Internet Research Needs Better Models

%

Sally Floyd and Eddie Kohler

October 28, 2002

Hot Topics in Networks

What do we mean by model?

%

- The world for a given simulation, experiment, or analysis
  - For networking, this might include topology, traffic mix, end-node protocols, queue drop policies, congestion, . . .
- Application-specific or research-area-specific
  - Each research topic needs its own set of models
  - Not “one model to rule them all, one model to bind them”
- Model should be as simple as possible, but no simpler
  - Invalid models give invalid results, as we will see . . .

Example: Models for research on router queue behavior

%

- Relevant fields: AQM, scheduling, Differentiated Services, . . .
- Important model features include:
  - Characteristics of congested links
  - Range of round-trip times
  - Traffic characterization (distribution of transfer sizes, . . .)
  - Reverse-path traffic
  - Effects of congestion elsewhere
- Models for other research topics will have different features

A fairly common model for router queue behavior research

%

- Characteristics:
  - Dumbbell topology with one congested link
  - One-way traffic
  - Simple traffic mix (either all long-lived or all Web traffic)
  - Limited range of round-trip times
  - Single congested link bandwidth
- How well does this match with reality?
  - Does any divergence from reality matter, or not?
  - Are we wasting research effort on bad models? . . .

Example: Bandwidth sharing

%

- Shows how tricky simulations can be [FJ 1992]
- How do TCP flows share a congested link?
- Model characteristics for a possible simulation experiment:
  - Dumbbell topology with one congested link
  - One-way traffic
  - All long-lived traffic, all packets the same size
  - Drop-tail queue management
  - Different round-trip times
  - Deterministic simulation (no randomness)

Result: Phase effects

%

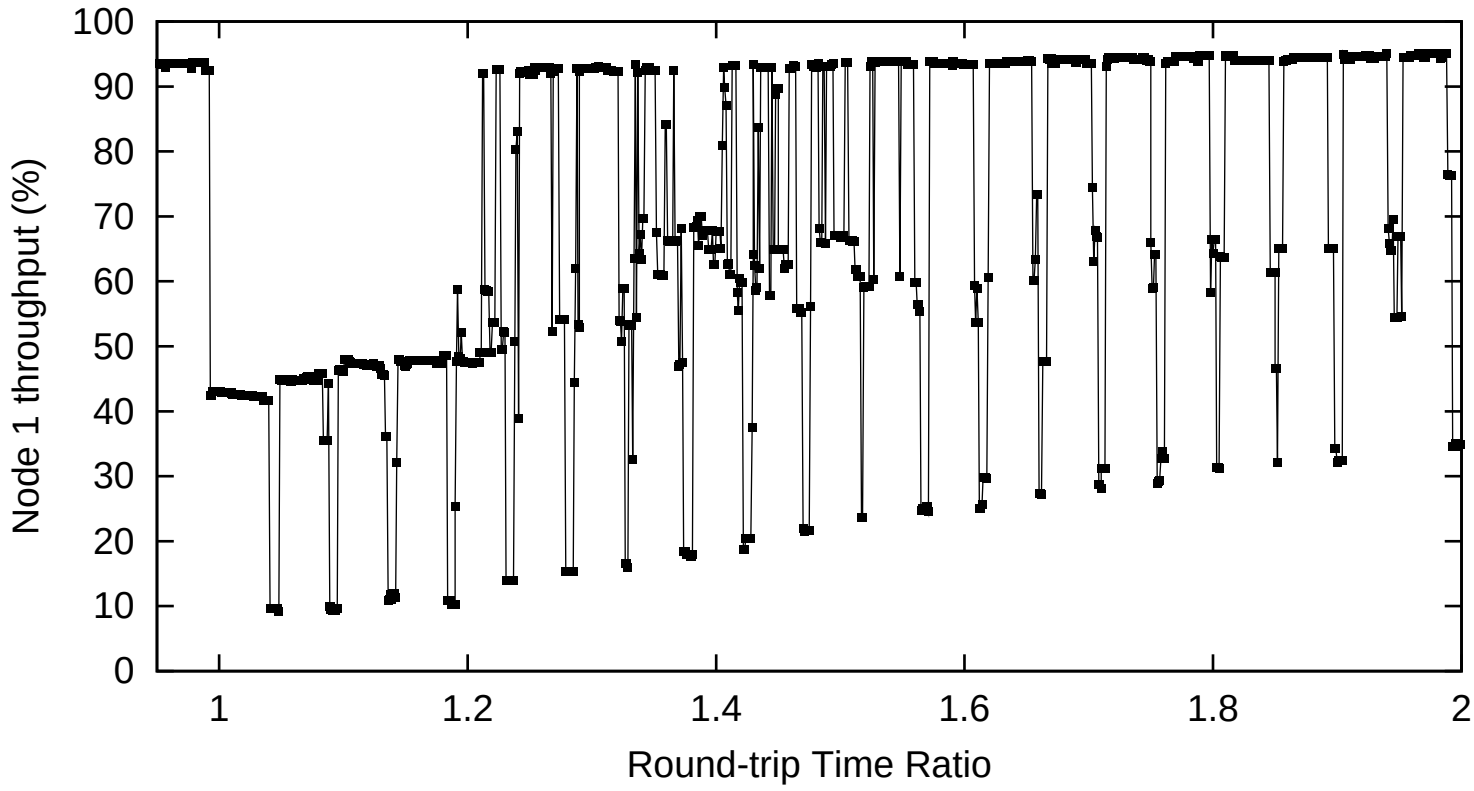

Flow 1's throughput as a function of the ratio of the two flows' round-trip times

What went wrong?

%

- Bad model! In particular:
  - Deterministic synchronization → *not realistic* (fix with randomness)
  - All long-lived traffic → *not realistic* (fix by adding Web traffic)
  - One-way traffic → *not realistic*, leads to all packets in queue having same size (fix by adding reverse-path flows)
- The rich behavior of phase effects is not relevant to the modern Internet.
- It doesn't lend any useful insight on the chaotic or fractal behavior of actual networks.
- This is not interesting to explore further!

Example: Router queue behavior with TCP/AQM

%

- Does TCP/AQM introduce undesired behavior at router queues, such as oscillations?
- Model characteristics for a possible simulation experiment:
  - Dumbbell topology with one congested link
  - One-way traffic
  - All long-lived traffic
  - RED active queue management
  - Same round-trip times
- This model is temptingly easy to analyze with control theory!

Result: Severe oscillations

%

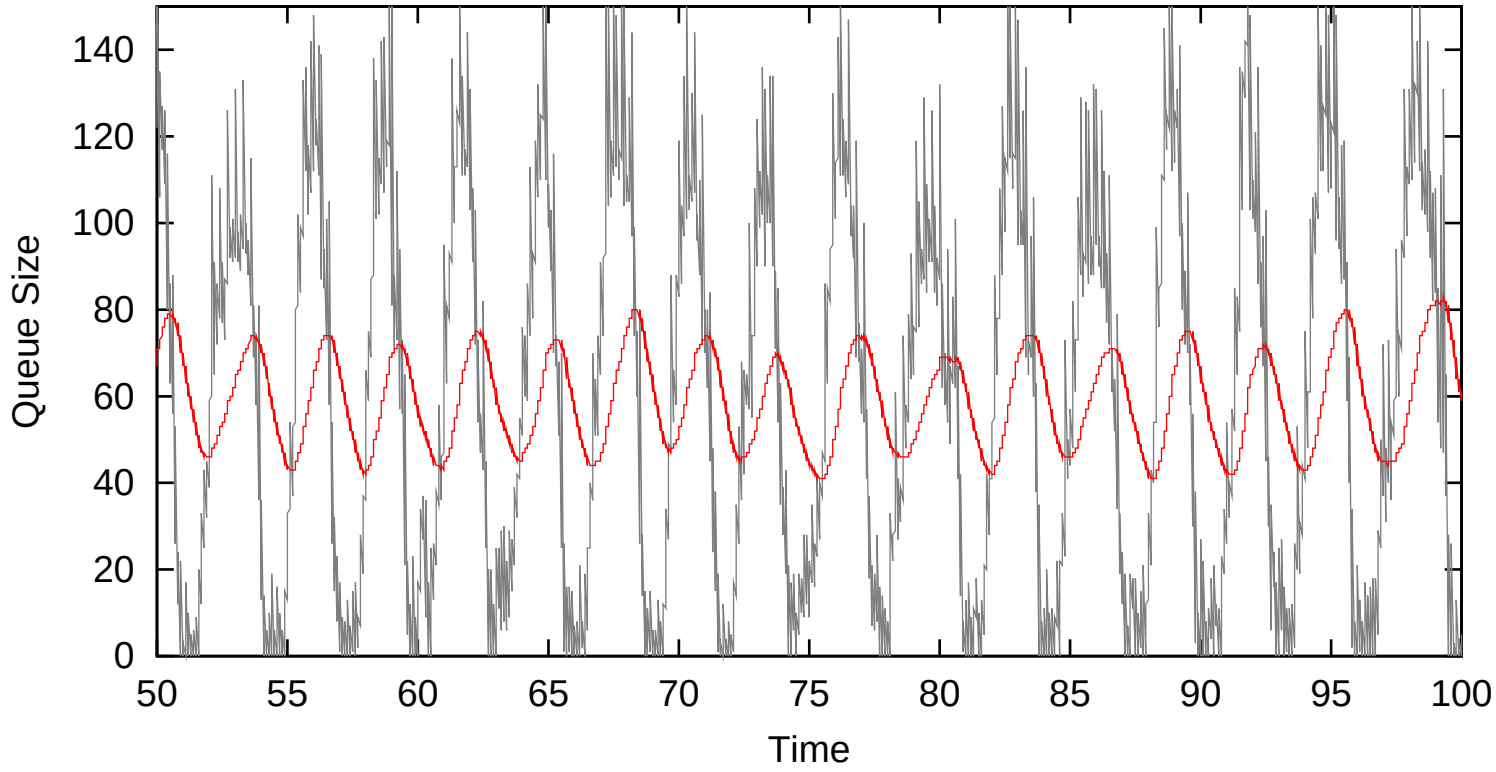

Long-lived traffic, 240 ms RTT

What happens if we vary the model?

%

- Wide range of round-trip times
- More short-lived traffic

Result: Less severe oscillations

%

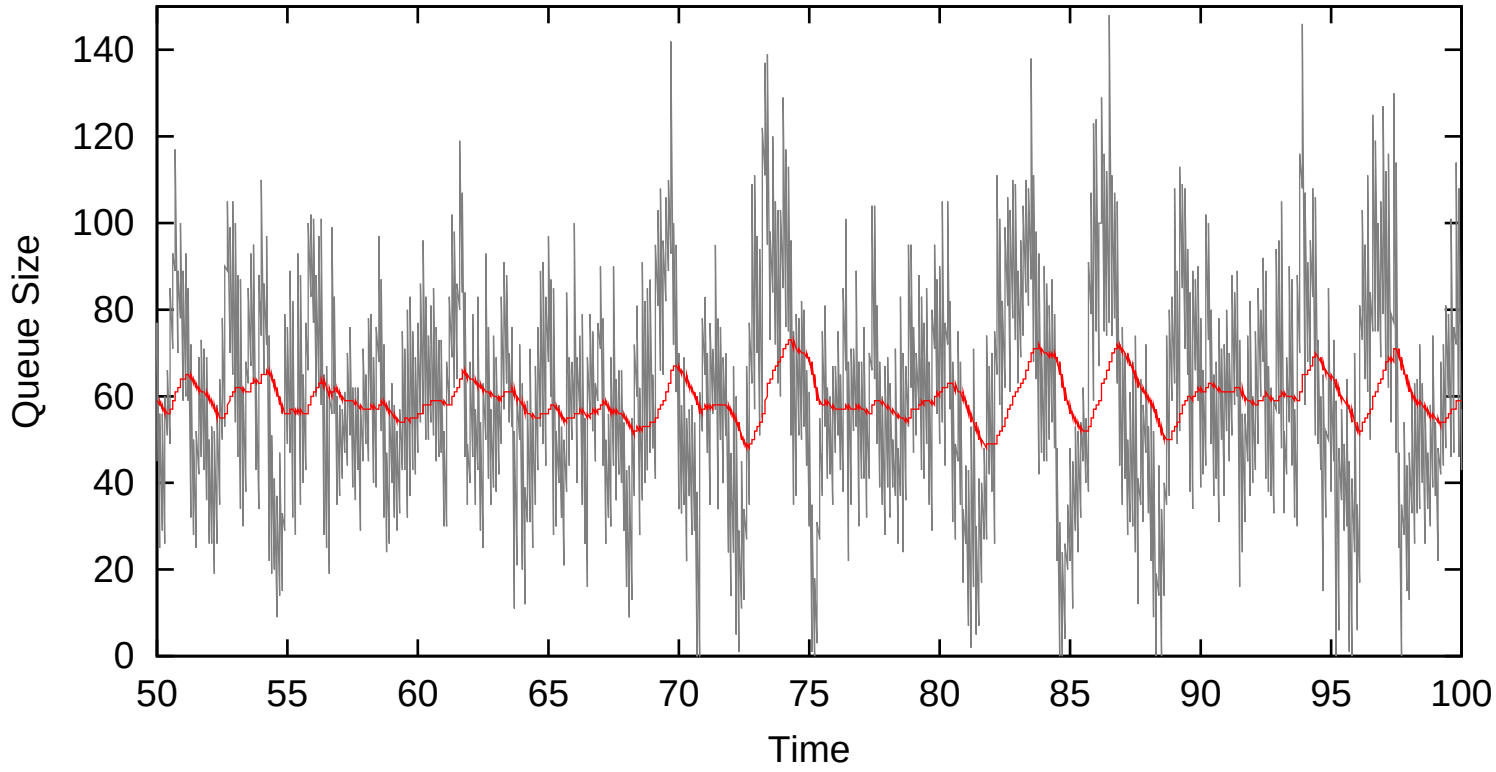

Long-lived traffic, 20–460 ms RTTs

Result: Less severe oscillations

%

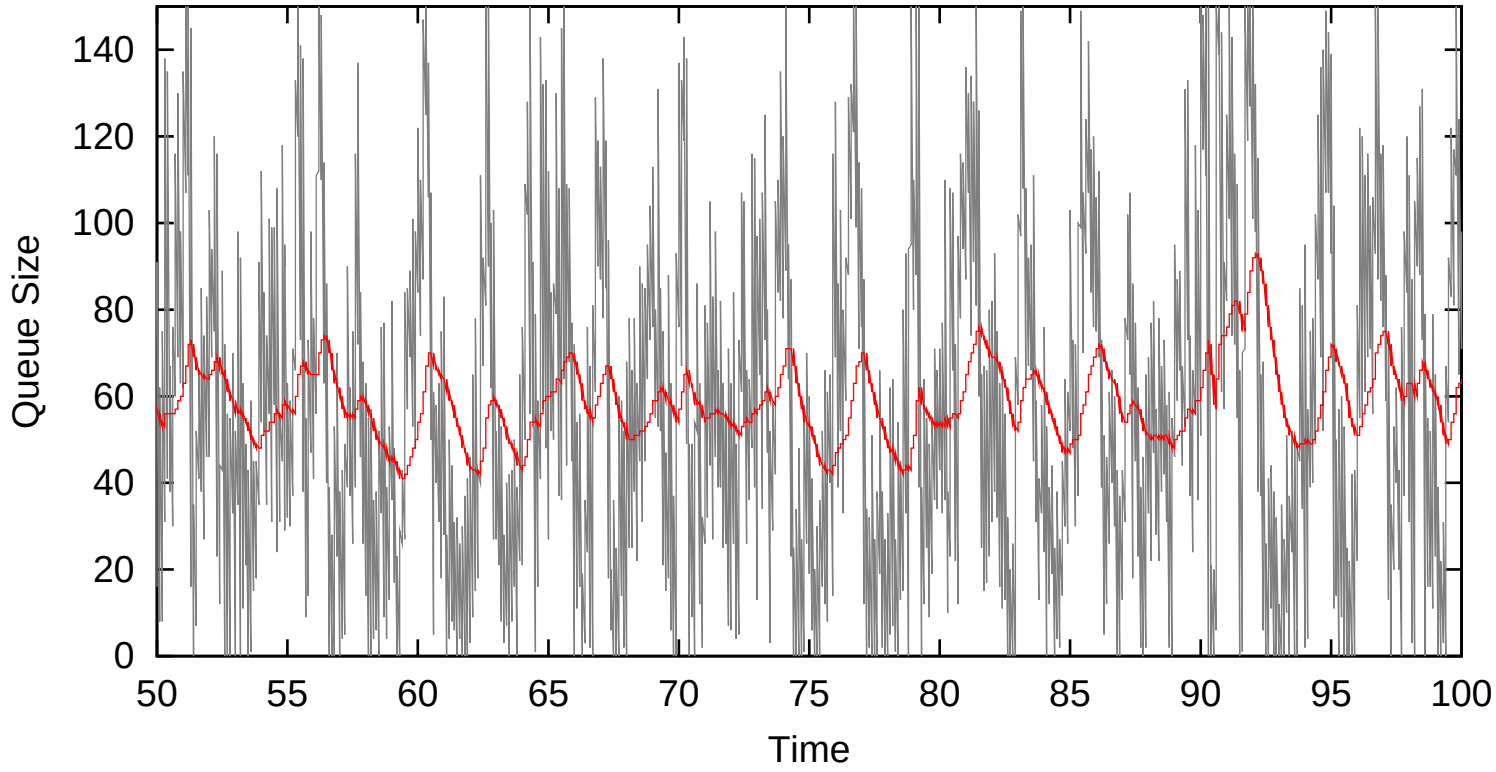

Mostly Web traffic, 20–460 ms RTTs

What is a realistic model?

%

- We actually don’t know!
  - We don’t know typical levels of congestion, typical bandwidths for congested links, . . .
- We can evaluate certain aspects of current models against measurement results, however:
  - Range of round-trip times
  - Flow sizes

Measurements: Range of round-trip times

%

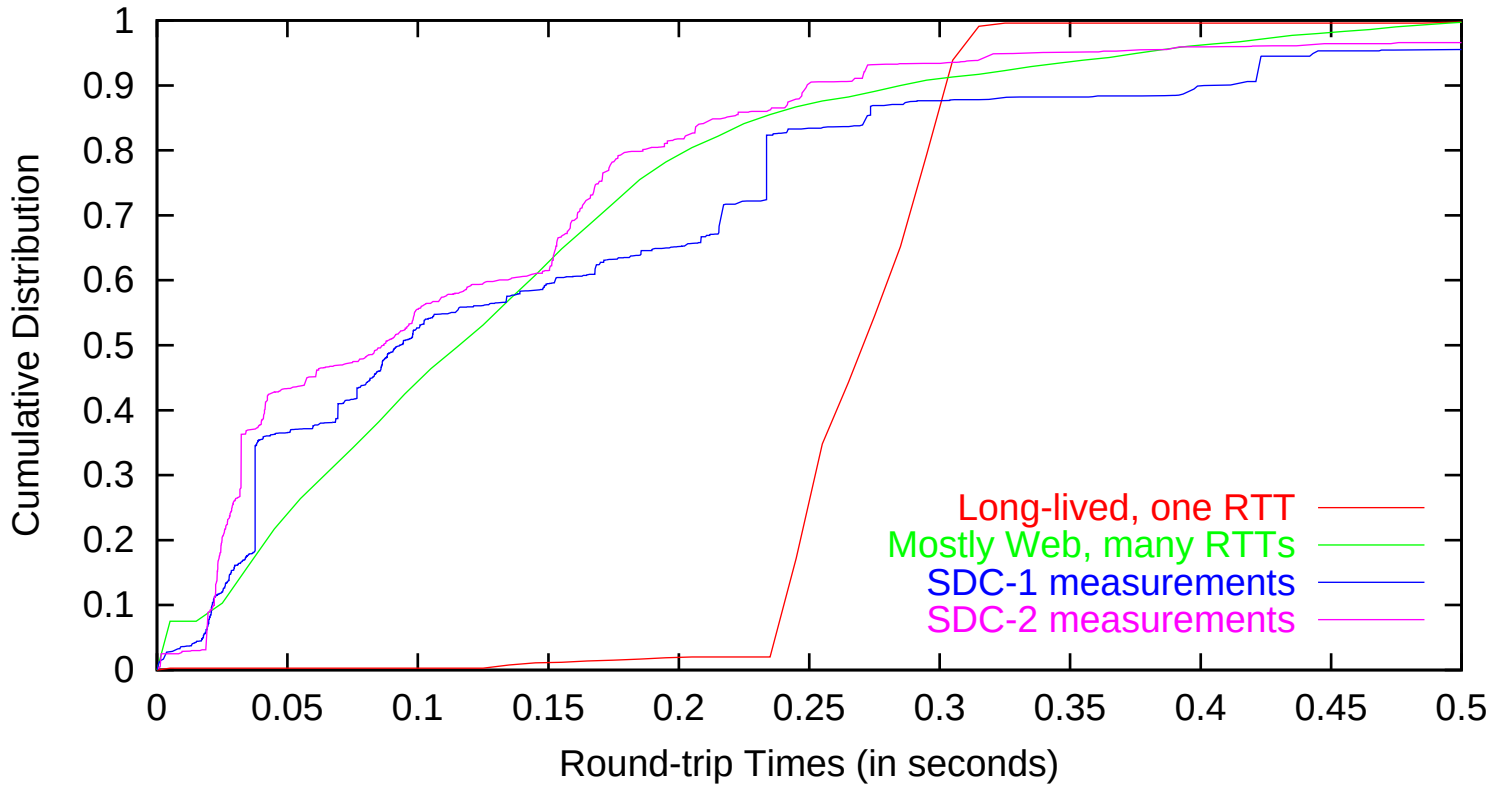

Distributions of packet round-trip times on the congested link of two simulations, with data measured on the Internet for comparison.

Measurements: Flow sizes (packet numbers)

%

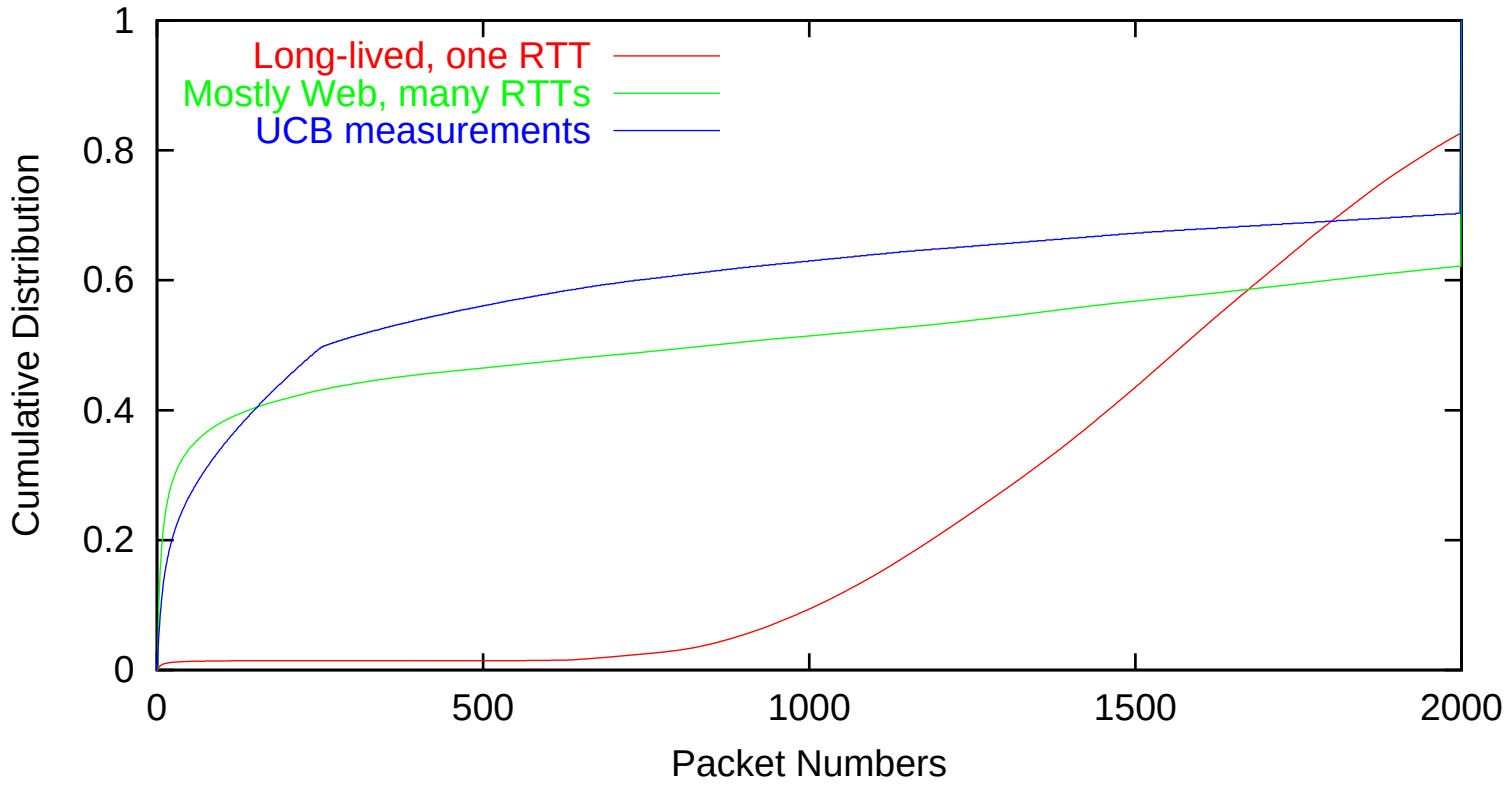

Distributions of packet numbers on the congested link over the second half of two simulations, with data measured on the Internet for comparison.

Conjecture: You can prove anything with simulations!

%

- To show oscillations with RED or Adaptive RED:
  - Use long round-trip times and mostly long-lived flows
  - (As we've shown, this is not realistic.*)
- To show poor performance with the Proportional-Integral Controller (PI) or Random Early Marking (REM):
  - Use mostly web traffic, or changes in the level of congestion over time
- To show poor performance with drop-tail or Adaptive Virtual Queues (AVQ):
  - Use packet drop rate as the key metric
- It would be nice for research to be grounded in something more solid . . .

More AQM modeling issues

%

- RED
  - Implicit model: low packet drop rates, narrow range of bandwidths for the congested link
  - Reality: high packet drop rates also occur, congested links cover wide range of bandwidths
  - Results: Poor performance at high packet drop rates, insufficient practical guidance on parameter settings

Modeling issues in transport protocols

%

- New research area, new model characteristics, same problems!
- TCP Reno
  - Implicit model: multiple drops in a single window mean more congestion
  - Reality: multiple drops commonly due to a single congestion event
  - Result: TCP Reno can behave poorly in practice
- TCP Vegas
  - Model: connection's sending rate affects its queue size
  - Implies low levels of statistical multiplexing
  - Reality: congestion also occurs at links with high statmux
  - Result: TCP Vegas can behave poorly in practice

We Need Better Models!

%

- The models that we use in simulations and experiments are not confirmed by experimental studies.
- The lack of tools for constructing and critically evaluating the models that we use is holding back the field.

Goals

%

- Models should be specific to the research questions being investigated.
- We need to understand how models' parameter settings affect experimental results.
- Modeling must go hand-in-hand with measurement.
- We want models that apply to the Internet of the future, as well as to the Internet of today.

Moving forward

%

- Start with specific research questions.
  - We (the network research community) need to understand which model parameters are critical to performance.
- For critical parameters, we need:
  - New measurements
  - A careful use of existing measurement studies to understand realistic settings . . .
  - . . . or new measurement tools
  - Not the same as measurement: need digested results
- Next, we need to begin a set of best practices for model construction.
  - A shared repository of models and simulation scenarios
  - Papers evaluating models

Related work

%

- There's lots!
  `http://www.icir.org/models/bettermodels.html`
- Model evaluation: [JRFGW 2001], [GV 2002], [LKJK 1999], ...
- Topology generation: Georgia Tech, Rocketfuel, ...
- Traffic characteristics: [JD 2002], ...
- Organizations: CAIDA, NLANR, IRMG, ...