Measurements, Models, and Simulation Scenarios for Internet Research

Sally Floyd and Eddie Kohler

NSF CISE/SCI PI Meeting

February 2004.

(Adapted from an earlier talk, "Internet Research Needs a Critical Perspective Towards Models", January 2004 IMA workshop on "Measurement, Modeling, and Analysis of the Internet.")

“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 guiding this research:

- 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 simpler.
- 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.
- Model #1: Mostly one-way traffic, small range of RTTs, long-lived and small flows but few medium-sized flows.
  - Result: High throughput and low delay is possible.
- Model #2: Two-way traffic, wide range of RTTs, wide range of flow sizes.
  - Result: Bursty traffic, throughput/delay tradeoffs.

Throughput vs. Queue Size

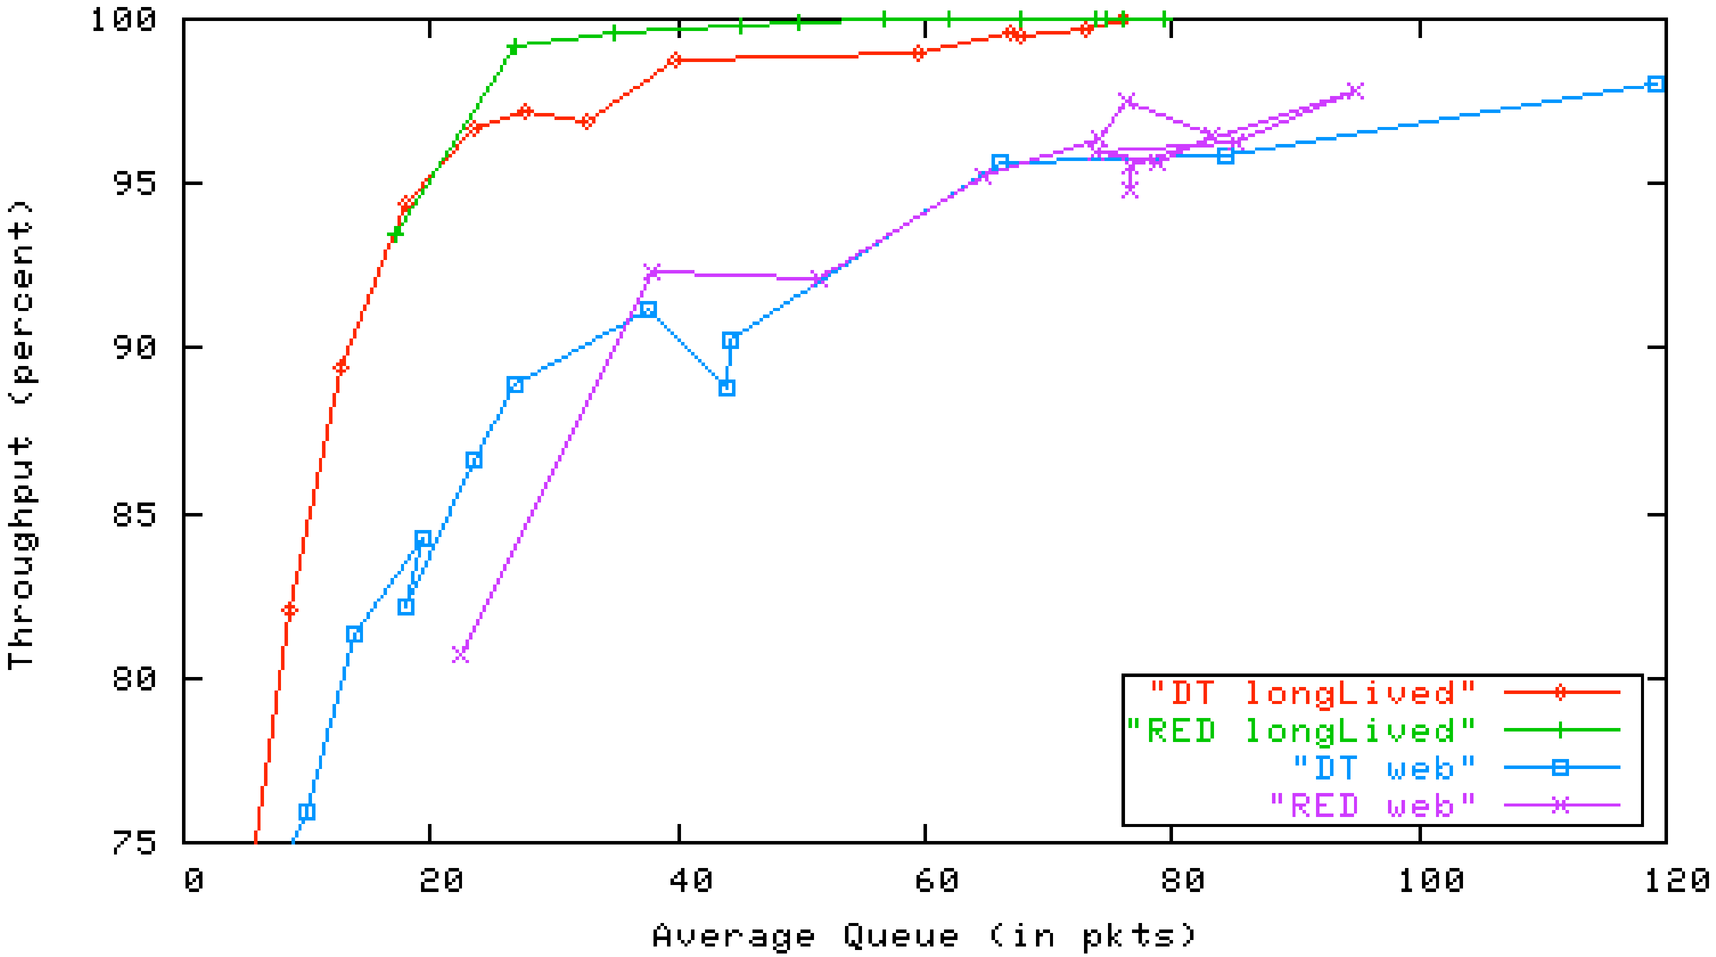

Packet Drop Rates

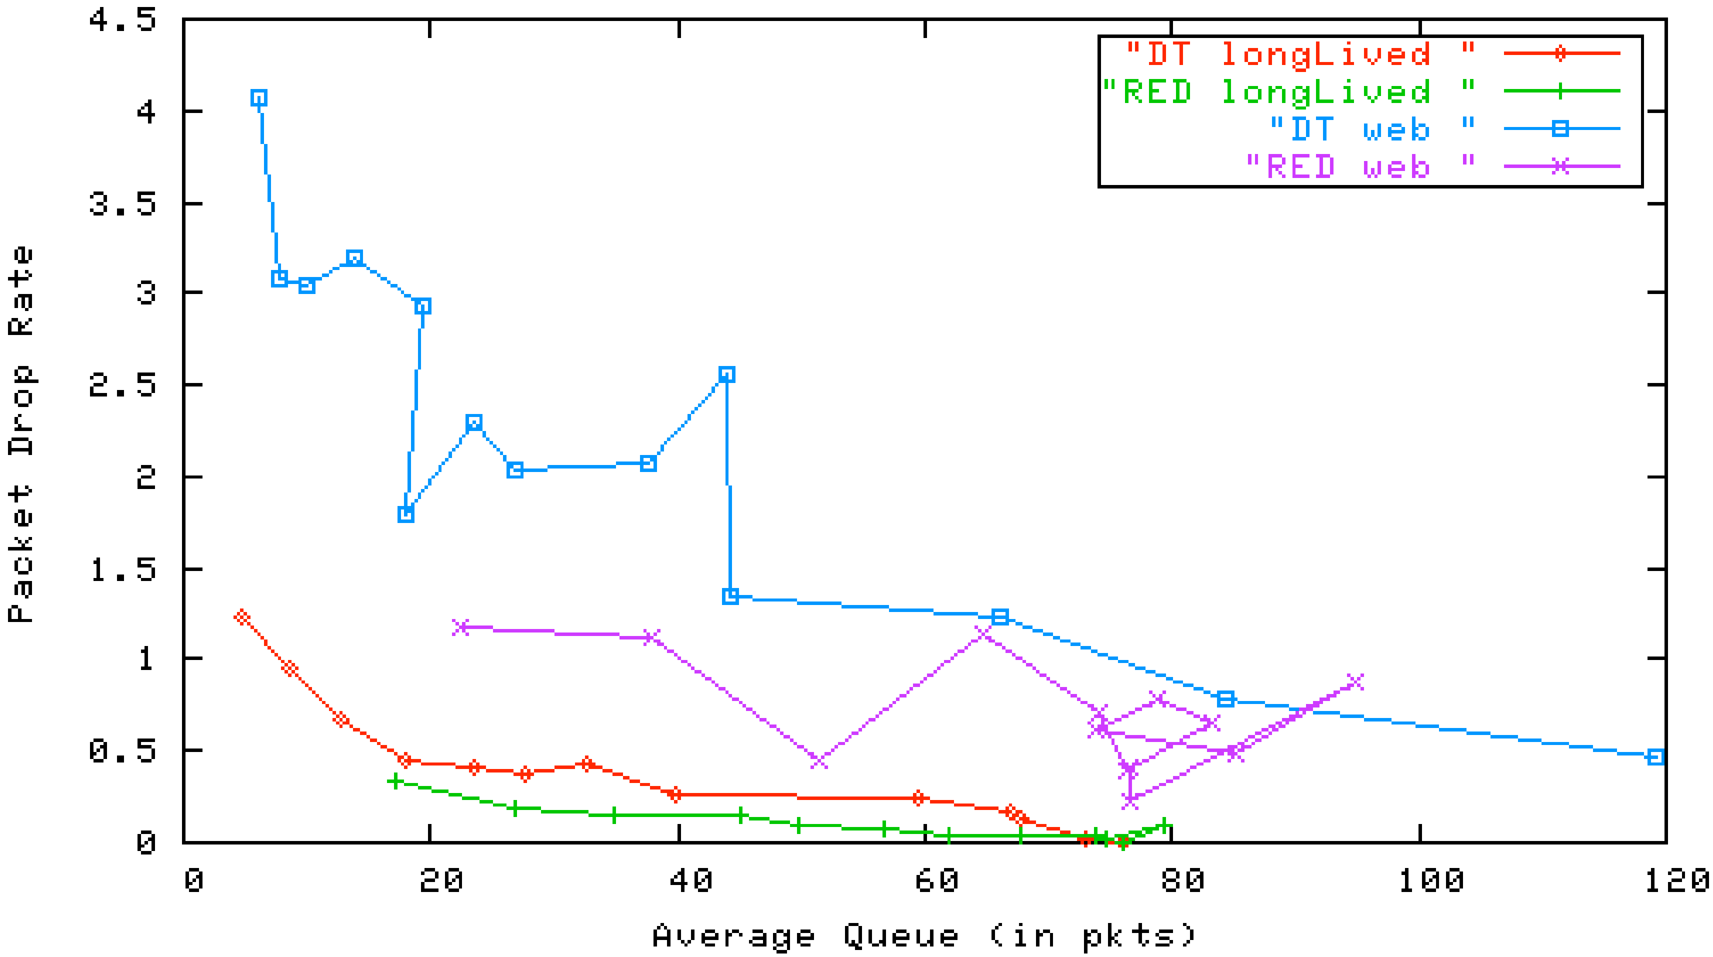

Topic: AQM Performance

- 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.
  - We have added simple tools to plot these distributions in NS simulations as well.

Distribution of Flow Sizes

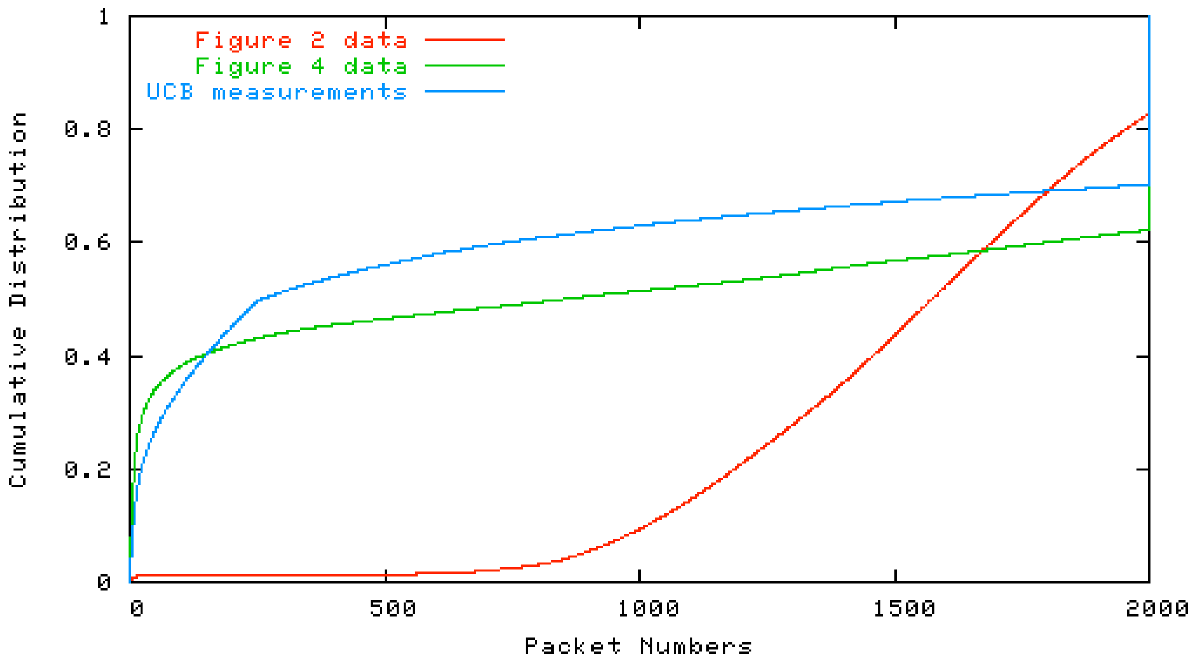

- Distributions of packet numbers on the congested link over the second half of two simulations, with data measured on the Internet for comparison.

Distribution of RTTs

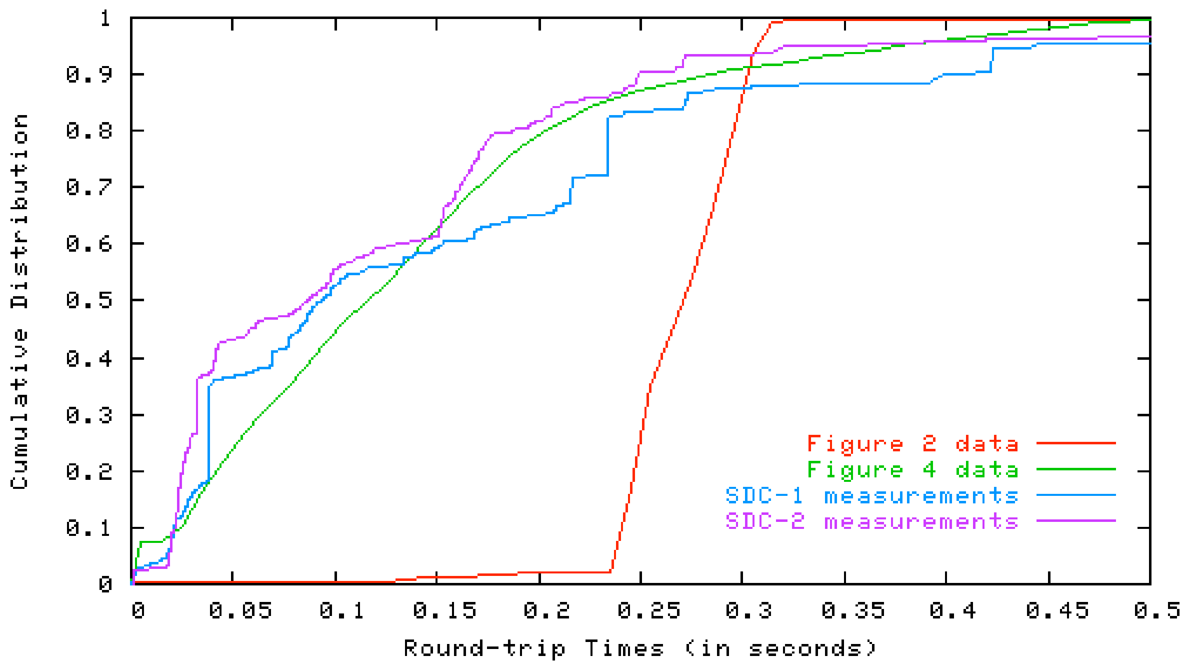

- Distributions of packet round-trip times on the congested link of two simulations, with data measured on the Internet for comparison.

Topic: AQM Performance

- Characteristics of aggregate traffic at congested links that we don’t understand very well:
  - Typical levels and patterns of congestion?
    - Congestion at access links, moderate levels of congestion?
    - Tools for measuring from TCP traces.
    - We also have some new tools and measurement results.
  - Reverse-path congestion?
    - Little is known.
  - How many flows are limited by end nodes or by access links?
    - Some measurements.

Topic: Evaluating assumptions with measurements

- How to answer these questions?
- A program of ongoing, large-scale, representative Internet measurement
- Different from application-directed measurement
  - Not just the available bandwidth, but the bottleneck capacity bandwidth
  - Not just the narrowest link, but any congested links on the path
  - Passive, trace-based → less intrusive, run on old traces to measure network evolution

Tools for measurements

- MultiQ: detects multiple bottleneck capacities and their order.
  - Building on a mature collection of tools for measuring bottleneck capacity (e.g., nettimer, pathrate).
- Mystery: robustly measures loss events, packet losses, and RTT changes.
  - Related tools: T-RAT, tcpanaly, etc.
- With passive measurements, multiple tools can be applied to each data set (and to old data sets).

Measurement studies

- Evolution of bottleneck capacity:
  - increased by an order of magnitude from 2002 to 2004
- Statistical multiplexing:
  - Level increased, from 2002 to 2004, so that fair-share bandwidth remained relatively stable.
- RTT changes around loss events.
- Loss event rate vs. bottleneck link capacity.

Study: Bottleneck capacity evolution

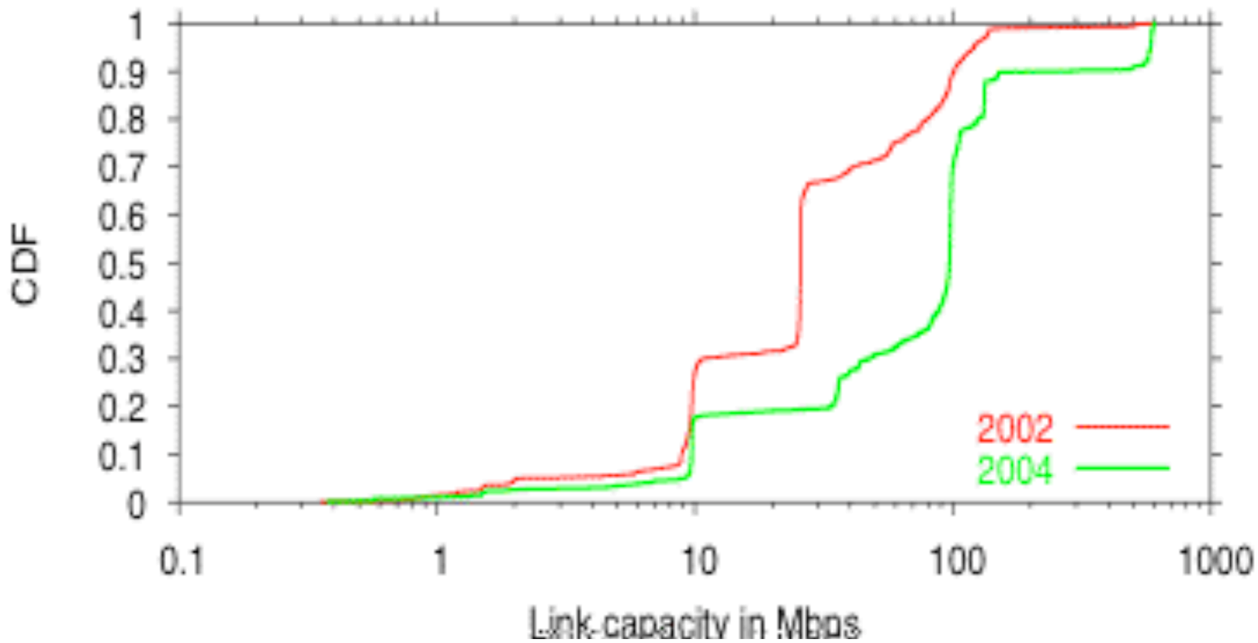

- CDF of bottleneck capacities in NLANR traces from 2002 and 2004
  - Median capacity goes up by 5x

Study: Loss rate vs. bottleneck capacity

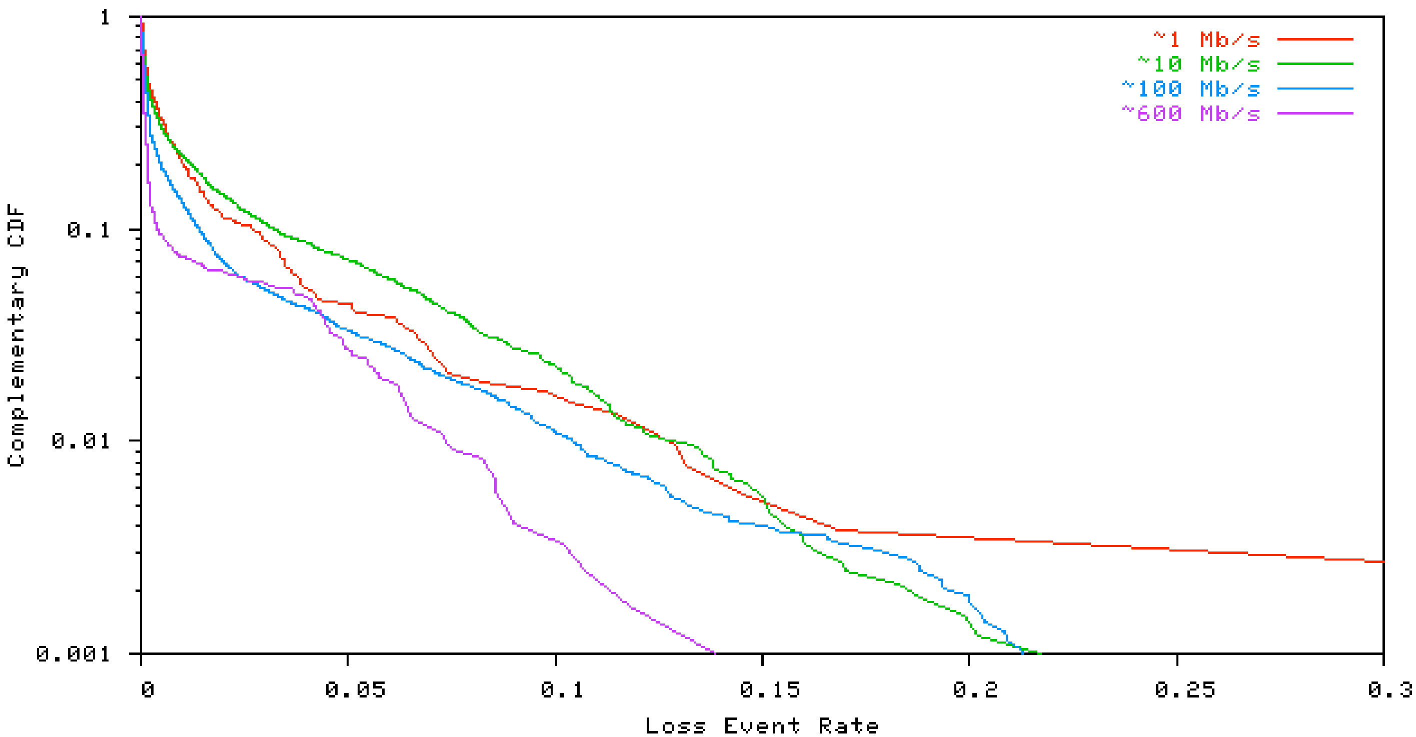

- CCDF of loss event rate (TFRC definition) for all flows with bottleneck capacity *c*
  - 10 and 100 Mb/s bottlenecks have same range of loss event rates

Topic: Dynamics of HighSpeed TCP, Scalable 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, some synchronization.
- Model #4: 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: Andrei Gurtov has added to NS's tools for modeling wireless links, with simulation scenarios for using these models.
- There is an interplay between wireless link mechanisms and transport protocols, with both changing and adapting to the other.
  - E.g., for exploring transport protocols over wireless links, one could look at:
    - older wireless link models with little FEC or link-level retransmissions;
    - or, more current models with link-level repair of corruption;
    - or, models of future wireless links?

Conclusions: Questions

- How do our models affect our results?
- How do our models affect the relevance of our results to the current or future Internet?
- What kinds of tools do we need to improve our understanding of models?

Papers:

- Sachin Katti, Charles Blake, Dina Katabi, Eddie Kohler, and Jacob Strauss, "M&M: Passive Measurement Tools for Internet Modeling", January 2004, under submission.
- A. Gurtov and S. Floyd, “Modeling Wireless Links for Transport Protocols”, November 2003.To appear in CCR
- S. Floyd and E. Kohler, “Internet Research Needs Better Models”, HotNets-I, October 2002.
- S. Floyd and V. Paxson, “Difficulties in Simulating the Internet” , Transactions on Networking, August 2001.

Simulation Scripts:

- Andrei Gurtov, "NS Simulation Tests for Modeling Wireless Links", directory tcl/ex/wireless-scripts in the NS simulator.
- Simulation scripts for distributions of packet numbers and flow sizes:
  – “http://www.icir.org/models/sims.html”.
- Simulation scripts for the distributions of packet numbers and flow sizes:
  – http://www.icir.org/models/sims.html”.

Webpages

- Internet Research Needs Better Models.
- Building Models for Aggregate Traffic on Congested Links.
- Network Simulators.
- Traffic Generators for Internet Traffic.
- Topology Modeling.
- Measurement Tools for Bandwidth Estimation, Estimating Loss Rates, etc.
- …
- From "http://www.icir.org/models/bettermodels.html".

Papers in Progress:

- Models for the Design and Evaluation of Active Queue Management.

- Models for the Design and Evaluation of Transport Protocols.

Extra Viewgraphs:

•

More on MultiQ and Mystery

- MultiQ:
  - The packet interarrival times at the receiver reflect the sizes of cross-traffic bursts at congested routers.
  - Modes in the distribution correspond to bursts of one or more 1500-byte packets.
- Mystery:
  - Uses ACK timing to distinguish false retransmissions (e.g., reordering, spurious timeouts) from true loss events.

Topic: The Evolvability of the Internet Infrastructure

- Research topics:
  - How do we understand the current limits to evolvability of the Internet infrastructure?
    - Evolvability for applications, qualities of service, forms of group communications, transport protocols, etc.
  - 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.
    - Fragility & complexity & robustness spirals.

Topic: The Evolvability of the Internet Infractructure

- What conceptual models do we use to help understand this?
- Standard models of complex systems have contributions, but also limitations:
  - Game theory;
  - Physics models;
  - Biological models of evolution;
  - Control theory and dynamical systems;

Topic: The Evolvability of the Internet Infrastructure

- Key aspects of conceptual models for this topic:
  - The layered IP architecture;
  - Feedback loops (e.g., TCP);
  - Change over time (e.g., overprovisioning);
  - Tussles: a decentralized system with many players (companies, ISPs, standards bodies, etc.);
  - Economic and political factors (e.g., pricing);
  - Chicken-and-egg deployment problems (e.g., ECN, IPv6, multicast, diffserv).