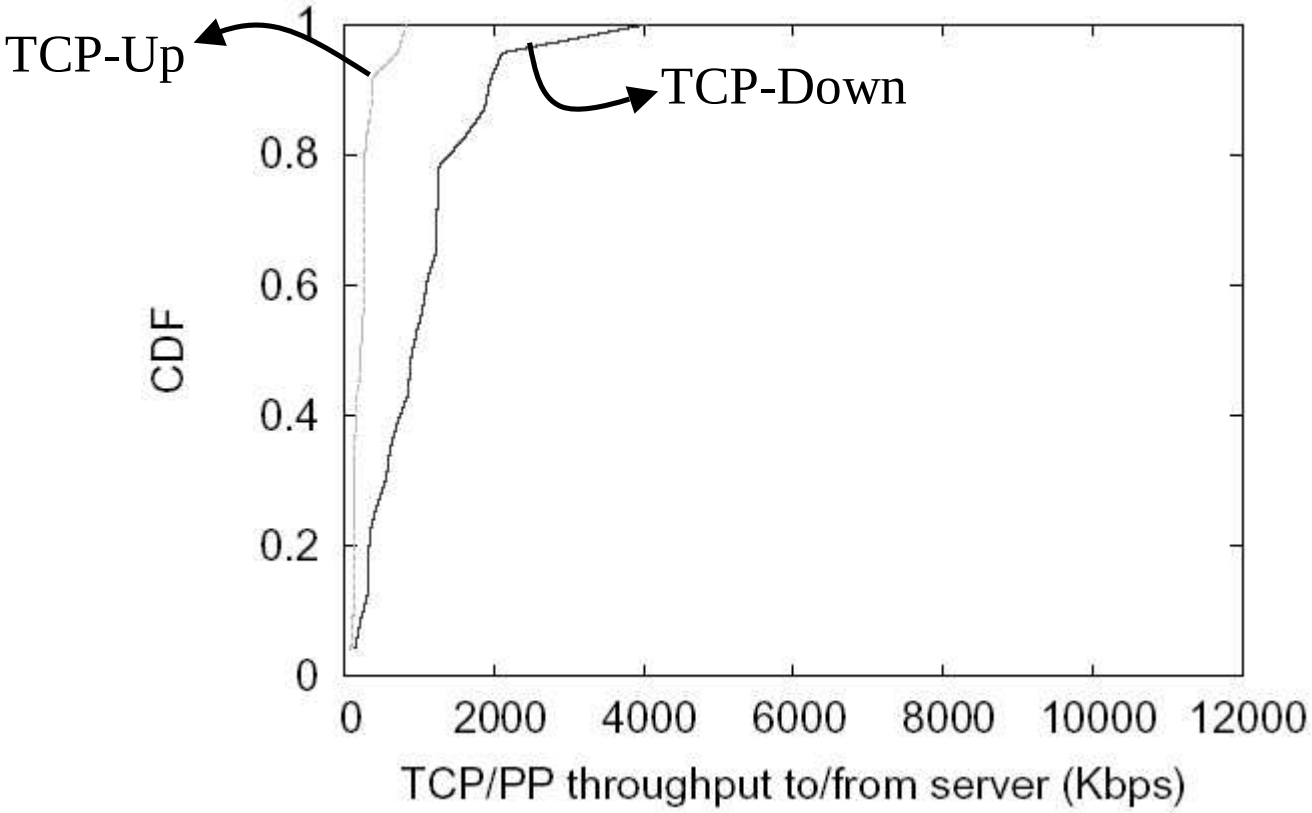

Prevalence of asymmetry is not surprising

#1: Impact of broadband link management

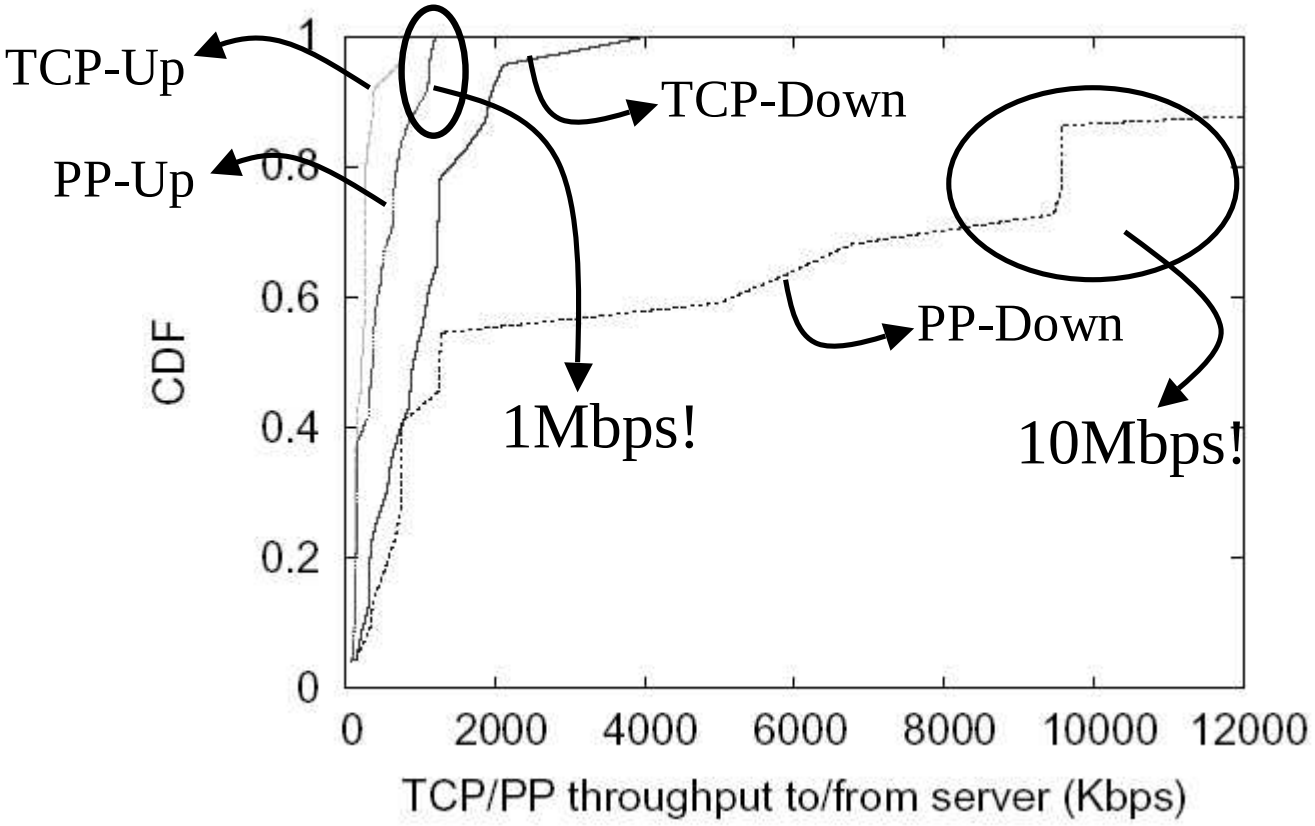

Packet-pair throughput >> TCP throughput

- Observed for cable modem hosts only

#1: Impact of broadband link management

Date & Time	Type	Bandwidth observed (kbps) Pair-1	Pair-2	Pair-3	Pair-4	Pair-5
9-18:15:9:26	PKTPAIR	748.86	744.33	465.19	259.63	242.69
9-18:16:4:16	PKTPAIR	749.01	744.28	531.71	253.68	237.58
9-18:16:47:59	PKTPAIR	751.63	743.00	436.02	245.44	247.86
9-18:15:9:42	TCP	242.62				
9-18:16:4:32	TCP	241.72				
9-18:16:48:16	TCP	241.88				

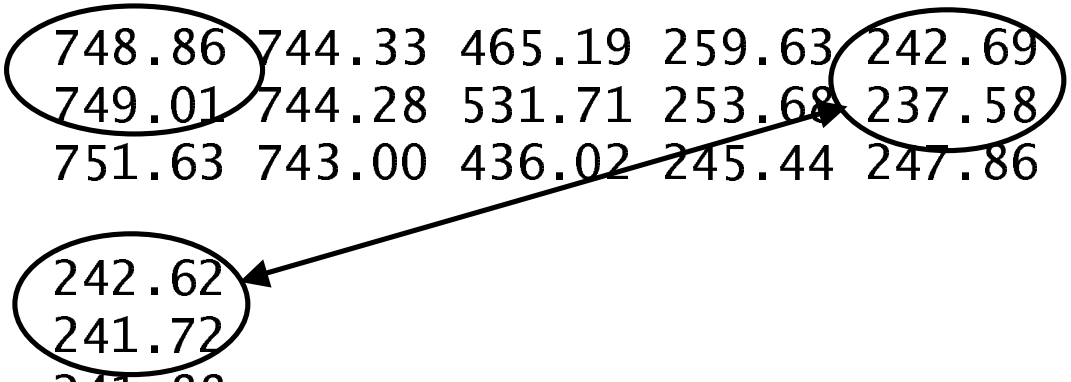

- Observed only for cable-modem hosts
  - Cable-modem routers perform token-bucket rate limiting
- Modification of measurement techniques
  - Drain the token bucket before packet-pair measurements

Measurement techniques have to be revisited for broadband hosts

Peer Selection

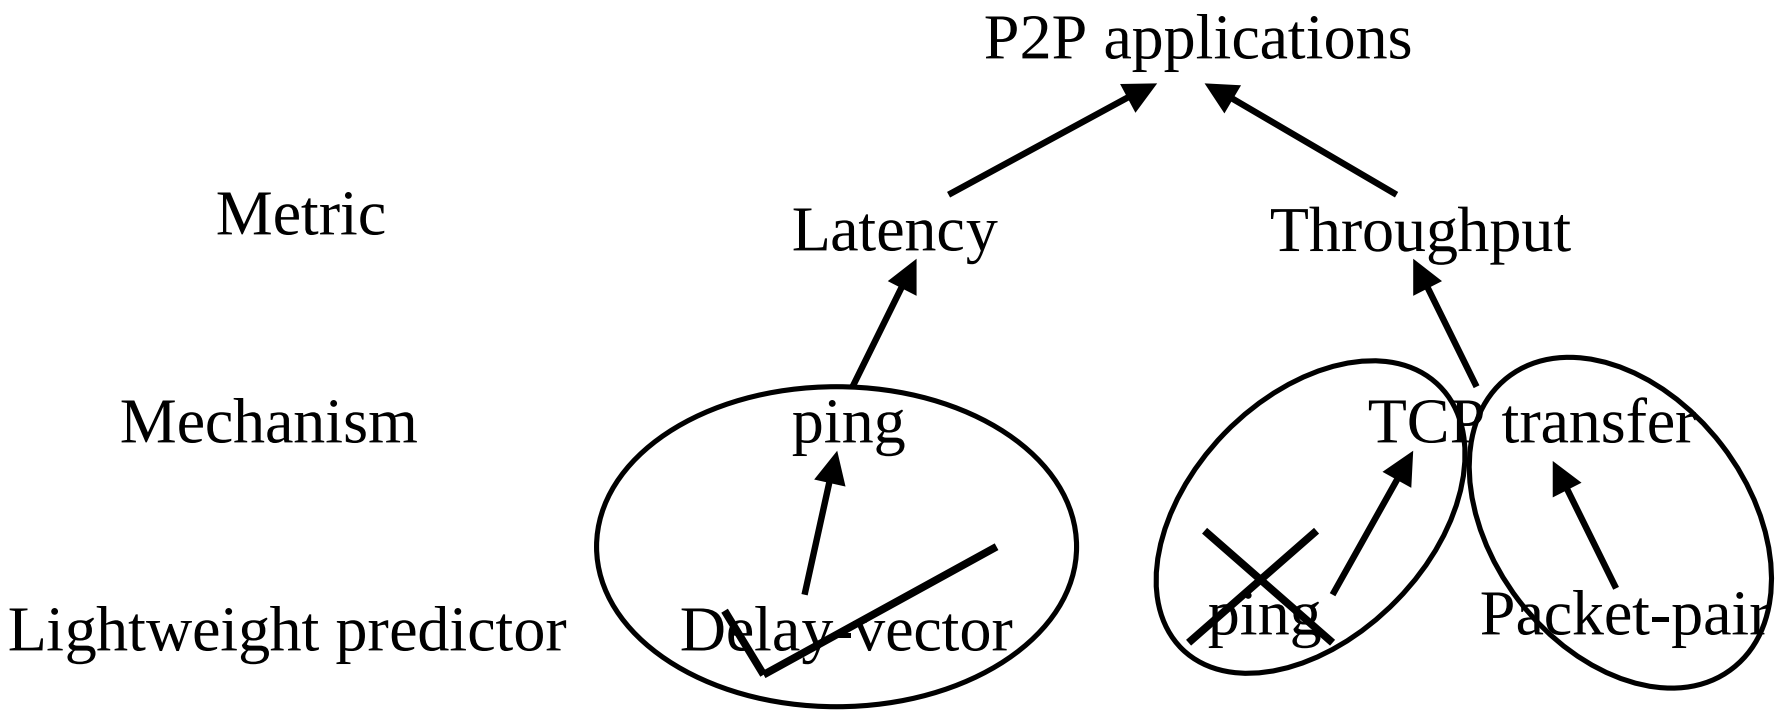

#2: Peer selection: Latency metric

- Delay-vector (coordinates) based approach
  - Motivated by GeoPing, GNP
  - Peers ping a set of landmarks and compute a delay vector

- Rank correlation between Euclidean distance and actual ping time is 0.73

- Ping time to selected peer is within 1.76 of the optimal in 90% of the cases

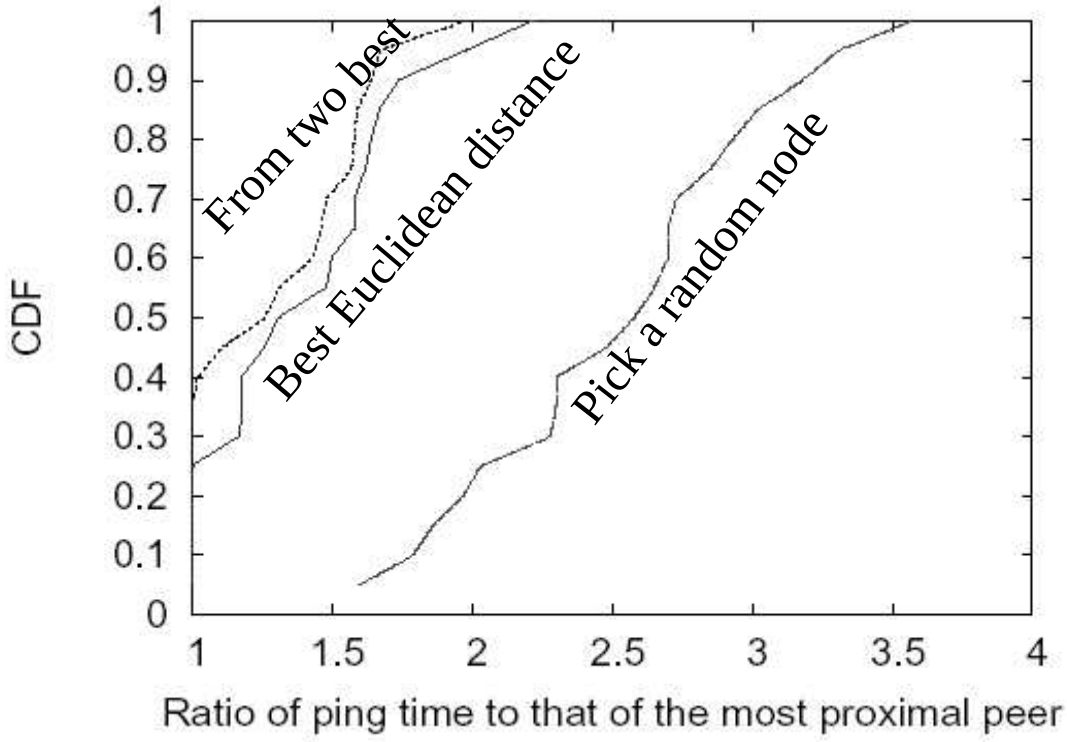

Delay-vector based approach performs well in finding proximate peers

#3: Peer selection: Throughput metric

- Common technique: Ping a set of hosts and pick the best

- P2P latency does not correlate well with P2P throughput
- Linear correlation of –0.14
  Rank correlation of –0.13

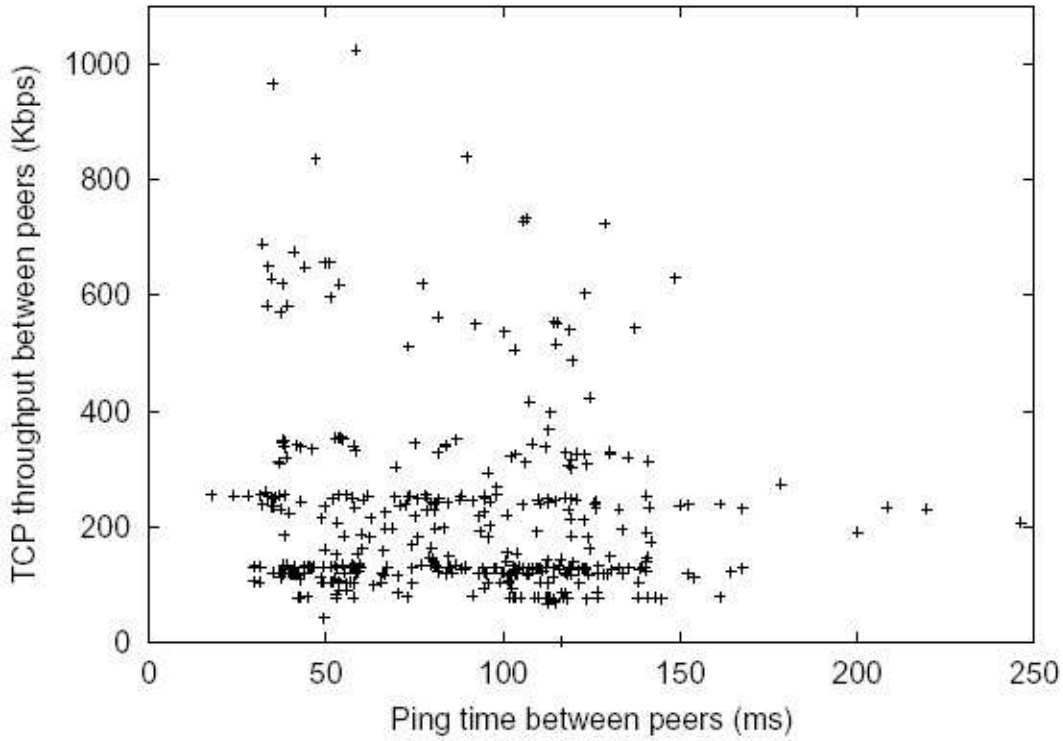

Latency is a poor predictor of TCP throughput (both cable and DSL)

#3: Peer selection: Throughput metric

- Using packet-pair to predict TCP throughput

  (i) Light-weight (ii) Low degree of statistical multiplexing

- Packet-pair is a good predictor for DSL hosts
  - Linear correlation of 0.79
    Rank correlation of 0.92
- Packet pair does not work well for cable hosts
  - Linear correlation of 0.33
    Rank correlation of 0.03

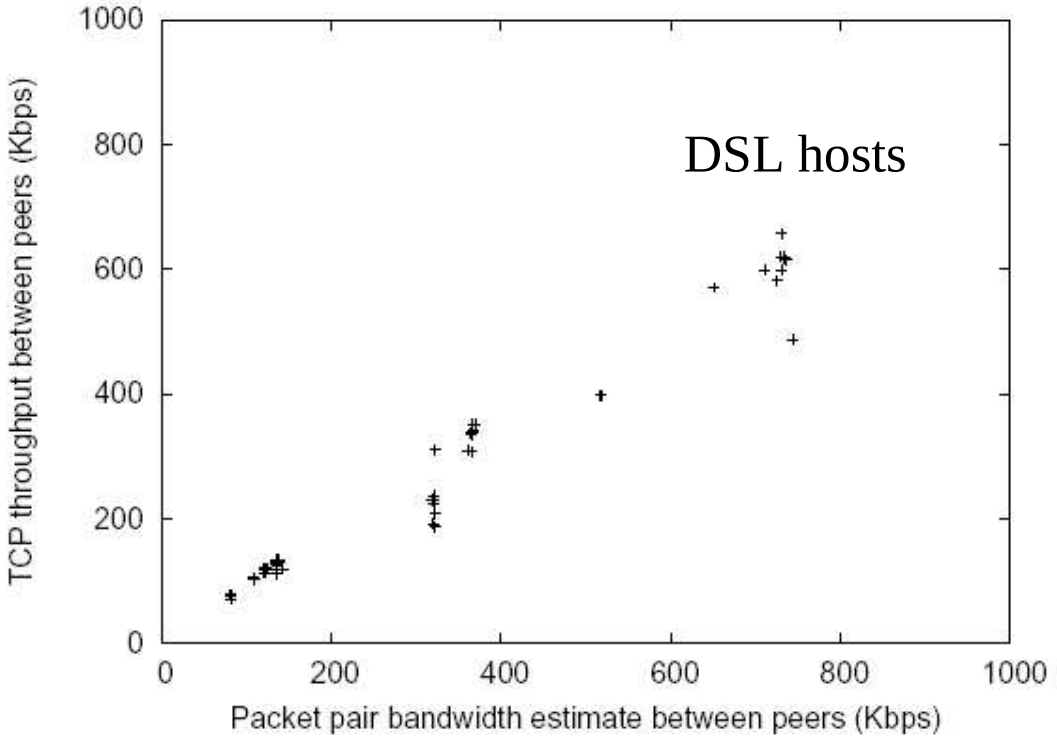

Packet-pair is a good predictor of TCP throughput for DSL

#4: Implications for Overlay Multicast

- Traditional goal: Mimic IP multicast
  - Minimize repeated traversal of physical links

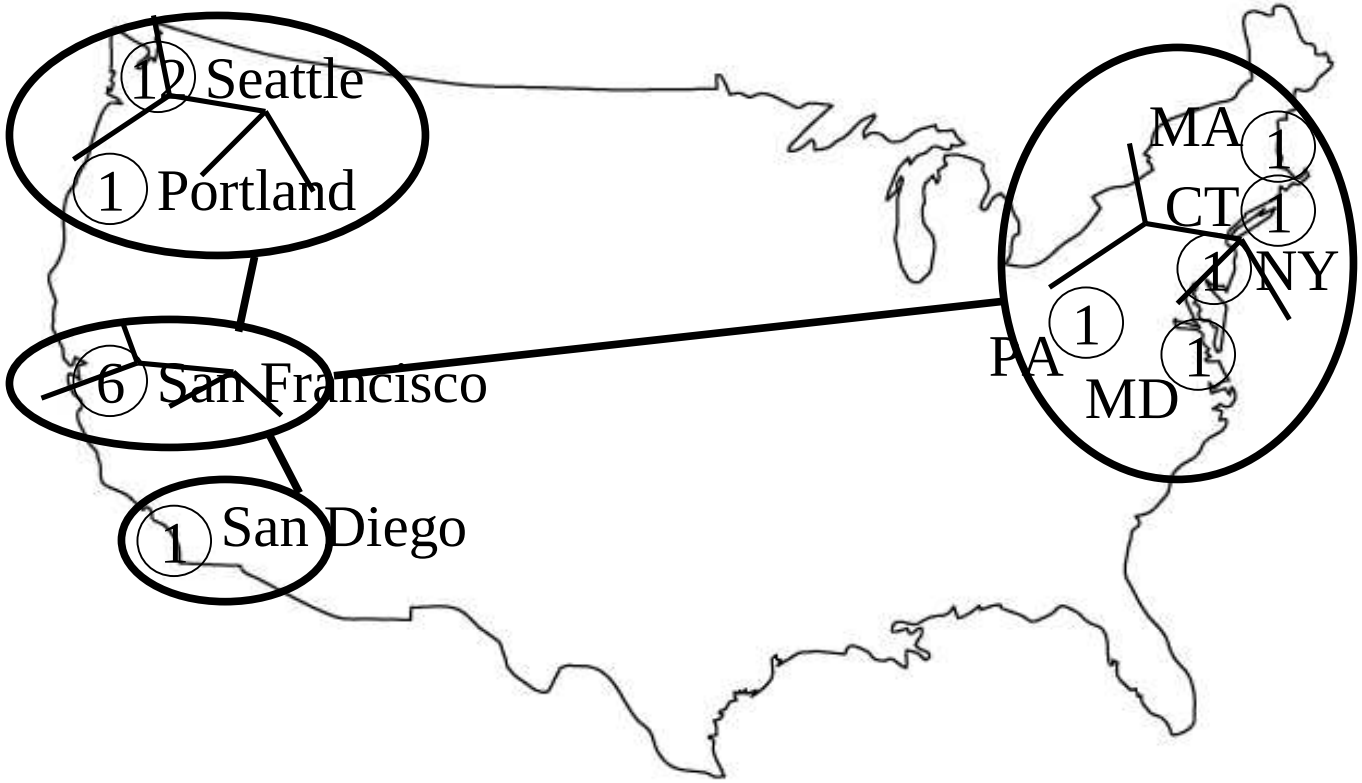

Geographic clustering: Approximation of network clustering

#4: Implications for Overlay Multicast

Multicast tree's root:
- Symmetric bandwidth: 750 kbps
- Location: Seattle

- For achieving delay less than 120ms, max stream is 148kbps

- Low upstream bandwidth limits out-degree considerably

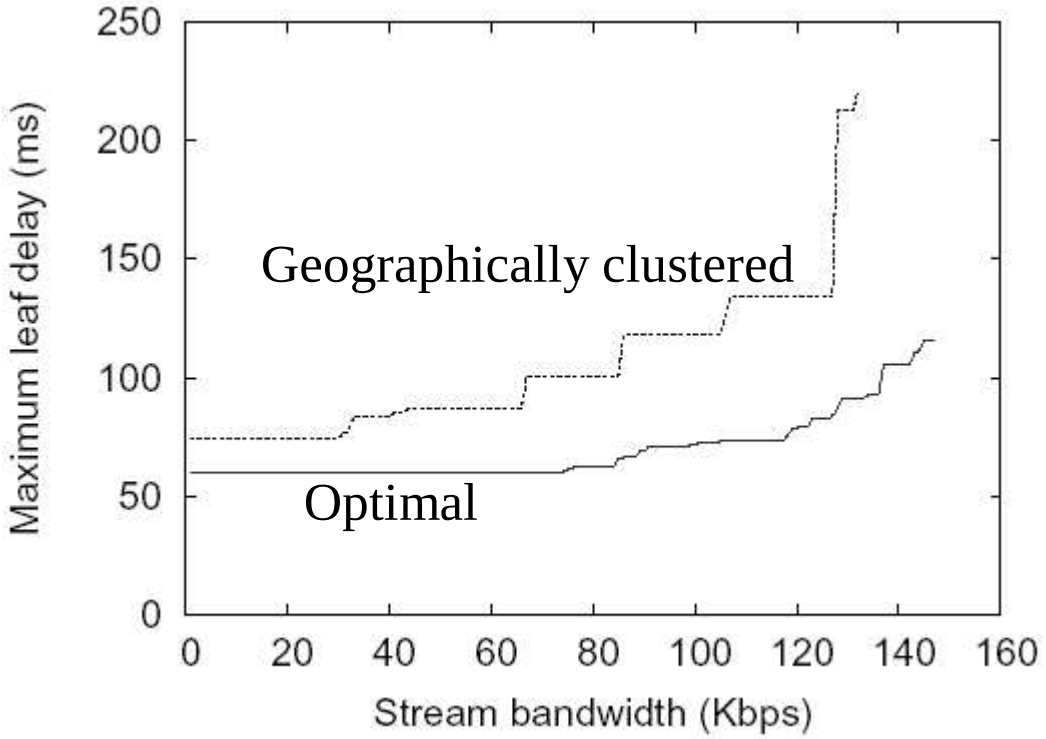

Locality-based heuristics for tree construction perform much worse

Conclusions

- Summary of results:
  - Traditional measurement techniques
    - For example, packet pair techniques need to be revisited
  - Well-accepted design techniques
    - Heuristics for peer selection and multicast tree construction might not work well
    - Some techniques like delay-vector for finding close hosts work well
- Limitations: Due to operational logistics, we had a modest set of 25 hosts to perform the study