Internet Research Needs a Critical Perspective Towards Models

– Sally Floyd

– IMA Workshop, January 2004

“Computer System Performance Modeling and Durable Nonsense”

- “A disconcertingly large portion of the literature on modeling the performance of complex systems, such as computer networks, satisfies Rosanoff's definition of durable nonsense.”

- "THE FIRST PRINCIPLE OF NONSENSE:

  For every durable item of nonsense, there exists an irrelevant frame of reference in which the item is sensible."

- "THE SECOND PRINCIPLE OF NONSENSE:

  Rigorous argument from inapplicable assumptions produces the world's most durable nonsense."

- "THE THIRD PRINCIPLE OF NONSENSE:

  The roots of most nonsense are found in the fact that people are more specialized than problems"

The quote is 25 years old!

- John Spragins, "Computer System Performance Modeling and Durable Nonsense", January 1979.

- R. A. Rosanoff, "A Survey of Modern Nonsense as Applied to Matrix Computations", April 1969.

The questions of this talk:

- Do we understand how our modeling assumptions affect our results?
- Do we know how our modeling assumptions affect the relevance of our results for the (current or future) Internet?
- What kind of tools do we need to help improve our understanding of models?

Assumptions:

- For each research topic, we want a model that is as simple as possible, but no simplier.
- Models underlie simulations, experiments, analysis, and pure thought experiments.
- For the fast-changing and heterogeneous Internet, determining the relevant model for a particular research question can be 95% of the work!

Topic: Active Queue Management Performance

- Research question: tradeoffs between throughput and delay.
- One model: One-way traffic, one RTT, long-lived and small flows but no medium-sized flows.
  - Result: High throughput and low delay is possible.
- Alternate model: Two-way traffic, range of RTTs, wide range of flow sizes.
  - Result Bursty traffic, throughput/delay tradeoffs.

Topic: AQM Performance

- Question: What do we know about the actual characteristics of aggregate traffic at congested links in the Internet?
  - Distribution of flow sizes?
    - Extensively studied.
  - Distribution of round-trip times?
    - Some measurements available.

Distribution of Flow Sizes

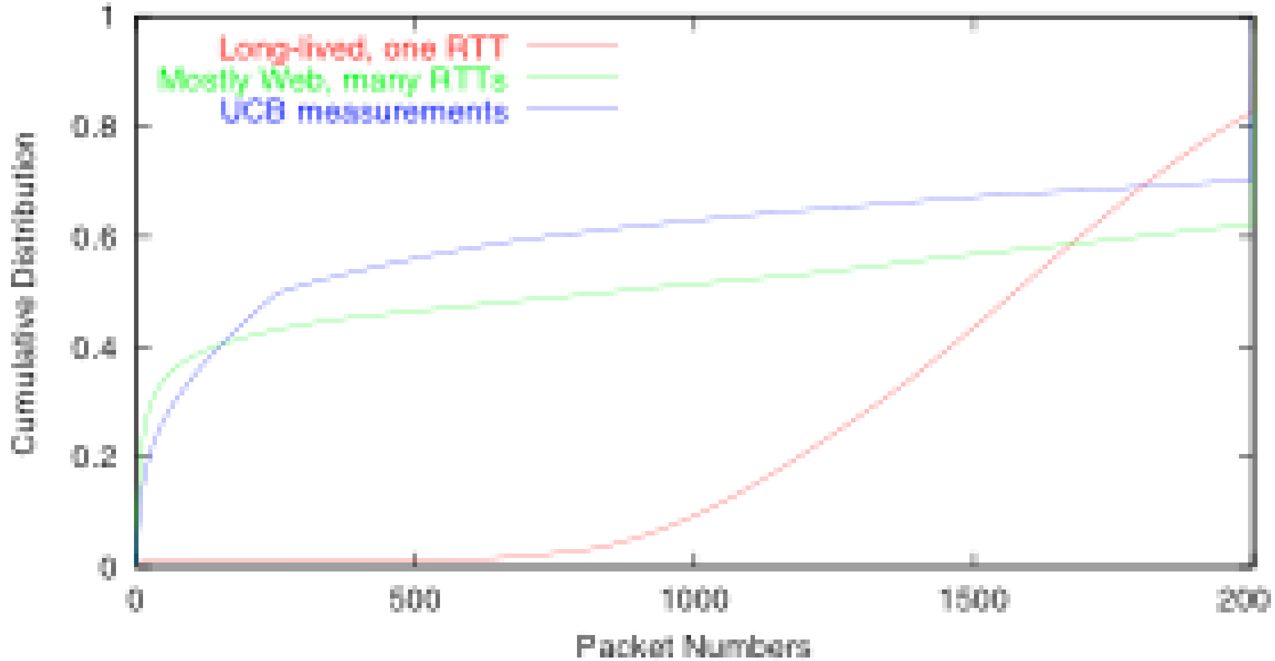

- Distributions of packet numbers on the congested link over the second half of two simulations, with data measured on the Internet for comparison.

Distribution of RTTs:

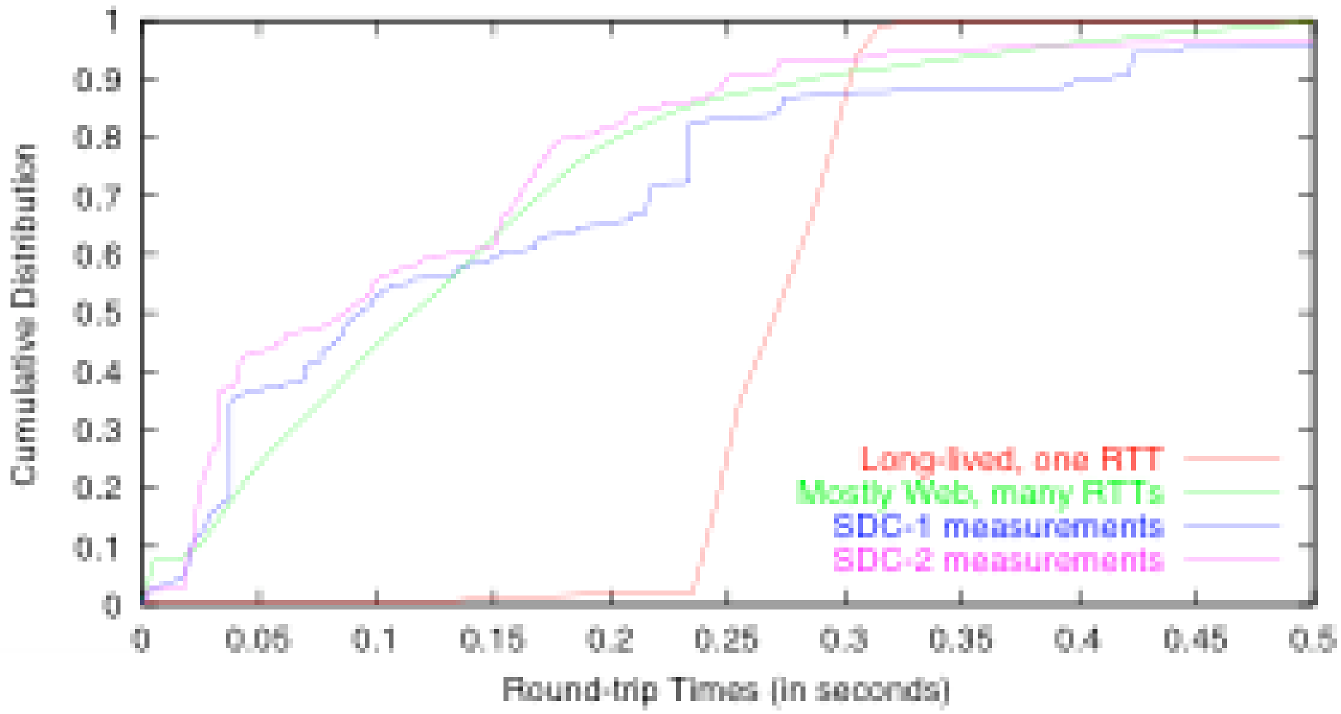

- Distributions of packet round-trip times on the congested link of two simulations, with data measured on the Internet for comparison.

Topic: AQM Performance

- Question: What do we know about the actual characteristics of aggregate traffic at congested links in the Internet?
  - Typical levels and patterns of congestion?
    - Congestion at access links, moderate levels of congestion?
    - Tools for measuring from TCP traces.
  - Reverse-path congestion?
    - Little is known.
  - How many flows are limited by end nodes or by other access links?
    - Some measurements.

Topic: Dynamics of HighSpeed TCP

- Research topic: convergence times (for new TCP flows competing against existing flows).
- Model #1: DropTail queues, global synchronization when packets are dropped.
- Model #2: DropTail queues, some synchronization, depending on traffic mix.
- Model #3: RED queues, no synchronization.
- Which model is the best fit for the current Internet? For the future Internet?

Topic: Transport protocol performance over wireless links

- Characteristics of wireless links that affect transport protocol performance:
  - Packet loss due to corruption.
  - Delay variation due to link-layer error recovery, handovers, and scheduling.
  - Asymmetric and/or variable bandwidth (e.g., satellite).
  - Shared bandwidth (e.g., WLANs).
  - Complex link-level buffering (e.g., cellular links).
  - Mobility.

Topic: Transport protocol performance over wireless links

- Tools: The NS simulator has tools for modeling wireless links; we (Andrei Gurtov) has added to them.
- There is an interplay between wireless link mechanisms and transport protocols, with both changing.
  - E.g., corruption is often repaired at the link layer.
- It is challenging to try to characterize relevant models for the current and future Internet.

Topic: The Evolvability of the Internet Infrastructure

- Research topics:
  - How do we understand the current limits to evolvability of the Internet infrastructure?
  - What would be the impact of different architectural changes on the evolvability of the Internet infrastructure?
    - E.g., security vs. evolvability
    - Communication between layers vs. evolvability.

Topic: The Evolvability of the Internet Infractructure

- What conceptual models do we use to help understand this?
- Standard models of complex systems have many limitations:
  - E.g., game theory;
  - Physics models;
  - Control theory and dynamical systems;
  - …

Topic: The Evolvability of the Internet Infrastructure

- Critical aspects of a conceptual model for this topic:
  - The layered IP architecture;
  - Changes over time (e.g., overprovisioning);
  - A decentralized system with many players (companies, ISPs, standards bodies, etc.);
  - Economic and political factors;
  - Chicken-and-egg deployment issues.

References:

- S. Floyd and V. Paxson, “Difficulties in Simulating the Internet” , Transactions on Networking, August 2001.
- S. Floyd and E. Kohler, “Internet Research Needs Better Models”, HotNets-I, October 2002.
- A. Gurtov and S. Floyd, “Modeling Wireless Links for Transport Protocols”, November 2003.To appear in CCR.

References, cont.

- S. Floyd, “Modeling the Internet as a Complex System”, viewgraphs, End-to-End Research Group, January 2003.