A Preliminary Analysis of TCP Performance in an Enterprise Network

INM/WREN'10

Boris Nechaev[1], Mark Allman[2], Vern Paxson[2,3], Andrei Gurtov[1]

[1]Helsinki Institute for Information Technology/Aalto University
[2]International Computer Science Institute
[3]University of California, Berkeley

27 April 2010

Background

- Enterprise traffic remains mostly unexplored
  - Logistically difficult to monitor
  - Enterprises are often viewed as working “well enough”
- Data:
  - Lawrence Berkeley National Laboratory
  - October 2005 – March 2006
  - Captured at switches, often switched to new set of port
  - 351 distinct hosts monitored (≈ 4% of total)
  - 292 million intra-enterprise TCP packets
  - Non-trivial calibration challenges (IMC'09 paper)

Connection status

- Focus only on intra-enterprise traffic
- Used Bro 1.5.1 to reconstruct connection status

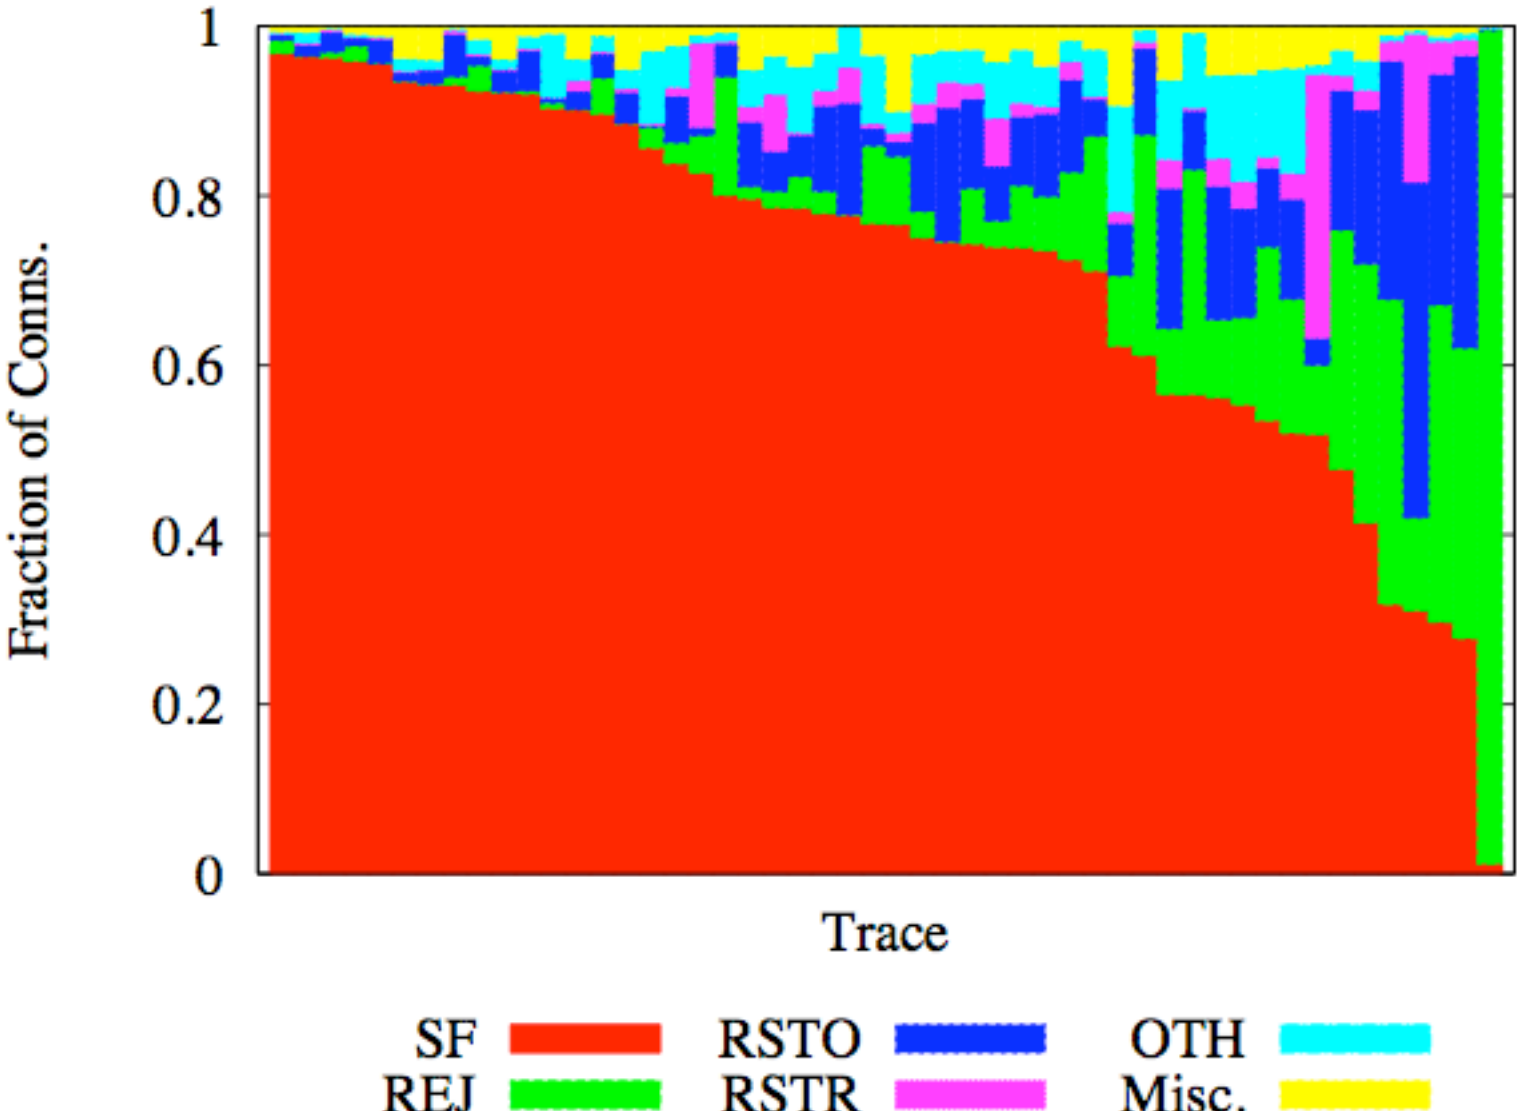

Connection status (cont'd)

- SF + RSTO + RSTR are “good” connections
  - 363K “good” connections (68%)
  - 50 GB of data transferred
  - Consider only these connections in further analysis
  - High percentage variability across traces
- REJ connections
  - Almost all originate at the same host
  - Scanning traffic
- OTH connections
  - Bro observed neither establishment nor teardown
  - Over 90% contain a single ACK or data packet

Connection characteristics

- 44% of connections stay inside the subnet
- Prevalent applications
  - Proportions of bytes/connections are unbalanced
  - Dantz backup: 27% bytes, 0.3% connections
  - HTTP: 9% bytes, 18% connections
  - NetBIOS-SSN: 1.5% bytes, 10% connections
- An application may show heavy tail in connection size or n

App.	Med.	99th	Max
Dantz	6.4 KB	233 MB	4 GB
ssh	5.5 KB	19 MB	2.6 GB
NFS	72 B	1.0 MB	1.1 GB
HTTP	1.9 KB	82 KB	835 MB
NetBIOS-SSN	2.0 KB	59 KB	137 MB
Warewulf	6.6 KB	52 KB	52 KB
Portmap	92 B	716 B	1.1 KB

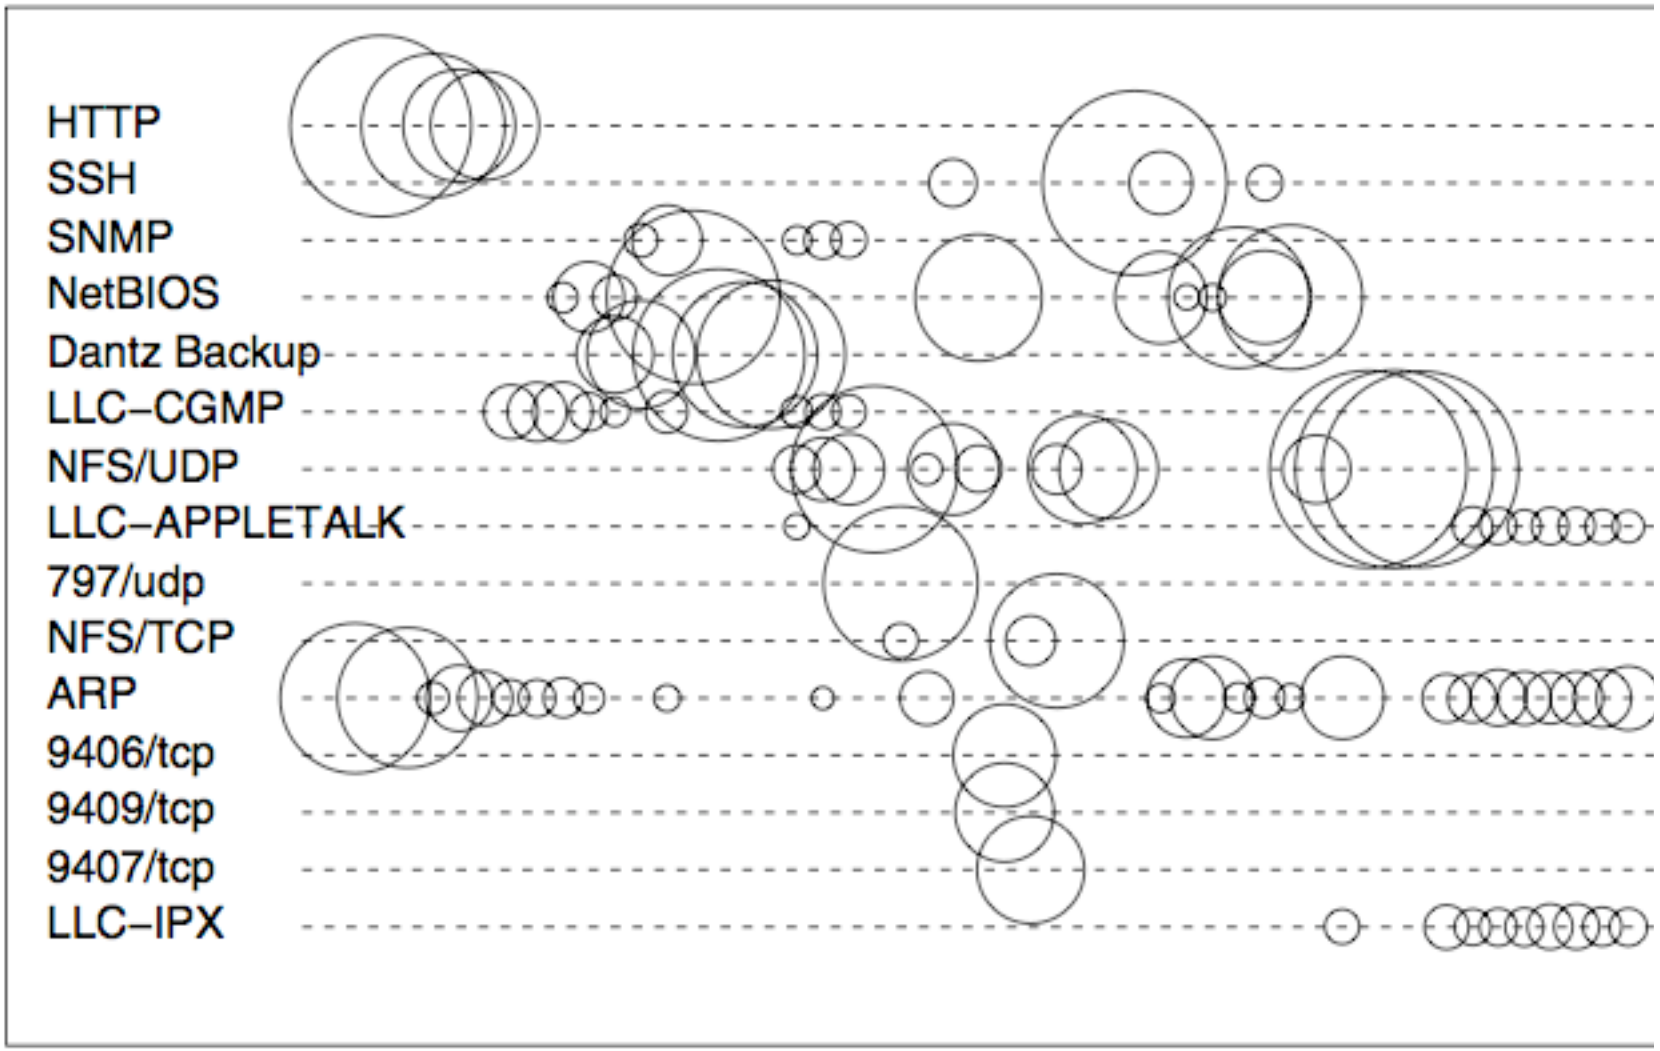

Prevalence in terms of pkts

(a) Subnet

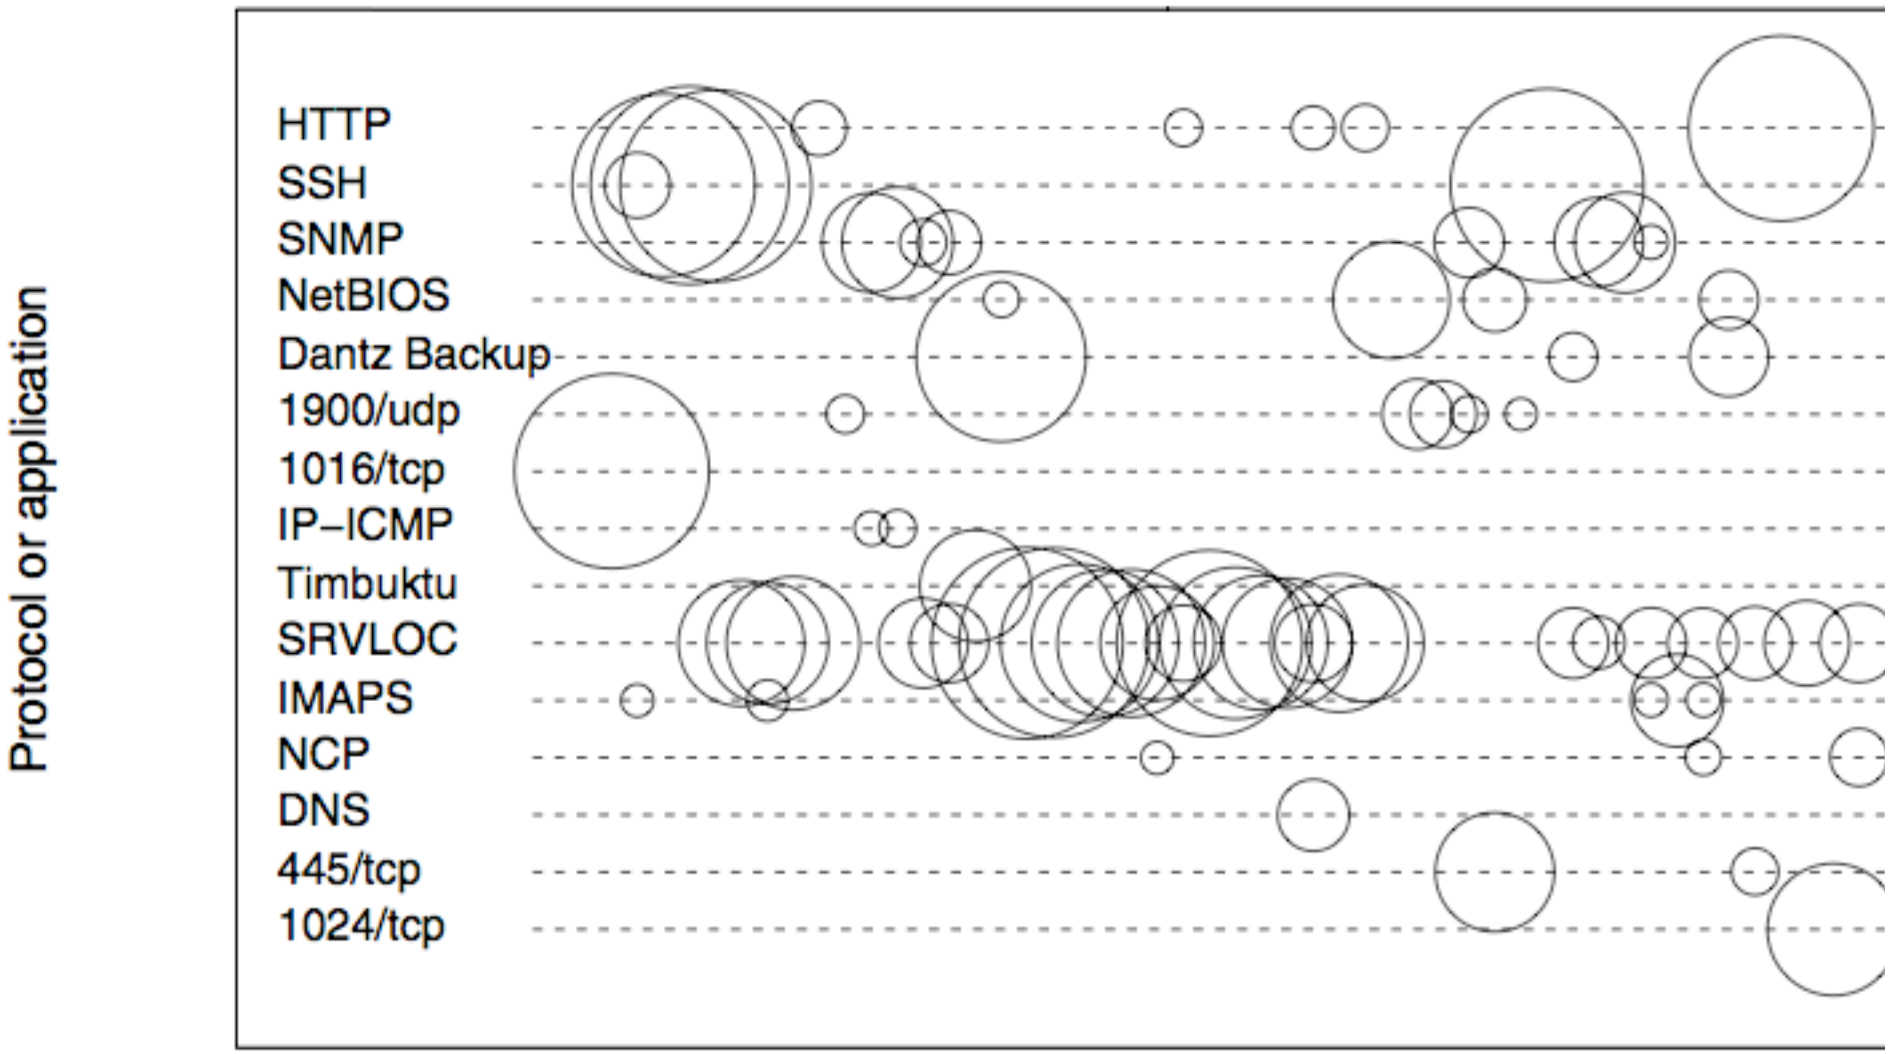

Prevalence in terms of pkts

(b) LBL

Connection characteristics (cont'd)

- Distribution of connection sizes (bytes)
- Ratio of originator data bytes to responder data bytes

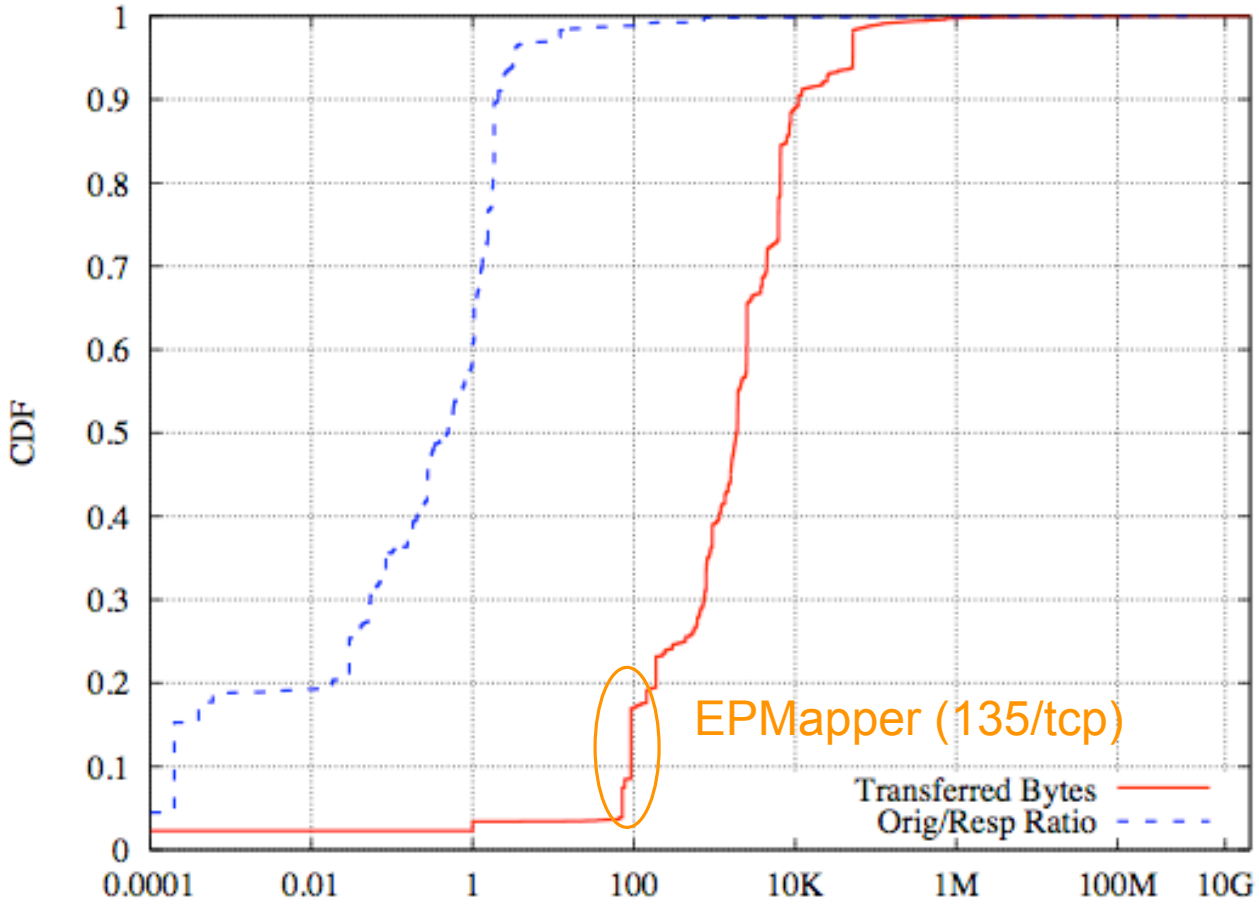

- Median transfer size ~2KB
- 90% of traffic comes from just 160 connections (out of 363l

Performance

- Very low number of packets with bad TCP checksum - 583
- 0.1% connections had packet reordering
- No replicated packets
- 0.5% connections experienced retransmissions
  - (Haven't done fully robust retransmission detection yet)
- Connection maximum flight sizes
  - Median: 214 bytes
  - 99th percentile: 5.3 KB
  - Bandwidth-delay product for 100Mb/s, 1 msec RTT: 12.5k
    - Do we see bandwidth underutilization?

Transfer rates

- Rate = (Total bytes in the connection) / (Duration)

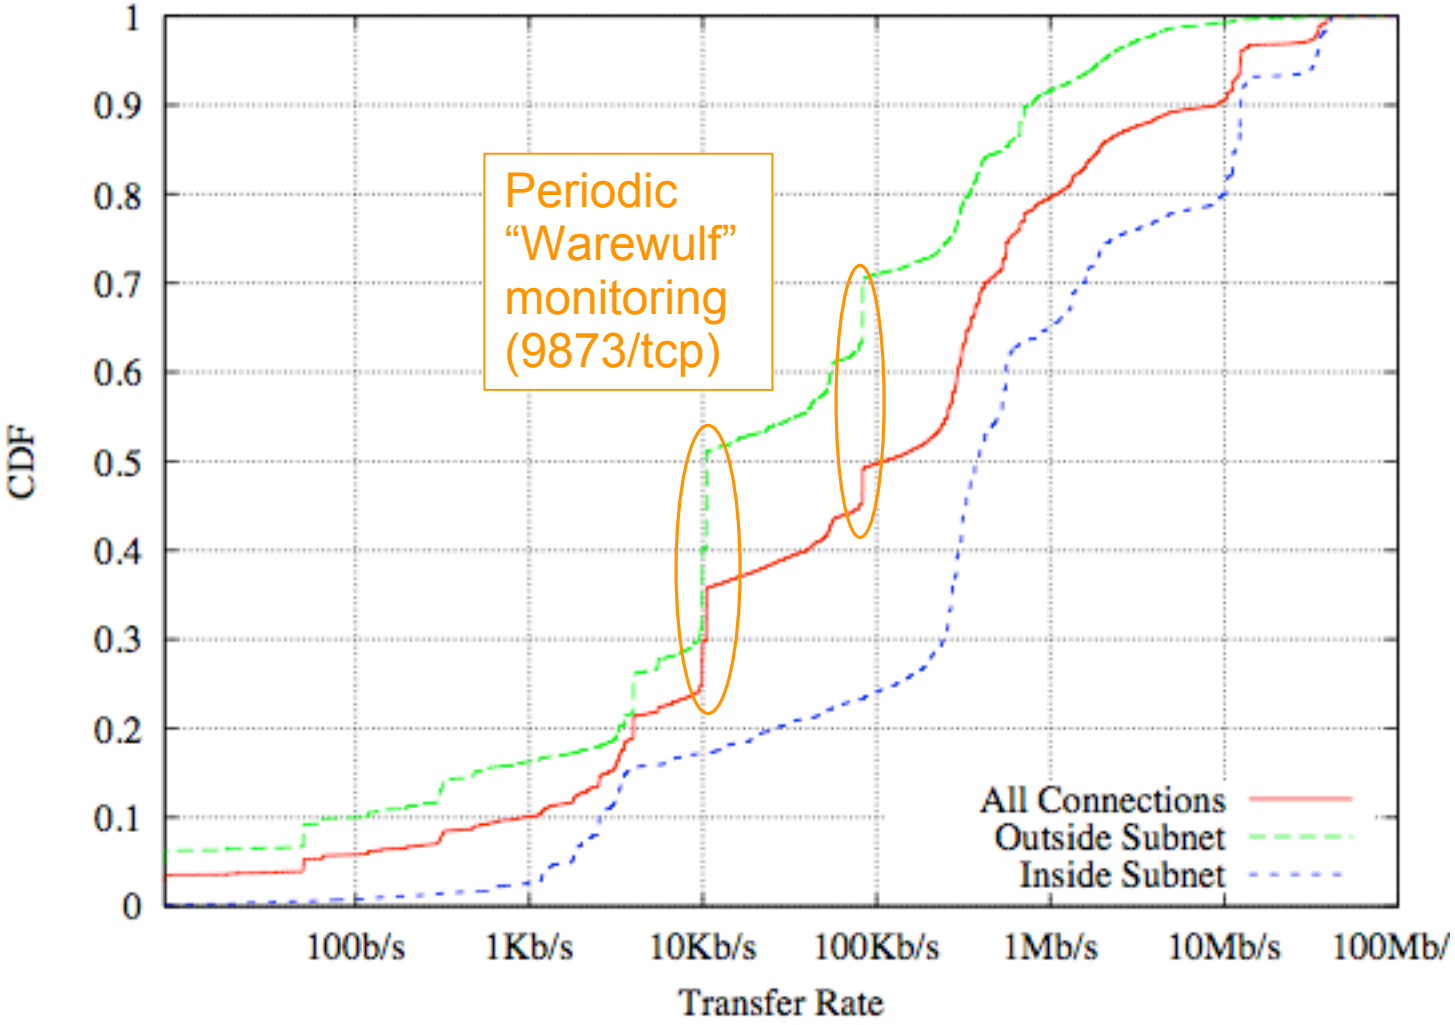

- Higher rates compared to WAN traffic studies
- Intra-subnet rates are 10 times higher than inter-subnet

Transfer rates (cont'd)

- 4 types of flows with 10 KB and 1 sec thresholds

Type.	Conns. (%)	Bytes (%)
Short-Small	57.2	0.6
Short-Large	2.6	0.8
Long-Small	31.8	0.8
Long-Large	8.4	97.8

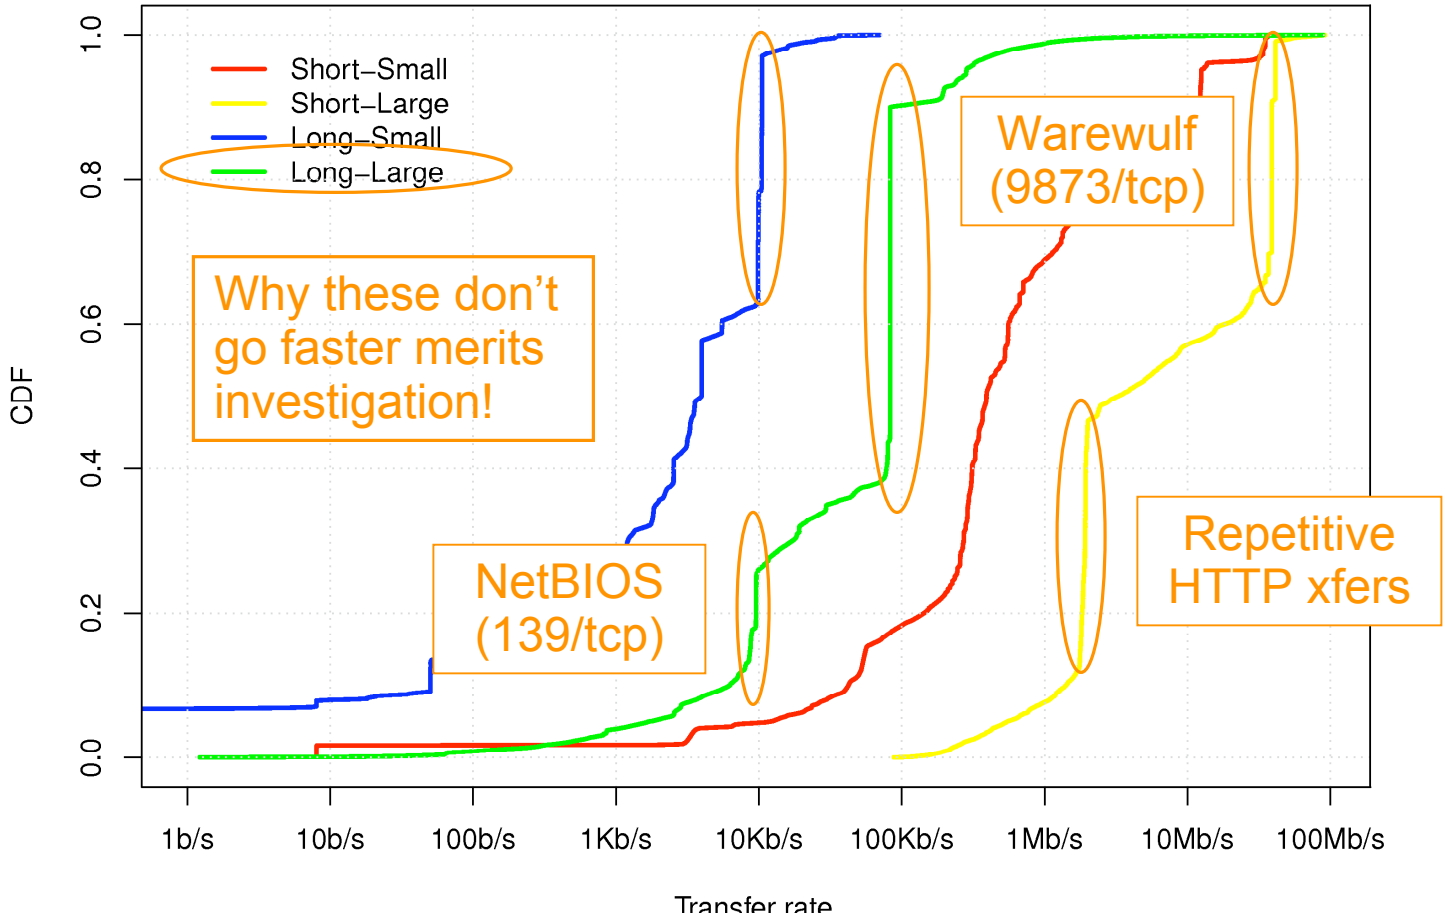

Summary

- Preliminary analysis of TCP performance
- Higher rates than in WAN
- Less loss than in WAN
- In general, enterprise connections appear to work well
  - Are flaws masked by high capacity and low delays?
- Next steps:
  - Analysis of packet latency dynamics
  - Assessment of loss & retransmission behavior
  - In-depth study of bandwidth utilization
  - Incorporation of a large new dataset
    - 1,000 end systems recorded 2009/2010