Maintaining a Critical Attitude Towards Simulation Results

Sally Floyd

WNS2 Workshop

Pisa, Italy

October 2006

http://www.icir.org/floyd/talks/WNS2-Oct06.ppt

Outline for talk:

- The simulator:
  - Validation tests as a tool for verifying functionality.
  - NS-2 and NS-3.
- Using simulations for network research:
  - What is the goal?
  - What are the metrics?
  - What are the simulation scenarios?
- Do the results need to be further validated?
- An example approach: TMRG
  - the Transport Modeling Research Group.

Question #1: the Simulator.

- (1) How much is the software in the simulator to be trusted?
- What is the responsibility of the researcher to verify for themselves that the software in the simulator performs as expected by the researcher?

A quote from the ns-2 web page:

- **“Read this first:**

  While we have considerable confidence in ns, ns is not a polished and finished product, but the result of an on-going effort of research and development. In particular, bugs in the software are still being discovered and corrected. Users of ns are responsible for verifying for themselves that their simulations are not invalidated by bugs. We are working to help the user with this by significantly expanding and automating the validation tests and demos.”

  “Similarly, users are responsible for verifying for themselves that their simulations are not invalidated because the model implemented in the simulator is not the model that they were expecting. The ongoing Ns Manual should help in this process.”

The Validation Tests in NS:

- The role of the validation tests:
  - Helping the programmer to verify that the protocol and protocol parameters work as intended.
  - By validating against saved output, verifying that the protocol still works as intended after subsequent changes to the simulator.
  - Visually helping the user to understand the behavior of the protocol and protocol parameters as implemented in NS.

Validation Tests:

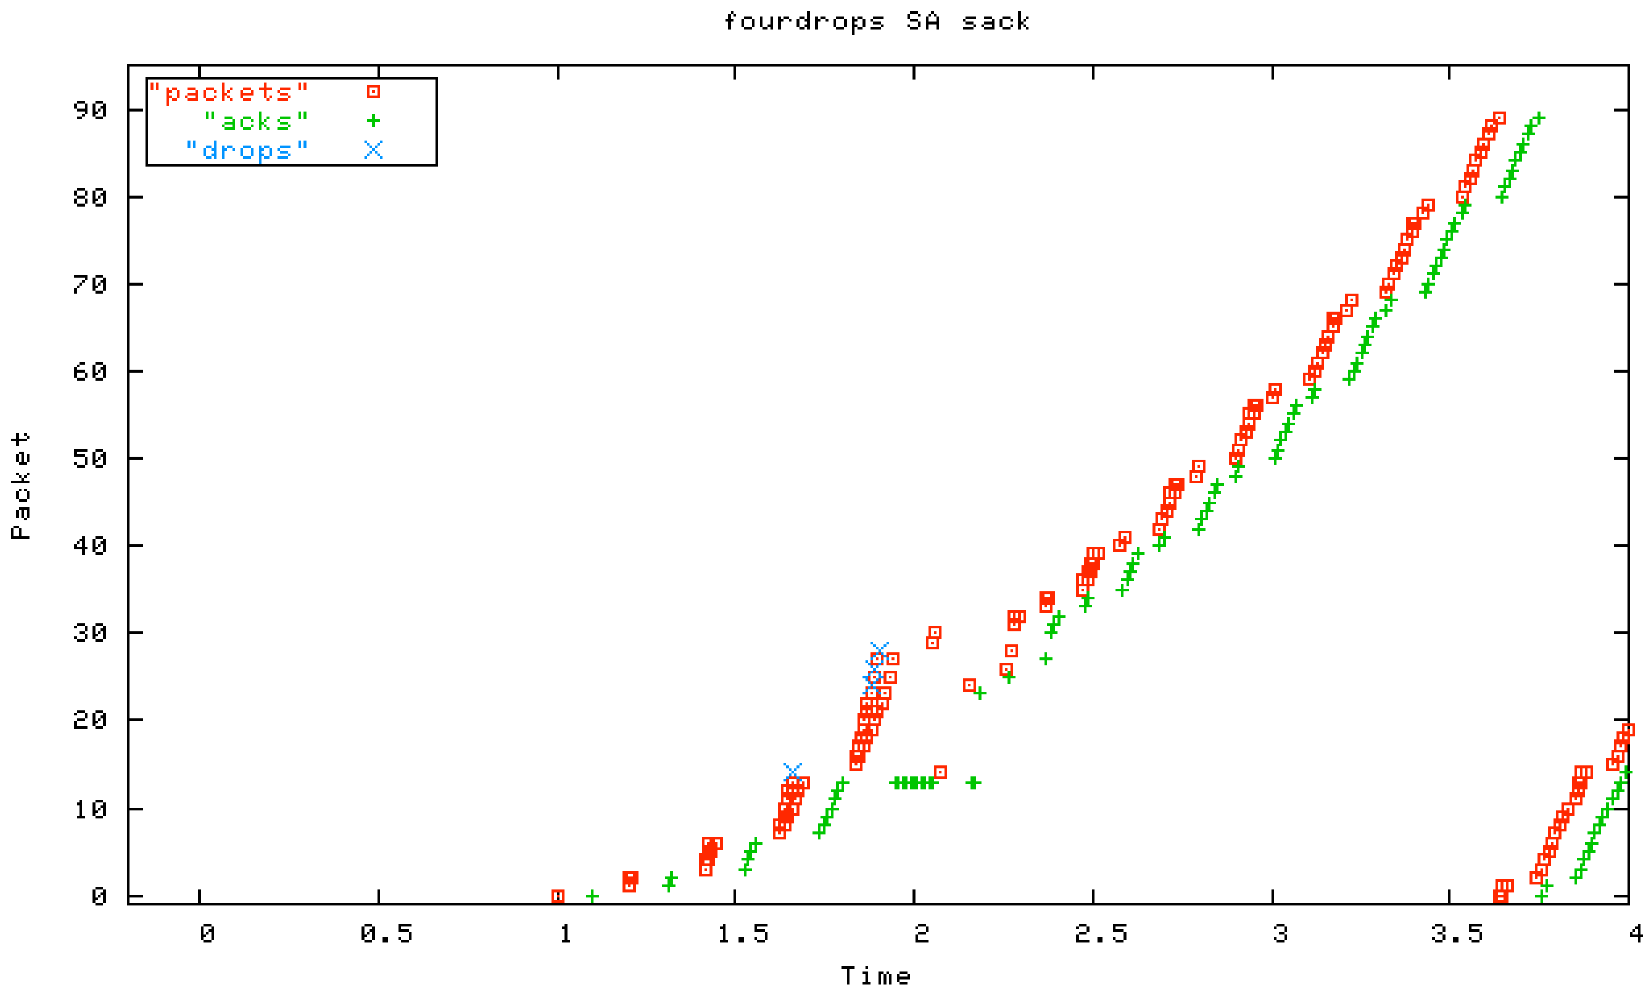

./test-all-tcpVariants fourdrops_SA_sack

Validation Tests in NS-2:

- simple tcp testReno newreno sack tcpOptions tcpReset \
- simple-full full testReno-full testReno-bayfull sack-full \
- tcp-init-win tcpVariants LimTransmit aimd greis rfc793edu rfc2581 rbp \
- sctp tcpHighspeed frto \
- friendly srm realaudio \
- ecn ecn-ack ecn-full quickstart \
- diffusion3 smac smac-multihop \
- manual-routing hier-routing algo-routing lan mcast vc session mixmode \
- red adaptive-red red-pd rio vq rem gk pi cbq schedule rr monitor jobs \
- intserv diffserv webcache mcache webtraf \
- simultaneous mip links plm linkstate mpls oddBehaviors \
- wireless-shadowing wireless-lan-aodv wireless-tdma wireless-gridkeeper \
- wireless-diffusion wireless-lan-newnode satellite WLtutorial \
- source-routing \
- misc tagged-trace message rng xcp wpan \
- energy snoop \
- packmime delaybox \

Validation Tests that Report Statistics:

./test-all-simple.stats

tcp0/time=5/cwnd=7.0000/ssthresh=7/ack=96/rtt=3

…

tcp 0 highest_seqment_acked 336

tcp 0 data_bytes_sent 375000

tcp 0 most_recent_rtt 0.200

…

fid: 0 per-link total_drops 13

fid: 0 per-link total_marks 0

fid: 0 per-link total_packets 360

fid: 0 per-link total_bytes 359040

…

aggregate per-link total_drops 19

aggregate per-link total_marks 0

aggregate per-link total_packets 463

NS-3: what do we need in a simulator?

- A faster simulator, smaller memory footprint.
- Improved emulation capability.
- Moving more easily between simulations and live experiments.
- More wireless models.
- IPv4 and IPv6 support, NATs.
- Integrate other open-source networking code.
- Better maintenance (validation, documentation).
- …

NS-3: the people

- Tom Henderson:
  - the lead PI for the NS-3 project.
- George Riley:
  - will talk about the NS-3 simulator later today.
- Mathieu Lacage:
  - will talk about Yet Another Network Simulator later today.

“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 over 25 years old!

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

- Rosanoff's definition of durable nonsense:
  - R. A. Rosanoff, "A Survey of Modern Nonsense as Applied to Matrix Computations", April 1969.

Question #2: Models

- (2a) What is the goal of the simulations?
- (2b) How does the researcher decide what metrics to use, and what range of simulation scenarios to explore?
- (2c) How does the researcher decide the models to use, in terms of topologies; traffic models; application, transport, and routing protocols; router mechanisms such as queue management; layer-two mechanisms; and the like.

What is the goal of the simulations?

- Evaluating new network protocols?
- Verifying analysis, possibly using a more realistic model?
- Exploring network dynamics?
- Exploring the relationship between parameter A and metric B?
- Helping network operators answer what-if questions?
- …

What models should be used?

- The simpliest model sufficient, but no simplier!
  - A simple topology with one-way traffic of long-lived flows all with the same RTT?
  - A complex topology aiming for full realism of the global Internet?
  - Or something in between?
- E.g., for evaluating TCP fairness as a function of RTT:
  - From "On Traffic Phase Effects in Packet-Switched Gateways", 1992:

Simulation with Two Long-lived Flows:

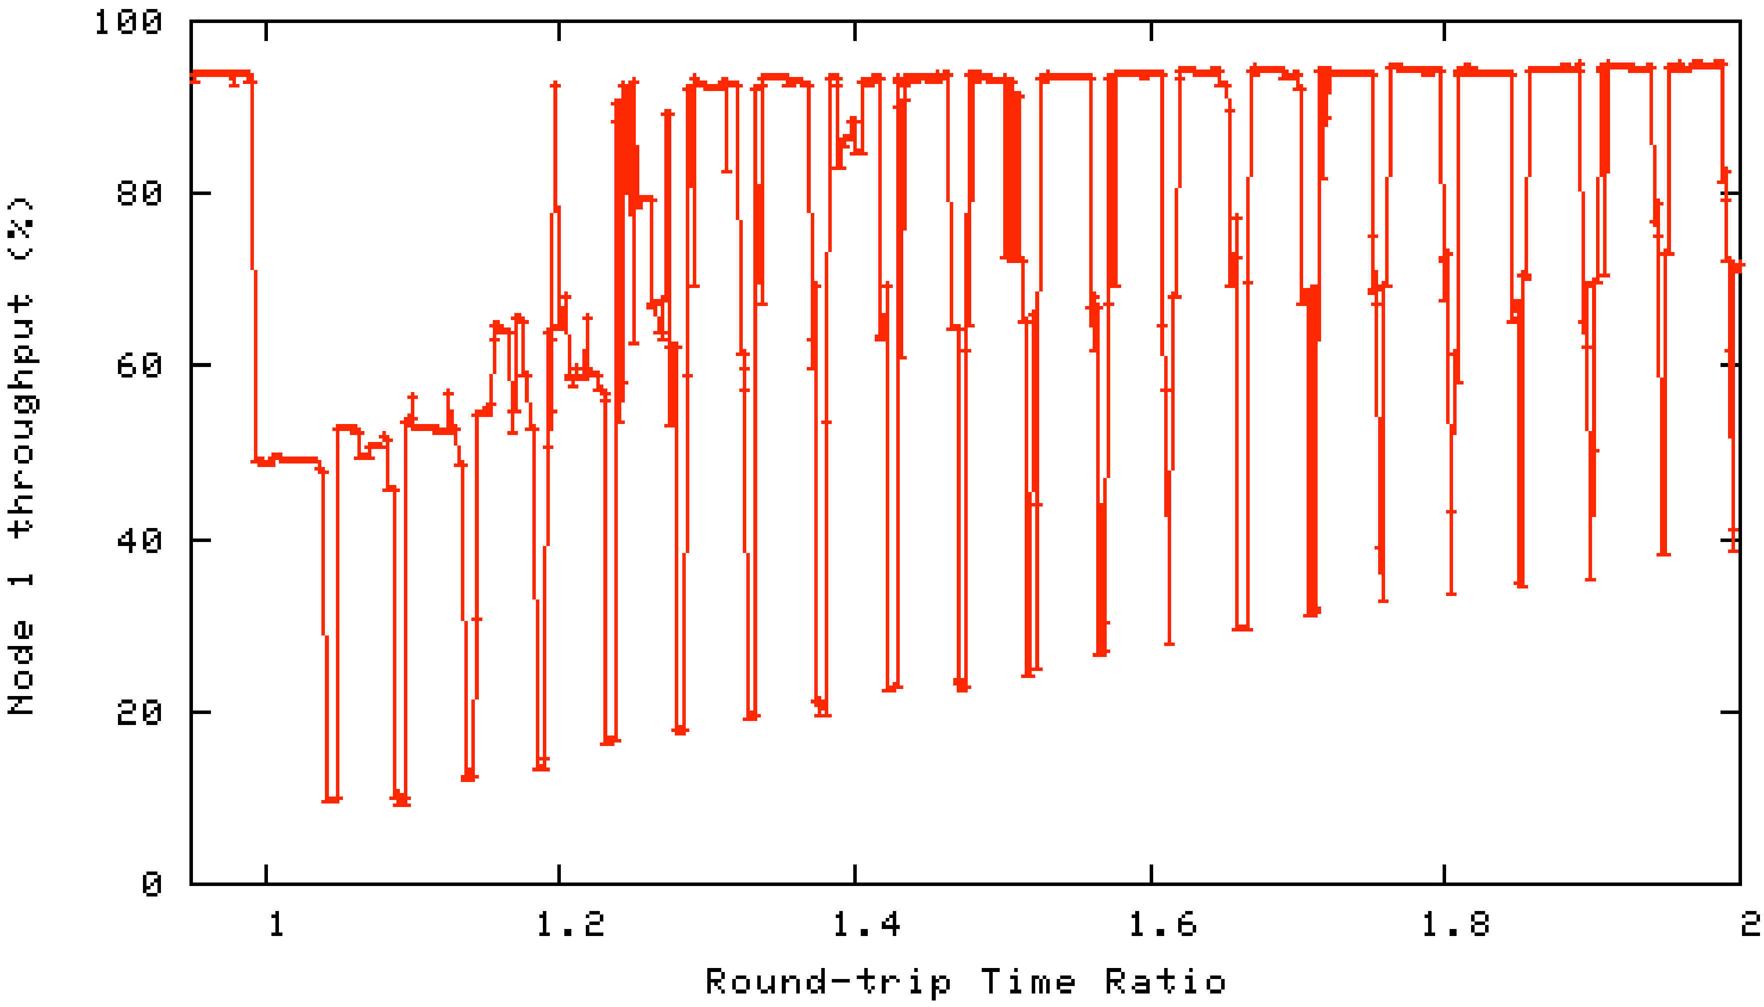

Two Long-lived Flows, with Telnet and Reverse-path Traffic:

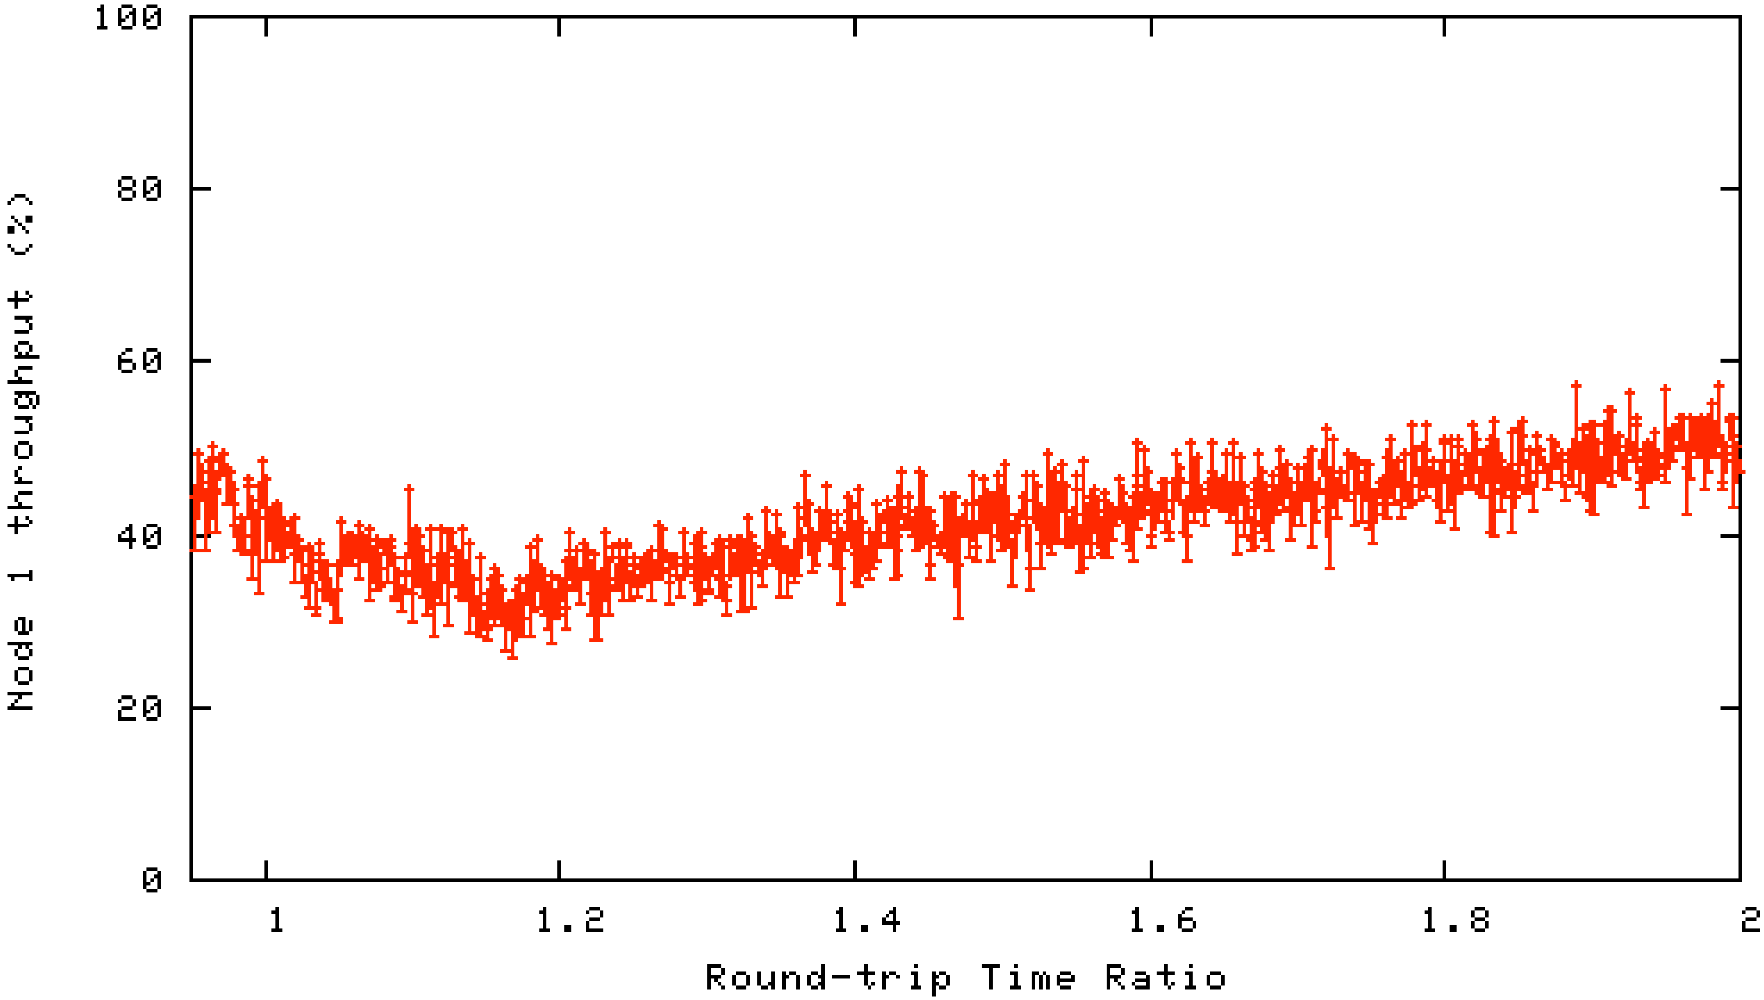

Use a range of scenarios.

- A range of:
  - Topologies?
  - Link bandwidths?
  - Levels of congestion?
  - Levels of statistical multiplexing?
- For evaluating protocols:
  - Look for weaknesses as well as strengths!
    - (As a scientist, not as a used car salesman…)
  - Look for the space of possible tradeoffs.

Take advantage of invariants:

- E.g., heavy-tailed distributions, where they have been verified (e.g., connection sizes, wait times).

Be aware of change:

- In applications, traffic patterns, protocols and router mechanisms, middleboxes, layer-two networks, metrics, etc…

Use results from measurement studies:

- For traffic, topologies, etc.

- From “Difficulties in Simulating the Internet”, 1997, 2001.

Question #3: Further Validation

- (3) To what extent do the simulation results need to be validated by analysis, experimental measurements, or performance results in the Internet?

Question #3: Further Validation

- Validation by experiments or real-world tests for:
  - More realism, in terms of router mechanisms, middleboxes, protocol implementations, link-layer dynamics, and the like.
  - Guarding against unknown effects of the protocol implementation and default parameters in the simulator.
- Simulations, experiments, and real-world tests are all useful for the unexpected behaviors that are discovered.
  - NS-3 should help with this.

But the role of simulations is also important:

- Evaluation of protocols for deployment in N years
  - I.e., not limited by the oddities of today's Internet, routers, transport protocols, etc.
  - In scenarios containing functionality not yet available in testbeds.
  - Evaluating protocols contributed by many different researchers.
- Analysis in a well-understood network model.
- Ease of use by a single researcher.

The Transport Modeling Research Group ("http://www.icir.org/tmrg/")

- Metrics:
  - Metrics for the Evaluation of Congestion Control Mechanisms, internet draft draft-irtf-tmrg-metrics-02.txt, June 2006.
- Tools:
  - Tools for Constructing Scenarios for the Evaluation of Congestion Control Mechanisms, internet draft draft-irtf-tmrg-tools-02.txt, June 2006.
- Next:
  - Best current practice sets of scenarios for simulation and experiments.

Metrics for the Evaluation of Congestion Control Mechanisms

- Throughput, delay, and packet drop rates.
- Response to sudden changes or to transient events; Minimizing oscillations in throughput or in delay.
- Fairness and convergence times.
- Robustness for challenging environments.
- Robustness to failures and to misbehaving users.
- Deployability.
- Security.
- Metrics for specific types of transport.

Throughput, delay, and drop rates:

- Tradeoffs between throughput, delay, and drop rates.
- The space of possibilities depends on:
  - the traffic mix;
  - the range of RTTs;
  - the traffic on the reverse path;
  - the queue management at routers;
  - …

Metrics for evaluating congestion control: response times and minimizing oscillations.

- Response to sudden congestion:
  - from other traffic;
  - from routing or bandwidth changes.
- Concern: slowly-responding congestion control:
  - Tradeoffs between responsiveness, smoothness, and aggressiveness.
- Minimizing oscillations in aggregate delay or throughput:
  - Of particular interest to control theorists.
- Tradeoffs between responsiveness and minimizing oscillations.

Metrics for evaluating congestion control: fairness and convergence

- Fairness between flows using the same protocol:
  - Which fairness metric?
  - Fairness between flows with different RTTs?
  - Fairness between flows with different packet sizes?
- Fairness with TCP
- Convergence times:
  - Of particular concern with high bandwidth flows.

Robustness to failures and misbehavior:

- Within a connection:
  - Receivers that "lie" to senders.
  - Senders that "lie" to routers.
- Between connections:
  - Flows that don't obey congestion control.
- Ease of diagnosing failures.

Metrics for evaluating congestion control: robustness for specific environments

- Robustness to:
  - Corruption-based losses;
  - Variable bandwidth;
  - Packet reordering;
  - Asymmetric routing;
  - Route changes;
  - …
- Metric: energy consumption for mobile nodes
- Metric: goodput over wireless links
- Other metrics?

Metrics for evaluating congestion control: metrics for special classes of transport

- Below best-effort traffic.
- QoS-enabled traffic

Metrics for evaluating congestion control: Deployability

- Is it deployable in the Internet?

Tools for Evaluating Scenarios in Simulations, Experiments, and Analysis: Characterizing Aggregate Traffic on a Link

- Distribution of per-packet round-trip times:
  - Measurements: Jiang and Dovrolis.
- Distribution of per-packet sequence numbers:
  - Measurements:distribution of connection sizes.
- Distribution of packet sizes.
- Ratio between forward and reverse path traffic.
- Distribution of per-packet peak flow rates.
  - Measurements:Sarvotham et al.
- Distribution of transport protocols.
- Typical bandwidth and packet drop rates for congested links.

Tools for Evaluating Scenarios in Simulations, Experiments, and Analysis: Characterizing Paths

- Synchronization Ratio.
  - Determined by queue management (Drop-Tail or RED), level of statistical multiplexing, traffic mix, etc.
- Drop or mark rates as a function of packet size.
  - Determined by queue structure
  - Affects congestion control for small-packet flows.
- Drop rates as a function of burst size.
- Drop rates as a function of sending rate.
  - E.g., determined by the level of statistical multiplexing.

The Effect of Background Traffic on Congestion Control Dynamics:

- A Step toward Realistic Performance Evaluation of High-Speed TCP Variants, S. Ha, Y. Kim, L. Le, I. Rhee and L. Xu, PFLDnet2006.
- The Effect of Reverse Traffic on the Performance of New TCP Congestion Control Algorithms for Gigabit Networks, S. Mascolo and F. Vacirca, PFLDnet2006.
- …
- Observations on the Dynamics of a Congestion Control Algorithm: the Effects of Two-Way Traffic, L. Zhang, S. Shenker, and D. Clark, SIGCOMM 1991.

Distribution of Flow Sizes

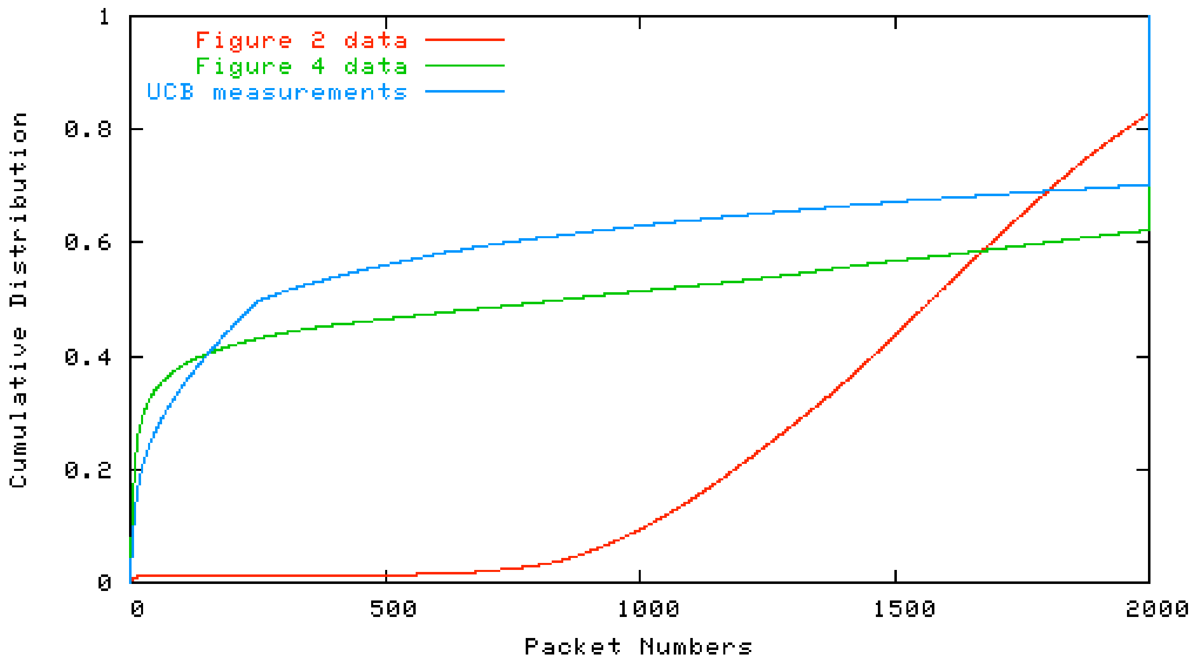

- Distributions of packet numbers on the congested link over the second half of two simulations, with data measured on the Internet for comparison.

Distribution of RTTs

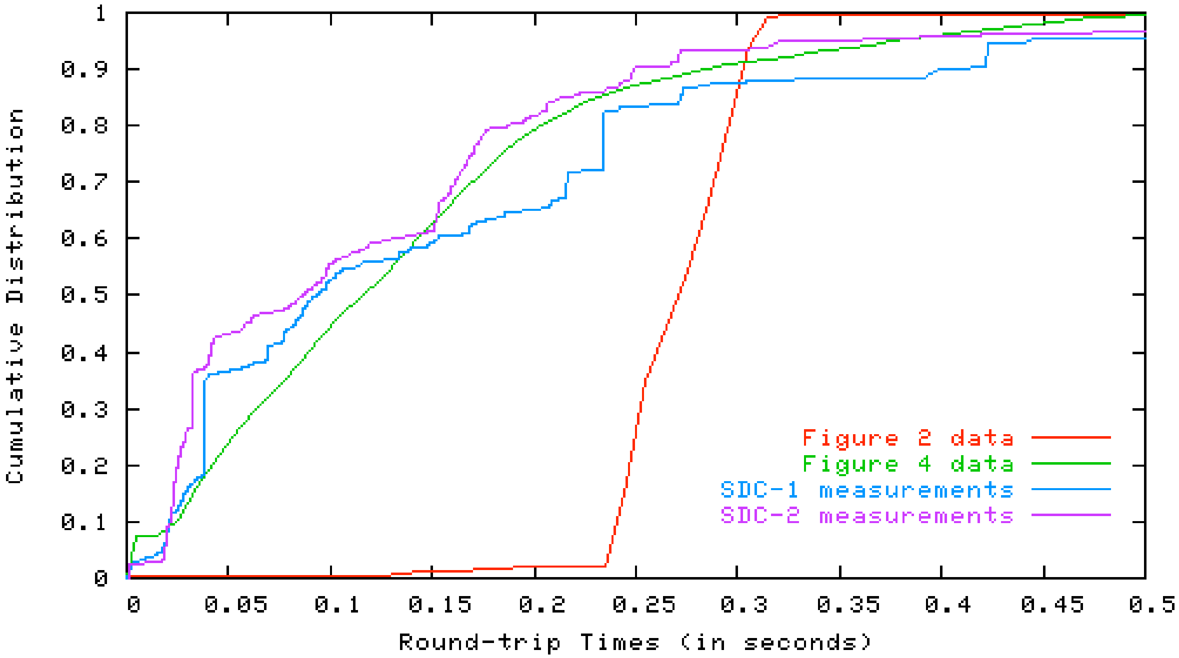

- Distributions of packet round-trip times on the congested link of two simulations, with data measured on the Internet for comparison.

References:

- On Traffic Phase Effects in Packet-Switched Gateways, S. Floyd and V. Jacobson, Internetworking: Research and Experience, 1992.
- "Difficulties in Simulating the Internet" , S. Floyd and V. Paxson, Transactions on Networking, August 2001.
- Internet Research Needs Better Models, Floyd and Kohler, Hotnets 2002.
- TCP Friendly Rate Control (TFRC) for Voice: VoIP Variant, Sally Floyd, internet-draft draft-ietf-dccp-tfrc-voip-02.txt, work in progress, July 2005.

TMRG References:

- TMRG Web Page: http://www.icir.org/tmrg/
- **Metrics for the Evaluation of Congestion Control Mechanisms**. S. Floyd, editor. Internet-draft draft-irtf-tmrg-metrics-02, work in progress, June 2006.
- **Tools for the Evaluation of Simulation and Testbed Scenarios**. S. Floyd and E. Kohler, editors. Internet-draft draft-irtf-tmrg-tools-02, work in progress, June 2006.

Extra Viewgraphs

Impact of Routing Events on End-to-End Internet Path Performance

- “Routing events contribute to end-to-end packet loss significantly.”
- SIGCOMM 06, Wang et al.

Systematic Topology Analysis

- Metrics for measuring graph properties:
  - Average node degree.
  - Degree distribution.
  - Interconnectivity among pairs of nodes with given degrees.
  - Interconnectivity among triples of nodes.with given degrees.
- SIGCOMM 06 paper, Mahadevan et atl.

Summary 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?

Metrics for evaluating congestion control: throughput, delay, and drop rates

- Throughput:
  - Router-based metric: link throughput.
  - User-based metrics:
    - per-connection throughput or file transfer times.
    - Throughput after a sudden change in the app's demand (e.g., for voice and video).
    - Fast startup.
- Delay:
  - Router-based metric: queueing delay
  - User-based metrics: per-packet delay (average or worst-case?)
- Drop rates.

Characterizing the end-to-end path: drop rates as a function of packet size

- Relevant for:
  - evaluating congestion control for VoIP and other small-packet flows.
  - E.g., TFRC for Voice: the VoIP Variant, draft-ietf-dccp-tfrc-voip-02.txt,
- Measurements:
  - compare drop rates for large-packet TCP, small-packet TCP, and small-packet UDP on the same path.
- There is a wide diversity in the real world:
  - Drop-Tail queues in packets, bytes, and in between.
  - RED in byte mode (Linux) and in packet mode (Cisco).
  - Routers with per-flow scheduling:
    - with units in Bps or in packets per second?