Congestion Control Principles

draft-floyd-cong-00.txt

Sally Floyd, Editor

IETF, Endpoint Congestion Management Working Group
November 1999

Outline of presentation:

- Current standards on end-to-end congestion control.
- The development of end-to-end congestion control in the Internet.
- The role of the standards process.
- Forms of end-to-end congestion control.
- TCP-specific issues.

Current standards on end-to-end congestion control:

- Standards on specific transport protocols:
  - E.g., TCP.

- Requirements for new transport protocols:
  - reliable multicast.

- Standards on communication between end-nodes and routers about congestion control or quality-of-service:
  - E.g., Explicit Congestion Notification (ECN),
  - differentiated services).

- This internet-draft: a general discussion of the role of the IETF in the standardization of congestion control mechanisms.

The development of end-to-end congestion control in the Internet:

- The prevention of congestion collapse.
- Fairness.
- Used by flows for their own purposes,
  - E.g., to maximize throughput, minimize delay and packet drops.

Congestion control for the prevention of congestion collapse:

- TCP in the early 80's:
  - TCP flow control to avoid overflowing receiver's buffer,
  - TCP's Go-Back-N retransmission.
  - FIFO scheduling, drop-tail queue management.

3

- A series of congestion collapses starting in 1986.

- Modern TCP retransmit timer and congestion control [Jacobson88]:
  - Packet drops as indications of congestion;
  - Additive Increase Multiplicative Decrease (AIMD), slow-start;
  - Exponential backoff of the retransmit timer.

Congestion control for fairness:

- For flows competing in a FIFO queue:
  - Compatible end-to-end congestion control mechanisms are required for some degree of fairness.

- Potential concerns about fairness:
  - Increasingly-aggressive, non-conformant TCP implementations;
  - A spiral of increasingly-aggressive transport protocols;
  - A spiral of increasingly-aggressive web browsers;
  - Best-effort traffic without end-to-end congestion control.

Congestion control for fairness, cont.:

- Terminology from RFC 2309:
  - TCP-compatible flow:
    - in steady-state, uses no more bandwidth than a conformant TCP under similar conditions.
  - unresponsive flow:
    - does not slow down in response to congestion.
  - responsive but not TCP-compatible:
    - responsive to congestion, but does not compete equally with TCP in a queue with FIFO scheduling.

Congestion control used by a flow for its own reasons:

- In an environment with per-flow scheduling or with small levels of statistical multiplexing:
  - A flow's delay and packet drop rate is in part a function of its own sending rate.

- In an environment with high levels of statistical multiplexing,
  - Tragedy of the commons is avoided because the "players" are not individuals but software vendors:
  - A flow's delay and packet drop rate is a function of the end-to-end congestion control provided by software vendors for all users.

A discussion of the role of the standards process:

Standardization of transport protocols, QoS mechanisms, ECN, etc.:

- Aspects traditionally subject to standardization:
  - Issues related to interoperability.
  - Mechanisms deemed critical to performance:

    [For standardized transport protocols, that is.]
    - Basic congestion control mechanisms;
    - Fairness with respect to other best-effort traffic;

- Traditionally not subject to standardization:
  - Implementation-specific issues.
  - Issues that do not affect interoperability,
    and do not significantly interfere with performance.

Forms of end-to-end congestion control:

- Concern: Avoiding congestion collapse from undelivered packets.
  - Solution: Congestion control to avoid a persistent high sending rate in presence of a high packet drop rate.

- Concern: Fairness with competing TCP traffic in a queue with FIFO scheduling?
  - Solution: TCP-compatible congestion control, such as:
    - AIMD with compatible increase/decrease parameters;
    - equation-based congestion control with a compatible equation;
    - layered multicast, receivers subscribing to layered multicast groups;
    - other forms...

TCP-specific issues:

- Slow-start:
  - Subject to standardization: Initial window.
  - Does not require standardization?:
    - Rate-based pacing;
    - Exiting slow-start early.

- AIMD:
  - Subject to standardization:
    - Increase/decrease constants;
    - Congestion control for return ACK path;
    - Window Increase based on byte-counting or ack-counting?
  - Does not require standardization?:
    - Interpretation of congestion window as sliding window, or limit to outstanding packets in the pipe.

TCP-specific issues, cont.:

- Retransmit timers:
  - Subject to standardization:
    - A proposal for more aggressive retransmit timers.
  - Does not require standardization?:
    - Modified mechanisms for setting retransmit timers that are not significantly more aggressive.

- Fast retransmit, fast recovery:
  - Subject to standardization:
    - Proposals for retransmitting packets based on one or two duplicate acknowledgements.
  - Does not require standardization?:
    - Proposals for sending a new packet based on one or two duplicate acknowledgements.