Network Dynamics and Scalable Congestion Control

Sally Floyd

November 19-20, 2001

Workshop on UltraLarge Networks

Questions of the workshop (from the web page):

- Can current networking approaches evolve to deal with large increases in scale, or are different, revolutionary paradigms required?

- Do we need new approaches for building models?

- Do we need new approaches for building simulation environments?

Topics:

- Evolution vs. the blank slate approach.

- Difficulties in modeling and simulating the Internet.

- Network dynamics, scalable congestion control, and other network issues relating to scale.

- An aside about layers of abstraction in NS.

Themes:

- The Internet is a work in progress, with no central control or authority, many players independently making changes, and many forces of change (e.g., new technologies, new applications, new commercial forces, etc.)

- So far, the success of the Internet has rested on the IP architecture's robustness, flexibility, and ability to scale, and not on its efficiency, optimization, or fine-grained control.

- While there is much to be learned from the "blank slate" approach to network design, there are also grave dangers. There is extraordinary power in the rough and discontinuous evolution of the current Internet infrastructure.

Topics:

- Evolution vs. the blank slate approach.
-
-
-

The benefits of a blank slate approach:

- Understanding mechanisms for accompishing a specific function.
- Thinking beyond the constraints of the current architecture.
- Addressing potential needs of the future.

(These things are done, in a different way, in the evolutionary approach as well.)

The benefits of an evolutionary and empirical approach:

- Remaining in touch with the heterogeneity of the Internet:
of applications, protocols, topologies, technologies, and more.

- Attention to the fully complexity of interactions:
among different mechanisms, technologies, and levels of the protocol stack.

- Attention to emerging forces:
new stresses, failure modes, sustaining technologies, and disruptive forces.

- Attention to the interplay between clean models and messy realities.

The interplay between clean models and messy reality:

- Simulation and analysis both require models.
  - For the evolutionary or the blank-slate approaches.

- Models are needed for investigating fundamental underlying dynamics
  - Without the bugs of real-world implementations.
  - Without the limited functionality of the real-world.
  - Without the unknown interactions of the real-world.

- Messy reality is needed also, however...
  - As a check on the inherent limitations of models.
  - The Internet is more heterogeneous that most things that we model.
  - The Internet is more fast-changing that most things that we model.

Topics:

-
- Difficulties in modeling and simulating the Internet.
-
-

Change and heterogeneity as conditions of the Internet:

- New link-level technologies: e.g., optical, wireless.

- Very high bandwidth in one place, and very low bandwidth in another.
  – Cheaper bandwidth leads to higher connectivity between ASes (Autonomous Systems).

- Changes in routers:
  – QoS, queue management, ECN (Explicit Congestion Notification).

- Changes to end-to-end congestion control mechanisms:
  – in TCP, and in new transport protocols.

- Changes in infrastructure:
  – web caching, content distribution.

- Changes in applications:
  – telephony, streaming multimedia, peer-to-peer, multicast.

Invariant properties of the Internet:

- 24-hour cycles in traffic patterns.
- Log-normal connection sizes (for the main body of the distribution).
- Heavy-tailed distribution of connection sizes.
- Poisson arrivals for start times of user sessions.
- Self-similarity in traffic patterns.
- Invariants in topology?
- Heterogeneity and change!
  - [Paxson and Floyd, “Difficulties in Simulating the Internet”, 2001]

Do we know the traffic dynamics and protocols in the current Internet?

• Measurements of response times and packet loss rates:
The Internet Traffic Report, the Internet Weather Report.

• Measurements of packet size distributions, protocol breakdown.

• Where is the congestion in the Internet?

• How is the traffic on a link characterized in terms of round-trip times, end-to-end congestion experienced by the packets on that link, etc.?

• We don't know much about the actual deployment of queue management mechanisms, traffic engineering, and a wide range of other issues.

– [Web Page on Measurement Studies, "http://www.aciri.org/floyd/ccmeasure.html".]

Topics:

•

•

• Network dynamics, scalable congestion control, and other network issues relating to scale.

•

Why do we need end-to-end congestion control?

- As a tool for the application to better achieve its own goals:
  – Minimizing loss and delay, maximizing throughput.

- To avoid congestion collapse:
  – Congested links could be busy sending packets that will only be dropped downstream.

- Fairness:
  – In the absence of per-flow scheduling.

Scaling issues in congestion control:

- Flash crowds (heavy traffic all going to the same web site).
- Distributed Denial of Service attacks.
  - Problem: Limiting the damage to the legitimate traffic to the site.
  - Problem: Protecting the rest of the Internet.
- Achieving very-high-speed TCP flows.
  - And the stresses on the rest of the system.
- Allowing best-effort traffic to start-up more quickly.
  - And allowing web mice to finish sooner.

The future of congestion control in the Internet: several possible views:

- View #1: Infinite bandwidth:
  - No congestion, no problems.
- View #2: The current world continued:
  - Best effort traffic, with co-operative end-to-end congestion control.
- View #3: Ubiquitous per-flow scheduling:
  - The game theory view, with users optimizing their own utility functions.
- View #4: Congestion-based pricing, differentiated services:
  - Control through pricing.
- View #5: Virtual circuits:
  - Back to the past?
- The darker views: Congestion collapse and beyond.

Other scale-related stresses in the evolution of the Internet:

- DNS: (semantic overloading, etc.)
- IPv6.
- Measurement.
- Middleboxes.
- Multicast.
- IP over Optical.
- Quality of Service:
  - Do we need it? When will we get it?

- Routing:
  - The size of the routing table.
  - The rate of updates to the routing table.
  - Convergence times after updates.

- Security.

- Web Pages:
  - “Papers on the Evolvability of the Internet Infrastructure”.
  - URL: http://www.aciri.org/floyd/evolution.html

  - “Papers about Research Questions for the Internet”.
  - URL: http://www.aciri.org/floyd/research_questions.html

Topics:

-
-
-
- An aside about layers of abstraction in NS.

Different layers of abstraction in NS.

- Session-level Packet Distribution.
  - Designed for multicast simulations over large topologies
  - Topology, routing, and queueing delays are abstracted out, and replaced by delay and loss modules.

- Abstraction techniques for routing.
  - Manual routing.
  - Algorithmic routing: Centralized route computation.
  - Nix-vector routing: Routes are computed on demand.
  - Hierarchical routing: For reducing the routing table size.

- Abstraction techniques for highly-multiplexed traffic.
  - Goal: a stream between a source and destination, as cross-traffic, that behaves roughly the same as many TCP flows.

References for abstraction in NS.

- Web Page on “Running Large Simulations in NS”
  - URL http://www.isi.edu/nsnam/ns/ns-largesim.html

- P. Huang et al,

“Enabling Large-scale Simulations: Selective Abstraction Approach to The Study of Multicast Protocols”.

- P. Huang et al,

“Minimizing Routing State for Light-Weight Network Simulation”.