Network Working Group S. Floyd
Internet-Draft M. Allman
Expires: June 2006 ICIR / ICSI
December 2006

Specifying New Congestion Control Algorithms
draft-floyd-tsvwg-cc-alt-00.txt

Abstract

The IETF's standard congestion control schemes have been widely shown to be inadequate for various environments (e.g., high-speed or wireless networks). Recent research has yielded many alternate congestion control schemes. Using these new congestion control schemes in the global Internet has possible ramifications to both the network and to traffic using the currently standardized congestion control. Therefore, the IETF must proceed with caution when dealing with alternate congestion control proposals. The goal of this document is to provide guidance for considering alternate congestion control algorithms within the IETF.

TO BE DELETED BY THE RFC EDITOR UPON PUBLICATION:

Changes from draft-floyd-cc-alt-00.txt:

* Changed the name to draft-floyd-tsvwg-cc-alt-00.txt.

* Added a bullet about incremental deployment. Feedback from Colin Perkins

* Clarified the fairness section; this section is not saying that strict TCP-friendliness is a requirement.

* Clarified that as an alternative to Full Backoff, a flow could stop sending when the packet drop rate is above a certain threshold.

* Clarified 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.

* Added a paragraph about Informational RFCs.

* Added a bullet about response to transient events, including routing events or moving from a private to a shared network.

END OF NOTES TO BE DELETED.

1. Introduction

This document provides guidelines for the IETF to use when evaluating suggested congestion control algorithms that significantly differ from the general congestion control principles outlined in [RFC2914]. The guidance is intended to be useful to authors proposing alternate congestion control and for the IETF community when evaluating whether a proposal is appropriate for publication in the RFC series.

This document does not give hard-and-fast rules for what makes for an appropriate congestion control scheme. Rather, the document provides a set of criteria that should be considered and weighed by the IETF in the context of each proposal. The high-order criteria for any new proposal is that a serious scientific study of the pros and cons of the proposal needs to have been done such that the IETF has a well rounded set of information to consider.

After initial studies, we encourage authors to write a specification of their proposals for publication in the RFC series to allow others to concretely understand and investigate the wealth of proposals in this space.

2. Status

Following the lead of HighSpeed TCP, alternate congestion control algorithms are expected to be published as "Experimental" RFCs until such time that the community better understands the solution space. Traditionally, the meaning of "Experimental" status has varied in its use and interpretation. As part of this document we define two classes of congestion control proposals that can be published with the "Experimental" status. The first class is algorithms that are judged to be safe to deploy in the global Internet and further investigated in that environment. The second class is algorithms that, while promising, are not deemed safe enough for deployment on the Internet, but are being specified to facilitate investigations via simulation and testbeds.

Each alternate congestion control algorithm published is required to include a statement in the abstract indicating whether or not the proposal is considered safe for use on the Internet. Each alternate congestion control algorithm published is also required to include a statement in the abstract describing environments where the protocol is not recommended for deployment even though it may not be unsafe for use.

As examples of such statements, [RFC3649] specifying HighSpeed TCP includes a statement in the abstract stating that the proposal is Experimental, but may be deployed in the current Internet. In contrast, the Quick-Start document [QuickStart] includes a paragraph in the abstract stating the the mechanism is only being proposed for controlled environments. The abstract specifies environments where the Quick-Start request would give false positives (and therefore would be unsafe to deploy). The abstract also specifies environments where packets containing the Quick-Start request could be dropped in the network; in such an environment, Quick-Start would not be unsafe to deploy, but deployment would still not be recommended because it could cause unnecessary delays for the connections attempting to use Quick-Start.

For researchers who are not ready to bring their congestion control mechanisms to the IETF for standardization (either as Experimental or as Proposed Standard), one possibility would be to submit an internet-draft that documents the alternate congestion control mechanism for the benefit of the IETF and IRTF communities. This is particularly encouraged in order to get algorithm specifications widely disseminated to facilitate further research. Such an internet-draft could be submitted to be considered as Informational, as a first step in the process towards standardization. Such a document would also be expected to carry an explicit warning against using the scheme in the global Internet.

3. Guidelines

As noted above, authors are expected to do a well-rounded evaluations of the pros and cons of proposals brought to the IETF. The following are guidelines to help authors and the IETF community. Concerns that fall outside the scope of these guidelines are certainly possible and these guidelines should not be used as a check-list.

(1) Fairness to Standard TCP, SCTP, and DCCP.

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 of 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.

We note that this bullet is not a requirement for strict TCP-friendliness as a prerequisite for an alternate congestion control mechanism to advance to Experimental. As an example, HighSpeed TCP is a congestion control mechanism that is Experimental, but that is not TCP-friendly in all environments.

(2) Using Spare Capacity.

Similar to point (1), 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.

For instance, consider the case where N flows cannot naturally use all the capacity and equally share one-third of the capacity (with each flow using 1/3N of the capacity). A single flow using an alternate congestion control scheme could use the unused two-thirds of capacity. However, the flow using alternate congestion control should not steal significant amounts of additional capacity from the N standard flows.

(3) 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.

(4) Investigating a Range of Environments.

Similar to the last criteria, proposed alternate congestion controllers should be assessed in a range of environments. For instance, proposals should be investigated across a range of bandwidths and round-trip times. 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.

(5) Protection Against Congestion Collapse.

The alternate congestion control mechanism should either stop sending when the packet drop rate exceeds some threshold

[RFC3714], or should include some notion of "full backoff". For "full backoff", at some point the algorithm would 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 either "full backoff" or a pause in sending comes into play will be algorithm-specific. However, this requirement is crucial to protect the network in times of extreme congestion.

If "full backoff" is used, this bullet does not require that the full backoff mechanism must be identical to that of TCP. As an example, this bullet does not preclude full backoff mechanisms that would give flows with different round-trip times comparable bandwidth during backoff.

(6) 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.

(7) 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.

(8) Responses to Sudden or Transient Events

The proposal should consider how the alternate congestion control mechanism would perform in the presence of transient events such as sudden congestion, a routing change, or a mobility event.

(9) Incremental Deployment.

The proposal should discuss whether the alternate congestion control mechanism allows for incremental deployment in the targeted environment. If the alternate congestion control mechanism is intended only for specific environments, the proposal should consider how this intention is to be carried out.

4. Conclusions

This document is intended as a guideline for researchers in bringing congestion control mechanisms to the IETF to be considered for Experimental status, and also as a guideline to the IETF in evaluating such proposals.

5. Security Considerations

This document does not represent a change to any aspect of the TCP/IP protocol suite and therefore does not directly impact Internet security. The implementation of various facets of the Internet's current congestion control algorithms do have security implications (e.g., as outlined in [RFC2581]). Alternate congestion control schemes should be mindful of such pitfalls, as well.

6. IANA Considerations

This document does not require any IANA action.

Acknowledgments

Discussions with Lars Eggert and Aaron Falk seeded this document. Thanks to Colin Perkins for feedback. This document also draws from [Metrics]. Thanks to Colin Perkins for feedback.

Normative References

Informative References

[RFC2581] M. Allman, V. Paxson, and W. Stevens, TCP Congestion Control, RFC 2581, Proposed Standard, April 1999.

[RFC2914] S. Floyd, Congestion Control Principles, RFC 2914, Best Current Practice, September 2000.

[RFC3649] S. Floyd, HighSpeed TCP for Large Congestion Windows, RFC 3649, September 2003.

[RFC3714] S. Floyd and J. Kempf, IAB Concerns Regarding Congestion Control for Voice Traffic in the Internet, RFC 3714, March 2004.

[Metrics] S. Floyd, Metrics for the Evaluation of Congestion Control Mechanisms. Internet-draft draft-irtf-tmrg-metrics-04, work in progress, August 2006.

[QuickStart] S. Floyd, M. Allman, A. Jain, and P. Sarolahti, Quick-Start for TCP and IP. Internet-draft draft-ietf-tsvwg-quickstart-07.txt, work in progress, October 2006. Approved for Experimental.

Authors' Addresses

Sally Floyd
Phone: +1 (510) 666-2989
ICIR (ICSI Center for Internet Research)
Email: floyd@icir.org
URL: http://www.icir.org/floyd/

Mark Allman

ICSI Center for Internet Research
1947 Center Street, Suite 600
Berkeley, CA 94704-1198
Phone: (440) 235-1792
Email: mallman@icir.org
URL: http://www.icir.org/mallman/

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.