Lecture 29

CS625: Advanced Computer Networks
Fall 2003

Tuesday, 04 November 2003

Bhaskaran Raman
CSE, IIT-Kanpur

http://www.cse.iitk.ac.in/users/braman/courses/cs625-fall2003/outline.html

Outline for today

- End-to-end Internet packet dynamics [P97]
- Why we don't know how to simulate the Internet [PF97]
- *Scribe for today?*

End-to-End Packet Dynamics

- End-to-end properties are measured
- How? Using TCP traces
- Advantages:
  - Directly performance indicators
  - Not too intrusive: built-in congestion control
- Disadvantages:
  - Need to separate protocol behaviour from network behaviour
  - Difficult/impossible to do frequency analysis

More on the Setup...

- 20,000 100KByte transfers between 35 sites
- Network Performance Daemons (NPDs) at these sites: N^2 effect
- Properties measured:
  - Network pathologies: reordering, replication, corruption
  - Bottleneck bandwidth
  - Packet loss
  - One-way delay

Packet Reordering

- Significant amount of reordering:
  - 2% of all packets in N1
  - 0.6% of all packets in N2
- Data packets reordered more than acks
- Site-specific behaviour: e.g., "ucol" had 15% reordered packets
- Asymmetric behaviour: e.g., only 1.5% packets reordered coming into "ucol"
- Reordering due to router fluctuation – also due to router forwarding lulls

Effect of Reordering

- TCP detects "spurious" packet loss
- Reordering time-scales:
  - Upto 81ms
  - Waiting 20ms detects 70% reordering in N1
  - Waiting 8ms sufficient for this in N2
- Define R(g:b) – ratio of good to bad reordering runs
  - 22 for N1, 300 for N2
  - When dup_thr = 4?
  - When dup_thr = 2?

Packet Replication and Packet Corruption

- Replication: negligible
- Packet corruption:
  - 1/5000 in both datasets
  - Lesser in ACKs
  - 16-bit checksum in TCP may not be enough!

Bottleneck Bandwidth

- Different from available bandwidth
- Estimated using self-interference: packet pair
  - Sender only
  - Receiver-based mechanism
- Sources of error:
  - Out-of-order, Clock resolution, Instability, Multi-channel routing
- Use packet-bunch-modes (PBM)
- Receiver support required for good estimates

Packet Loss

- 2.7% in N1, 5.2% in N2
- Quiescent and busy periods
- Correlation over time periods of hours!
- Data loss is much more than ACK loss
- Loss rate is asymmetric/uncorrelated
- Loss outages span several time scales
- Efficacy of TCP retransmissions:
  - Retransmissions: unavoidable, due to coarse feedback (mostly this), due to bad RTO

Simulating the Internet

- Heterogeneity:
  - Topology, link properties, protocol differences
  - Problems with trace-driven simulations
- Changes over time:
  - Pricing structure, routing protocols, topology, caching, different killer application
- Scale: limits to simulation
- How to cope?
  - Model invariants (difficult to find)
  - Explore parameter space

Further Topics

- Web caching