

Perspectives on Transport Benchmarking

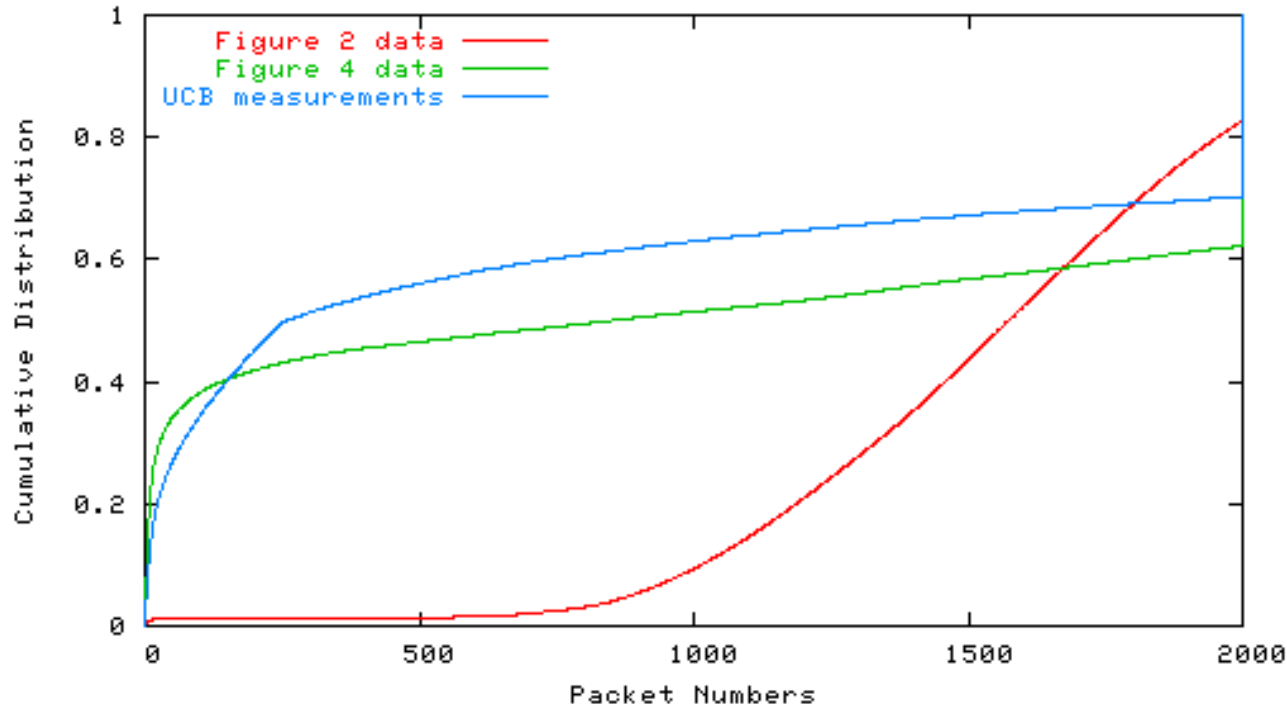
Panel for PFLDnet 2006

Nara, Japan

February 3rd, 2006

So, what is the problem?

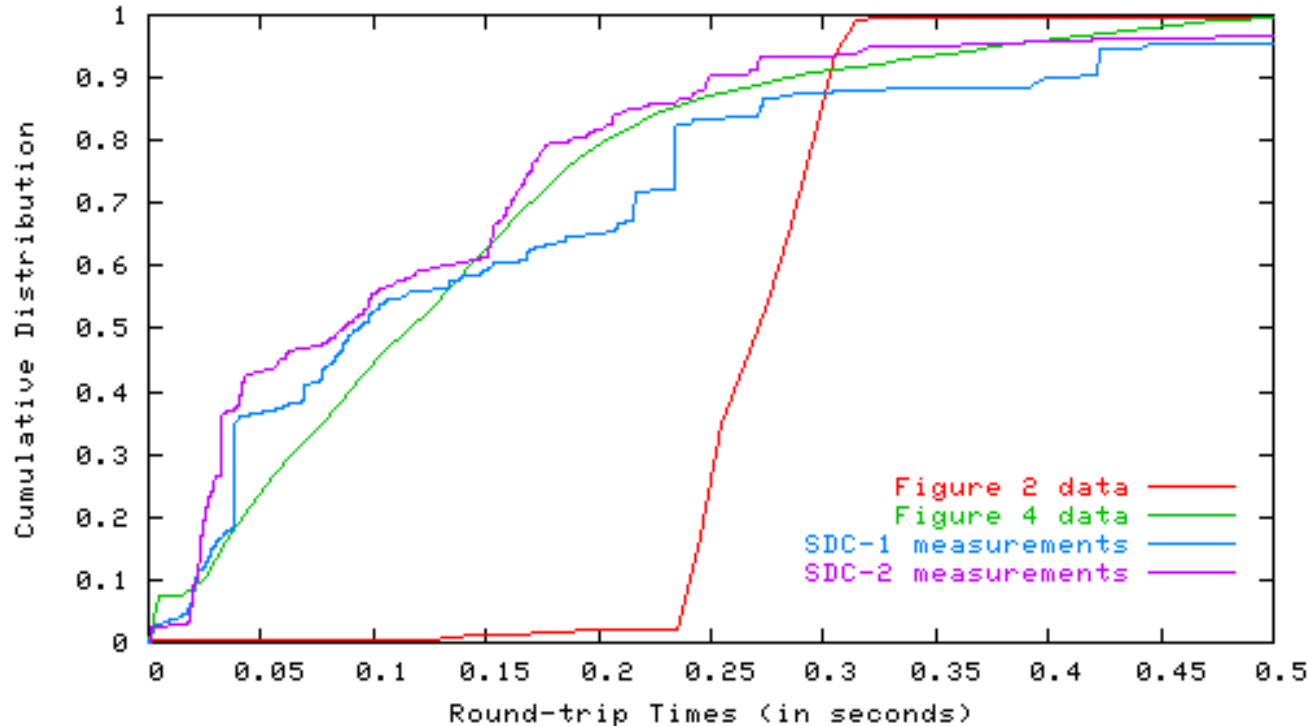
Distribution of Flow Sizes



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

[Floyd and Kohler, 2002]

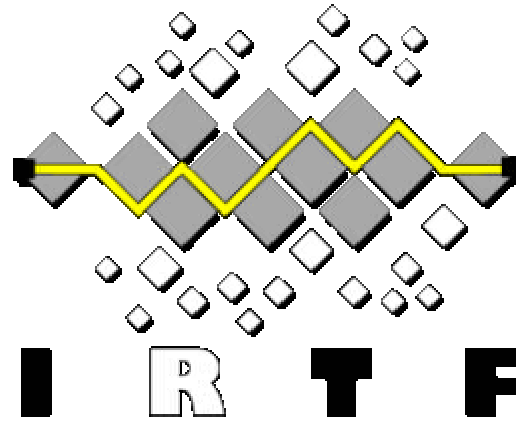
Distribution of RTTs



- **Distributions of packet round-trip times** on the congested link of two simulations, with data measured on the Internet for comparison.
[Floyd and Kohler, 2002]

Panel

- Goal: *to get multiple perspectives on transport benchmarking*
- Composition
 - Richard Hughes-Jones, University of Manchester
 - Lachlan Andrew, University of Melbourne
 - Injong Rhee, NC State University
 - Doug Leith, Hamilton Institute
 - Aaron Falk, USC Information Sciences Institute
- Agenda
 - 15 minutes (firm) for each speaker
 - Questions at end of panel



Transport Activities in the *Internet Research Task Force*

Aaron Falk
IRTF Chair
USC Information Sciences Institute

Internet Research Task Force

- Promotes research relating to the Internet using small, focused research groups (RGs)
- RGs are very autonomous
 - Most have open membership, various mtg schedules
- Oversight provided by the IRTF Chair and Internet Architecture Board (IAB)
- RG output
 - Standardization in IETF
 - Informational or experimental RFC publication
 - Academic papers (by participants)
- <http://www.irtf.org>

Internet Congestion Control Research Group

- Chartered late 2005, chaired by Mark Handley and S. Keshav
 - Just kicking off (1st messages today)
- *Goal: to move towards consensus on which technologies are viable long-term solutions for the Internet congestion control architecture, and what an appropriate cost/benefit tradeoff is.*
 - *Scope includes LFNs, global Internet, realtime apps*
- Open membership, 2-3 meetings per year (planned)

Strawman ICCRG Strategy

- ⇒ 1. Challenge document
- Current TCP limitations, long term requirements
2. Solution Proposals
- *Your solution here*
3. Discussion
- Themes, best-practices, consolidation, missing solutions
4. Consensus Formation
- Make recommendations

Transport Modeling Research Group

- Chartered in Fall 2005, chaired by Sally Floyd
- *Goal: improve methodologies for evaluating congestion control mechanisms*
 - Motivated by the lack of good models underpinning analysis, simulation, and testbed experiments
- Open membership, unlikely to meet f2f
- Two drafts published:
 - *draft-irtf-tmrg-metrics-01.txt*
 - *draft-irtf-tmrg-tools-01.txt*

“Metrics for the Evaluation of Congestion Control Mechanisms”

- Goal: enable a broad view of strengths and weaknesses of newly proposed mechanisms
- Scope:
 - Define a consistent set of metrics
 - Consider relationships between metrics
 - E.g., delay vs. throughput
 - Describe network-, flow-, and user-based interpretations
 - Note application and transport goals
 - For some metrics, e.g., fairness, there is no agreement on goals

Metric Overview

- Throughput, Delay, & Loss Rate
- Response Times & Minimizing Oscillations
- Fairness & Convergence
- Robustness for Challenging Environments
- Robustness to Failures and Misbehavior
- Deployability
- Metrics for Specific Types of Transport
- User-based Metrics

This is a richer set of metrics than have been presented this week.

“Tools for the Evaluation of Simulation and Testbed Scenarios”

- Goal:
 - Enable construction of better scenarios, based on better underlying models.
 - Needed for analysis, simulation, and testbed experiments
- E.g., measurement tools allowing characterization of link aggregates & end-to-end paths
- If successful will allow comparison of key characteristics found in models vs. real world

Traffic Characteristics

- Link aggregate characteristics
 - Per-packet RTT distribution
 - Packet sequence number distribution
 - Packet size distribution
 - Forward vs. reverse traffic
 - Peak flow rates
 - Transport protocols
- End-to-end path characteristics
 - Synchronization ratio
 - Drop rates as a function of packet size.
 - Drop rates as a function of burst size.
 - Drop rates as a function of sending rate.
 - Degree of packet drops.
 - Range of queuing delay.

This is a richer set of scenarios than have been presented this week.

Conclusions

- We're not ready to pick a winner
- Indeed, we still need firm definitions for some metrics and scenarios
- The TMRG is the place to do the work on baselining metrics and benchmarking
- The ICCRG is the place to do the work on comparing protocols

Thank you