

Observations on  
Network Research  
and the  
Evolution of the Internet Infrastructure

Sally Floyd  
SIGCOMM  
August 2007

URL “<http://www.icir.org/floyd/talks.html>”

# Goals of the Network Research Community

- The evolution of the Internet infrastructure is not the only goal for the network research community.
- Other goals:
  - Science.
  - Education.
  - Clean-slate architectures.
  - Research on applications, layer two, etc.
  - Research on other infrastructures.
  - ...
- But the evolution of the Internet infrastructure is the goal that I have been interested in.

## This interest in the evolution of the Internet infrastructure is widely shared in SIGCOMM

- SIGCOMM 2007 **workshops**:
  - Mobility in the Evolving Internet Architecture,
  - IPv6 and the Future of the Internet,
  - Large Scale Attack Defense
- SIGCOMM 2007 **papers**:
  - Routing, naming and addressing, security, network management, Internet topologies, transport protocols...
- Other SIGCOMM **conferences and publications**:
  - Internet Measurement Conference,
  - CCR,
- ...

Of course, understanding or affecting the evolution of the Internet infrastructure is not easy:

- **Chicken-and-egg** deployment problems.
- **Tussles** in cyberspace.
- Problems with **middleboxes**:
  - Security vs. evolution.
  - Routers/middleboxes that drop packets.
- Problems of **incremental deployment**:
  - Old routers, middleboxes, and end hosts.
- **Unexpected interactions** between layers.
- The model of a single, coherent Internet architecture doesn't match the real world.

## And yes, the Internet architecture and the Internet infrastructure have serious problems:

- Possible failure scenarios:
  - The Internet taken over by **attackers**?
  - A malicious or accidental **routing break-down**?
  - Running out of **IPv4 addresses**?
  - **Congestion collapse**?
  - Widespread **spam**? (oops, already happened...)
- Researchers **can't get their favorite mechanisms ubiquitously deployed**?
  - **It is good for the barrier for deployment to be high.**
  - It is not good if the barrier is unsurmountable, however.

This doesn't mean that one abandons research on the evolution of the Internet infrastructure!

- The underlying architecture still has life in it.
- Important work to be continued:
  - Research on naming and addressing, routing, transport, ...
  - Research on interactions between layers:
    - Apps, transport, routing, layer two.
  - Research on how things play out.
    - Learning from the real world.
  - Many other things.

# When Research Meets the Real World

- Research meeting the real world:
  - IP Multicast.
  - Intserv, diffserv.
  - ECN.
- How well do we as a community understand the problems?
- We can learn from what has taken off, and why, as well as from what hasn't taken off.
  - RTP, SDP, RSVP-TE, but not for IP multicast.

# The rest of the talk:

## Some observations from my own work.

- Some research is not about making changes, but about contributing to our collective understanding:
  - Wide-area traffic: the failure of Poisson modeling
  - Difficulties in simulating the Internet.
  - The role of end-to-end congestion control.
  - Measurements on transport protocol deployment, etc.
  - Models for evaluating congestion control mechanisms.
- Some research is about changes that end up being deployed:
  - Early Packet Discard for ATM networks.
  - TCP: NewReno, SACK, higher initial windows, Limited Retransmit.
  - Active Queue Management.



# The rest of the talk:

## Some observations from my own work.

- Some research is about changes that \*might\* end up being deployed, one day:
  - ECN?
  - Quick-Start?
  - TFRC?
- Some research is about mechanisms that didn't end up being deployed:
  - Reliable multicast?
  - Controlling High-Bandwidth Flows at the Congested Router?
  - Pushback and aggregate congestion control?
  - All of this research can still be useful!

# ECN (Explicit Congestion Notification): An example of cross-layer communication.

- **ECN**
  - Transport to routers: ECN-Capable packet.
  - Routers to transport: Congestion Experienced.
- **Deployment problems:**
  - Chicken-and-egg deployment problems.
  - Middleboxes that drop TCP SYN packets attempting to negotiate ECN.
  - Routers that crash when an ECN-Capable data packet arrives.
  - Other router bugs in responding to ECN.
- **ECN is not a recent idea!**
  - TCP and Explicit Congestion Notification, CCR, 1994.
  - RFC 2861, “A Proposal to add ECN to IP”, 1999.
  - RFC 3168, K.K. Ramakrishnan, Sally Floyd, and David Black, 2001.

# Quick-Start:

## An example of cross-layer communication

- **Quick-Start with TCP**, setting the initial window:
  - Sender's desired sending rate in IP option in SYN packet.
  - Routers decrement a TTL counter, if OK.
  - The receiver sends feedback to the sender in SYN/ACK.
  - The sender knows if all routers on the path approved.
- **Potential problems:**
  - some tunnels are not compatible;
  - some middleboxes drop packets with IP options;
  - chicken-and-egg deployment problems.
- **Deployment:** maybe in a controlled environment.
- RFC 4782, A. Jain, S. Floyd, M. Allman, and P. Sarolahti

# TMRG: Models used in Simulations, Experiments, and Analysis for Evaluating Transport Protocols

- “Internet Research Needs Better Models”, w/Eddie Kohler, 2002.
- Models for evaluating transport protocols:
  - Metrics.
  - Tools for evaluating models:
    - Characterizing traffic:
      - Connection sizes, packet sizes, RTTs, ...
    - Characterizing end-to-end paths:
      - Congestion, layer-two interactions, routing changes, ...
  - Test suites for simulators and testbeds.
- Sally’s pet peeve:
  - We are scientists, not used-car salesmen (or saleswomen).

From: Eddie Kohler

Subject: a poem on incrementality

## Diligence Is to Magic as Progress Is to Flight

[by Marianne Moore]

With an elephant to ride upon--"with rings on her fingers and bells on her toes",  
she shall outdistance calamity anywhere she goes.  
Speed is not in her mind inseparable from carpets. Locomotion arose  
in the shape of an elephant; she clambered up and chose  
laboriously. So far as magic carpets are concerned, she knows  
that although the semblance of speed may attach to scarecrows  
of aesthetic procedure, the substance of it is embodied in such of those  
tough-grained animals as have outstripped man's whim to suppose  
them ephemera, and have earned that fruit of their ability to endure blows  
which dubs them prosaic necessities--not curios.

# Acknowledgements

- Almost all of the work mentioned in this talk involved **collaborations with other researchers**. I will not attempt to list them all here.
- I thank people for **feedback** on these slides.

## A partial list of collaborators for some of the work mentioned in this talk:

- Mark Allman
- Steve Bellovin
- Kevin Fall
- Mark Handley
- Tom Henderson
- Van Jacobson
- Amit Jain
- Eddie Kohler
- Ratul Mahajan
- Steve McCanne
- Jitendra Padhye
- Vern Paxson
- K.K. Ramakrishnan
- Allyn Romanow
- Pasi Sarolahti
- Scott Shenker
- Lixia Zhang
- ...

Extra slides:



# Open issues: Cross-layer Communication

- Many possibilities, many drawbacks.
- Possibilities:
  - Lower layers to transport: E.g., link changes.
  - Transport to lower layers: E.g., preferences for packet handling.
  - Two-way communication: E.g., ECN, Quick-Start, XCP.
  - Transport protocols and performance-enhancing middleboxes.
- Some of the drawbacks:
  - Interactions with routing;
  - Unexpected router and middlebox behaviors;
  - Tunnels; ...
- Guideline: robustness over efficiency!
- Transport-layer Considerations for Explicit Cross-layer Indications, draft-sarolahti-tsvwg-crosslayer-01.txt, (a survey document), P. Sarolahti, S. Floyd, and M. Kojo.

# Problems:

## Disconnects between Areas

- Disconnect between routing and transport?
- Disconnect between economics and Internet architecture?
- Disconnect between IETF and research community?

# Open Issues: The IRTF (Internet Research Task Force)

- EME (End Middle End Research Group)
  - Explicitly dealing with NATs and middleboxes.
- ICCRG (Internet Congestion Control RG)
  - Open Research Issues in Internet Congestion Control
- RRG (Routing Research Group)
  - Focus: designing an alternate architecture to meet challenges of scalability, mobility, multi-homing, and inter-domain traffic engineering.
- TMRG (Transport Modeling Research Group)
- ...

## Open issues:

# The Routing and Addressing Problem (ROAP)

- Problems of scaling:
  - Increasing the degree of interconnection.
  - Increasing the information load.
- Problems of mobility and multihoming.
- The problem of identifier/locator separation.
- Architectural goal:
  - “A network should be able to implement reasonable internetworking choices without unduly impacting another network's operation.”

## References:

- TRMG:
  - URL: <http://www.icir.org/tmrg/>
  - Internet Research Needs Better Models:
    - URL: <http://www.icir.org/models/>

# References:

- Cross-layer communications:
  - Interactions of Link-Layer and Transport Protocols, "<http://www.icir.org/link-triggers.html>".
  - RFC 4907, “Architectural Implications of Link Indications”, B. Aboba, 2007.

## References:

- ROAD:
  - Geoff Huston, More ROAP: Routing and Addressing at IETF68, The Internet Protocol Journal, V. 10, No. 2.