Network Measurement

CSE 561 Lecture 10, Spring 2002.
David Wetherall

Overview

- Motivation – why measure?
- What would we like to measure?
- Measurement approaches / methodologies
- Challenges

- Bolot93 – Delay and Loss
- Leland93 – Self-Similar Traffic

Motivation

- Operational needs
  - Is something broken? What is it?
  - Do I need more / better X? When will I need it?
- Research needs
  - How is the Internet really configured?
  - How well does it do X?
  - How do people use it?
  - What is the trend for X?
  - How do these answers impact application / protocol design?
- Underlying assumption
  - We don't really understand how networks work / are used
  - If we did, then we could use simulation or analytic means

What to measure?

- **Channel properties**
  - Corruption, sequencing, duplication
  - Latency, loss, jitter, bandwidth
- **Topological properties**
  - Network connectivity (physical, L2, L3, AS-level)
  - Routing protocol dynamics
- **Application properties**
  - Traffic composition
  - Request distributions, actor locations

Methodologies

- **Active vs Passive**
  - Active: send a probe into the network and see what happens
    - E.g., ping, traceroute to measure paths
  - Passive: observe existing traffic to determine result
    - E.g., Web traces to measure caching behavior
    - E.g., TCP traces to measure bandwidth etc.
    - E.g., RouteViews peers with routers to observe BGP routes

- **Observation vs inference**
  - Few things can be observed; statistical inference is key
  - E.g. ping: round-trip time is observed, packet loss in inferred

Challenges – The Experiment

- **You can't measure what you want where you want!**
  - No centralized points of control / administration
  - No / little cooperation from intermediate systems
  - Little / no cooperation from end systems: one or two armed?

- **Result is a need/emphasis on creative inference**
  - Interior properties based on E2E observations
  - E2E properties using "stealth"end-system support
    - e.g., Sting hides via TCP, King via DNS

Challenges – Getting Good Data

- **Representative data**
  - Internet is huge and heterogeneous
  - Good trace data is hard to come by / protected
- **Technical difficulties**
  - High speed passive measurement is hard
  - Active probes treated differently from normal data (ping)
  - Privacy concerns; encryption obscures structure
  - Asymmetry; may only be able to monitor one direction
  - Repeatability

Challenges – Using the Data

- **Good metrics & statistics**
  - What to measure
    - Flow vs bytes vs packets
  - How to summarize sample data?
    - Mean, median, standard assumptions, heavy-tails, etc…
  - Validation
    - How do you know you didn't make a mistake?
- **Uncertain predictive power**
  - Adaptive on short-time scales; changing on longer ones
  - How valuable is yesterdays' measurement?

Bolot93 – E2E Delay and Loss

- Characterizes E2E packet delay and loss
  - Active E2E observations only, no network access
- Infers properties of the path from measurements
  - By relating properties of measurements to analyses

- Bottleneck bandwidth falls out
  - Probe / ACK compression seen too
- Dependent / grouped losses fall out
  - On top of random background losses

Other Inference Techniques

- Bottleneck bandwidth (pchar, bprobe, nettimer, clink)
- Available bandwidth (treno, ?)
- Path loss, reordering in both directions (sting)
- Loss before or after bottleneck (paxson)
- Queuing delays (Vegas?)
- Location of congestion (Andy?)
- Topology (traceroute, ally, Neil, 561?)
- Link weights (Ratul)
- Routing policy (Gao, 561?)
- Latency between arbitrary points (king)

- Where will we be in a few years?

Leland93 – Self Similar Traffic

- Meticulous analysis of traffic timings
  - Shows traffic is self-similar (bursty across a wide range of timescales)
  - Burstiness (Hurst parameter) gets worse with load!

- What does this mean?
  - Aggregated traffic does not get smooth; departure from telco design

- Intuitive construction
  - Combine ON / OFF sources with heavy-tailed periods
  - Result is self-similar traffic

In a similar vein …

- File / flow sizes are heavy-tailed
- Document popularity is Zipf
- Web transfer times, rates are heavy-tailed

- These have implications for system design
  - Average doesn't characterize much
  - A small number of flows carry most of the bytes!
  - Exploit for load-sensitive routing, penalty boxes, …
  - (Cooperative) caching is of limited benefit

Summary

- Network measurement is easy to do … but hard to do right!

- Need to be creative about collecting data and inferring quantities
- Need to be careful about collection and analysis methodology
- Need to consider the underlying causes

- Two kinds of results
  - Lots of raw results: “Good data outlives bad theory”
  - A few important conclusions: “Web page popularity is heavy tailed, so the benefit of caching is limited”