

Specifying New Congestion Control Algorithms

Sally Floyd and Mark Allman
draft-floyd-cc-alt-00.txt

November 2006

TSVWG

Slides: <http://www.icir.org/floyd/talks.html>

What is the problem?

- The problems:
 - There are many proposed congestion control mechanisms.
 - Some TCP implementations use congestion control that has not been through IETF process.
 - E.g., Linux and BIC TCP.
- Goals:
 - Encourage new congestion control mechanisms to go through IETF review.
 - Give guidelines for considering congestion control mechanisms for Experimental status.

Experimental status:

- Experimental RFCs for congestion control would indicate, in the abstract, the status:
 - Safe to deploy in the global Internet, or not?
 - Environments where the protocol is not recommended?
- Examples:
 - RFC 3649, HighSpeed TCP: safe to deploy.
 - Quick-Start: proposed for controlled environments.

Guidelines:

- Fairness to TCP?
- Using spare capacity?
- Difficult environments?
- Investigating a range of environments.
- Full backoff?
- Fairness within the proposed mechanism?
- Performance with misbehaving nodes and attackers?

Fairness to TCP:

- “In environments where standard congestion control is able to make reasonable use of the available bandwidth the proposed change should not significantly change this state.”
- “For instance, in a situation where each of N flows uses $1/N$ the network capacity, a new congestion control scheme should not significantly deviate from this state. For instance, a flow using an alternate congestion controller that took half the capacity and left each of the remaining N flows with $1/2N$ of the capacity would be suspect.”

Using spare capacity:

- “Alternate congestion control algorithms may take up spare capacity in the network, but may not steal significant amounts of capacity from flows using currently standardized congestion control.”

Difficult Environments.

- “An assessment of proposed algorithms in difficult environments such as paths containing wireless links and paths with reverse-path congestion. In addition, proposed algorithms should be evaluated in situations where the bottleneck has high and low levels of statistical multiplexing.”

Investigating a Range of Environments.

- “A particularly important aspect of evaluating a proposal for standardization is in understanding where the algorithm breaks down. Therefore, particular attention should be paid to extending the investigation into areas where the proposal does not perform well.”

Full Backoff

- “All alternate congestion control algorithms ultimately should include some notion of "full backoff". That is, at some point the algorithm should reduce the sending rate to one packet per round-trip time and then exponentially backoff the time between single packet transmissions if congestion persists. Exactly when this "full backoff" comes into play will be algorithm-specific. However, this requirement is crucial to protect the network in times of extreme congestion.”

Fairness within the Alternate Congestion Control Algorithm.

- “In environments with multiple competing flows using the alternate congestion control algorithm, the proposal should explore how bandwidth is shared among the competing flows.”

Performance with Misbehaving Nodes and Outside Attackers.

- “The proposal should explore how the alternate congestion control mechanism performs with misbehaving senders, receivers, or routers. In addition, the proposal should explore how the alternate congestion control mechanism performs with outside attackers. This can be particularly important for congestion control mechanisms that involve explicit feedback from routers along the path.”

Changes still to make:

- Add a bullet about **incremental deployment**.
- Clarify the **fairness** section. This is not saying that strict TCP-friendliness should be a requirement.
- Clarify that as **an alternative to Full Backoff**, a flow could stop sending when the allowed sending rate is below a certain threshold.
- Clarify that the Full Backoff bullet does not require that **different flows with different round-trip times** use the same criteria about when they should back off to one packet per round-trip time or less.
- ...