# PROJECT SUMMARY: Pragmatic Internet Measurement: Empirical Results and Tools for Modeling and Simulation

Networking researchers have unspoken understandings—mental models—of the Internet's important properties. The scenarios used in simulations and experiments reveal aspects of these mental models, including one or more of the following: Flows live for a long time and transfer a lot of data. Simple topologies, like the infamous "dumbbell" with one congested link, are sufficient to study any traffic property. All flows on that congested link share a small range of round-trip times. All traffic across the link is one-way; reverse-path traffic is never congested.

All of these properties can deeply affect simulation and experimental results, and therefore our evaluations of research. But none of them are justified by measurement studies. Some are actively wrong. Lack of good measurements and measurement-based intuitions, and lack of tools for evaluating measurement results and applying their results to current simulations, are holding back the field.

We propose a program of pragmatic measurement studies that will discover what the Internet is really like. Our particular focus is properties that researchers use to evaluate potential network changes via simulation and experiments. We will build on existing measurement results, generate new empirical results, create tools and repeatable processes that facilitate ongoing measurement, and boil our results down to concise intuitions and simulation-ready models. We will further illustrate how each aspect of the models can influence the results of simulations and experiements.

Packet traces will lie at the core of our techniques. [??] Fundamental Internet properties at least potentially measurable from traces [??] that remain ill understood in practice, or that are poorly reflected in our current modes, include:

- Where is congestion on the Internet? How often does congestion occur at routers, if at all? Why does congestion occur?
- What range of round-trip times might one expect to see on a single link?
- How many congested links might a flow experience? Do congested flows often experience reverse-path congestion? Or network-induced burstiness? How prevalent is ack compression?
- How do link properties, such as expected congestion, range of round-trip times, and so forth, differ on: home access links, server access links, backbone links, peering points, transoceanic links, wireless links, satellite links, links outside of North America?
- How much do these link properties affect end-to-end traffic dynamics? In particular, how important are transoceanic links to overall Internet behavior?
- Should our simulation topologies take all of this into account? If so, then how?
- How are the answers to these questions evolving over time?

Part C

# PROJECT DESCRIPTION

## 1 Introduction

[Fragments of text that might go somewhere:]

This work is motivated by the difficulties we have experienced in our own research in determining the relevant aspects of simulation scenarios that need to be explored. We need both better models, and better tools for evaluating our own and others' models. This work is motivated by a continuing frustration with poor quality of models in general in Internet research, and by the need for a critical approach in evaluating the models that we use in simulations and experiments.

Restricting ourselves arbitrarily to papers that have appeared in recent years in SIGCOMM, we can observe the following models being used: Dumbbell topologies (cite our own work); All long-lived flows, all with the same round-trip times; One-way traffic; Reno TCP.

In this proposal we intend to start from what modelers would or should want to know, for specific research questions, and to see what is there.

| Research Problems | Typical Models | Supporting Measurements |
|---|---|---|
| AQM, scheduling,<br>differentiated services. | A dumbbell topology, with<br>aggregate traffic. | Characteristics of congested links,<br>Range of round-trip times,<br>Traffic characterization<br>(distribution of transfer sizes, etc.)<br>Reverse-path traffic,<br>Effect of congestion elsewhere. |
| Unicast congestion control. | A single path,<br>with competing traffic. | Characteristics of links,<br>Queue management along path,<br>Packet-reordering behavior,<br>Packet corruption on a link. |
| Multicast congestion control. | A single multicast group<br>in a large topology | Router-level topologies,<br>Loss patterns,<br>Traffic generation by group members. |
| Routing protocols. | A large topology | Router-level topologies,<br>AS-level topologies,<br>Loss patterns |
| Caching infrastructures and CDNs. | ? | ? |
| Peer-to-peer systems. | ? | Topologies,<br>Routing of P2P systems,<br>P2P traffic patterns. |
| Controlling DDoS attacks. | Models of large topologies<br>with aggregate traffic | Topologies,<br>Attack patterns. |
| Web cache performance | A single cache with<br>many clients and servers,<br>as in Web Polygraph. | Detailed client behavior,<br>server behavior. |

We note that we aren't attempting to construct a single model of the global Internet, with a single set of simulation scenarios, for use by all researchers. Each kind of research needs its own kinds of models, and its own kind of supporting measurements for building and evaluating those models. Thus, a model and related set of simulation scenarios is needed for research and evaluation of congestion-related mechanisms such as scheduling or queue management at a queue; currently researchers often use a dumbbell topology, with edge links that give a range of round-trip times, and with a range of forward and reverse-path traffic. A completely different model and set of simulation scenarios is used to evaluate multicast congestion control mechanisms, one stressing the role of topology and of both unicast and multicast routing. For developing and evaluating routing protocols, models are used that take into account the AS-level structure as well as the router-level structure of the topology. For evaluating the performance of unicast transport protocols such as TCP, models and simulation scenarios are needed that can control such aspects of the end-to-end path as reordering behavior, variable delay, asymmetric bandwidth, etc.

One measurement/modeling issue is that of constructing and evaluating models for congestion-related mechanisms at a queue in a router: the research could be about AQM; ECN; diff-serv and QoS; aggregate-based congestion control; fairness; etc. For that, there are a lot of things about the aggregate traffic at a link that we don't know. Assuming that we are generally interested in links that might at some point have congestion, we don't even know which links we are interested in. Access links? Transoceanic links? Links at peering points? Then we really need to know not only the obvious things like ranges of RTTs, but also things like how the aggregate traffic at that link is affected by burstiness or bottlenecks from elsewhere in the network.

Another measurement/modeling issues are those that come from characterizing the address space. Other possibilities for measurement/modeling issues include modeling issues needed for research on large-scale traffic patterns, and modeling issues needed for research on middleboxes.

# 2 Problems with Unrepresentative Models

We are working on the problem of the models used in research because we believe that it makes a big difference.

For example, for exploring oscillations with AQM, it can make a big difference whether the assumed environment is one of a simple topology with one-way traffic composed of long-lived flows all with the same round-trip times, or whether the environment is assumed to be more diverse and complex.

These simulations: Packet drop rate from 2 to 4 percent. Throughput 86 - 95 percent. NO ECN. Buffer was Congested link was 120 ms. Change to 60 ms.

On a somewhat more subtle level,

# 3 Measurement properties for queuing models

Areas of research: queue management, scheduling, differentiated services, aggregate-based congestion control.

What we need: models to use in simulations and experiments to construct realistic aggregate link traffic. This concerns not only the long-range dependence or lack of long-range dependence in the aggregate traffic, which is an area that is well-studied.

Factors such as the range of round-trip times, the distribution of flow sizes, the effects of congestion elsewhere in the network (on both the forward and reverse paths) can all significantly affect performance.

## 3.1 Where are the congested links?

For some of this research, all that is relevant is models of those links that sometimes experience congestion. As a first question, we don't even know which links those are, and what are the patterns of congestion. That

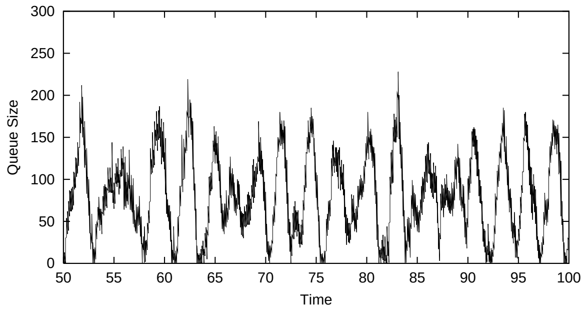


**Figure 3**: Long-lived and web traffic.

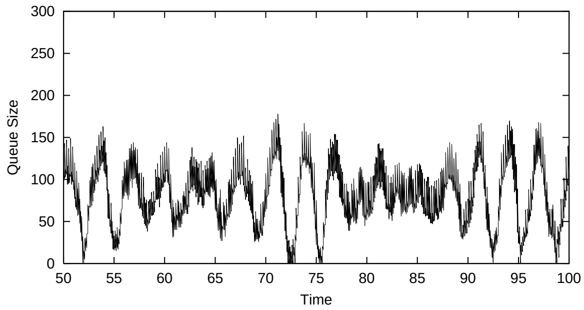


**Figure 1**: One-way, long-lived, same RTT.

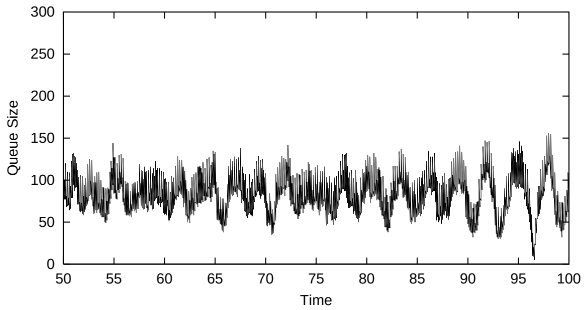


**Figure 4**: ECN.

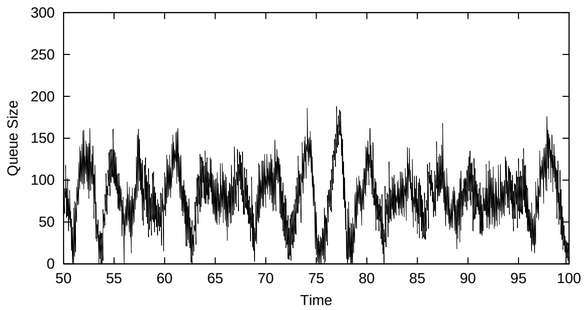


**Figure 2**: A wide range of RTTs.

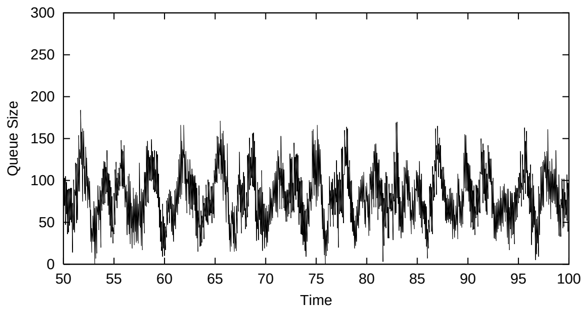


**Figure 5**: Diverse RTTs and traffic mix, ECN.

is, where does the congestion occur? Edge networks? Peering points? Transoceanic links? Links outside of North America? What are the congested links like, in terms of bandwidth, delay, level of statistical multiplexing, etc.

A possible approach would be to first develop some understanding of where congestion is in the network. We might want separate models of different kinds of congested links: access links, campus links, transoceanic links, links to public or private peering points, etc.

Barford and Plonka in [2] have begun work on characterizing network anomolies such as network operation anomolies, flash crowd anomolies, and network abuse anomolies from flow-based measurements. [2] also surveys the existing literature in detecting DDoS attacks.

## 3.2 The range of round-trip times

The range of round-trip times is clearly a critical modeling issue. Measurements on the range of round-trip times on a link goes back to the 1996 Olympic Web server [1, 5], and techniques for estimating the round-trip times from packet traces are discussed by Jiang and Dovrolis [3]. However, many more measurements

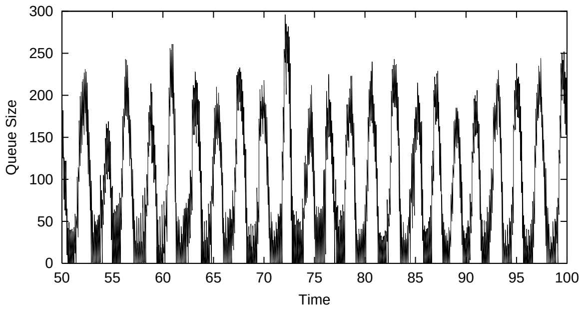


**Figure 6**: Simple scenario with reverse-path congestion.

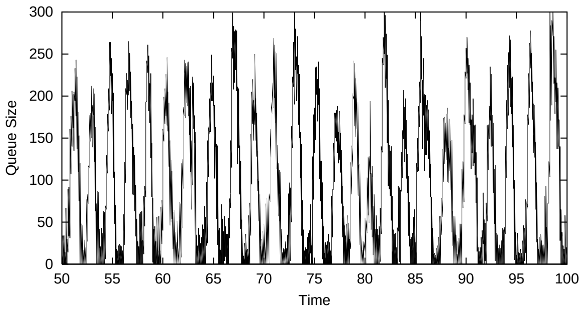


**Figure 7**: Complex scenario with reverse-path congestion.

are needed.

One possible metric for evaluating the distribution of per-packet round-trip times on the link would be by looking that the median round-trip time and the lower and upper quartiles. That is, half of the packets on the link will have round-trip times between the lower and upper quartiles. The ratio between the upper and the lower quartile gives a possible metric for the range of round-trip times on that link. We would conjecture that there could be a useful relationship between some aspects of performance of AQM (e.g., stability concerns) and this metric.

A second task related for round-trip-time distributions would be generating models to use in simulations and experiments that resulted in a suitable range of round-trip times for packets on the link in question.

### 3.3 The per-packet distribution of flow sizes

Similar metrics (median and upper and lower quartile) have been used for characterizing the distribution of per-packet flow sizes [4, 9]. Again, more measurement is needed in this area, and more evaluation of the importance of this distribution in the performance of AQM.

### 3.4 Small-scale burstiness

The delay-throughput tradeoffs are in part determined by the small-scale burstiness of traffic at the link, and this small-scale burstiness is in turn determined by such factors as the queueing and drop rates experienced by the data and acknowledgement packets elsewhere along their paths. The drop rates for drop-tail queues are also in part determined by this small-scale burstiness. The related work in this area would be found by asking Vern, I think.

### 3.5 Levels of congestion

The "level of congestion" can be quantified by the packet drop rates at the queue. The Internet Traffic Report, the Internet Weather Report, the Internet End-to-end Performance Monitoring Group, and other sites [4] have long had measurements of the packet loss rates of pings to various routers in the Internet. Most of this measurement is necessarily about end-to-end path properties, rather than direct router measurements.

### 3.6 Well-travelled territory: self-similarity

There is of course a large body of work on characterizing the self-similarity of aggregate traffic on a link, but we are not really intending to tread in that well-traveled space.

## 4 Related work

[8] discusses some of the difficulties encountered in simulating the Internet, in terms of addressing the heterogeneity, scale, and rate of constant change in the Internet. For example, there is heterogeneity in topology and link characteristics, in transport and routing protocols, in traffic characteristics, and in levels of congestion. [8] also posits some invariants in modeling the Internet that seem relatively stable; these include a diurnal pattern of activity, self-similarity in aggregate traffic, Poisson session arrivals, a log-normal distribution of connection sizes, heavy-tailed distributions, a Pareto distribution for Telnet packet generation, and some topological invariants of the global topology determined by the clustering of populations into cities, with the minimum delay between cities determined by the speed of light.

The web page on measurements about congestion [4].

Areas that we will not address in our work include the following well traveled areas: modeling for wireless networks; characterizing path properties [7]; examining the long-range dependence, or lack of long-range dependence, in aggregate network traffic [6]; evaluating the degree distribution of network topologies.

We would probably also avoid the MPLS/traffic engineering ranges of research issues, though they do have modeling issues that would draw on previous work on characterizing address structures.

**Part D**

# REFERENCES CITED


[1] H. Balakrishnan, V.N. Padmanabhan, S. Seshan, S. Stemm, and R.H. Katz. TCP Behavior of a Busy Web Server: Analysis and Improvements. Mar. 1998. http://citeseer.nj.nec.com/balakrishnan98tcp.html.

[2] Paul Barford and David Plonka. Characteristics of Network Traffic Flow Anomalies. *ACM SIGCOMM Internet Measurement Workshop 2001*, 2001. http://www.icir.org/vern/imw-2001/proceedings.html.

[3] Hao Jiang and Constantinos Dovrolis. Passive Estimation of TCP Round-Trip Times. http://www.cis.udel.edu/~dovrolis/Papers/rtt.ps.gz.

[4] Web Page. Measurement Studies of End-to-End Congestion Control in the Internet, . http://www.icir.org/floyd/ccmeasure.html.

[5] Web Page. Questions about Round-trip Times, . http://www.icir.org/floyd/rtt-questions.html.

[6] Web Page. Self-similarity and Long Range Dependence in Networks, . http://www.cs.bu.edu/pub/barford/ss_lrd.html.

[7] Vern Paxson. Measurements and Analysis of End-to-End Internet Dynamics. Ph.D. dissertation, 1997, http://www.icir.org/vern/papers.html.

[8] Vern Paxson and Sally Floyd. Difficulties in simulating the Internet. 2001.

[9] Anees Shaikh, Jennifer Rexford, and Kang G. Shin. Load-Sensitive Routing of Long-Lived IP Flows. 1999. http://citeseer.nj.nec.com/shaikh99loadsensitive.html.