Internet Engineering Task Force
INTERNET-DRAFT
Intended status: Informational
Expires: 18 February 2008

S. Floyd
M. Allman
ICIR
18 August 2007

Comments on the Usefulness of Simple Best-Effort Traffic
draft-floyd-tsvwg-besteffort-01.ps

Status of this Memo

Copyright Notice

Abstract

This document presents some observations on "simple best-effort" traffic, defined loosely for the purposes of this document as Internet traffic that is not covered by Quality of Service mechanisms, congestion-based pricing, cost-based fairness, admissions control, or the like. One observation is that simple best-effort traffic serves a useful role in the Internet, and is worth keeping. While differential treatment of traffic can clearly be useful, we believe such mechanisms are useful as **adjuncts** to simple best-effort traffic, not as **replacements** of simple best-effort traffic. A second observation is that for simple best-effort traffic, some form of rough flow rate fairness is a useful goal for resource allocation, where "flow rate fairness" is defined by the goal of equal flow rates for different flows over the same path.

Table of Contents

Changes from draft-floyd-tsvwg-besteffort-00.txt:

* Added a sentence about peer-to-peer traffic opening many simultaneous connections in a session. From Tim Shephard.

* Added a discussion on the control of attacks, flash crowds, and the like. Feedback from Tim Shephard.

* Clarified the requirements of cost-based fairness in terms of the economic infrastructure. From feedback from Bob Briscoe:

* Clarified the definition of simple best-effort traffic. Feedback from Bob Briscoe.

* Added a citation to a paper by Jim Roberts on "Internet Traffic, QoS, and Pricing".

* Added a discussion of whether either the transport protocol (TCP vs. UDP) or the application should affect the fairness goals for simple best-effort traffic. Added a discussion of the precision of fairness. Feedback from Mitchell Erblich.

* Added a sentence to the discussion of byte vs. packet fairness about the bytes in packet headers. Feedback from Mitchell Erblich.

1. Introduction

This document gives some observations on the role of simple best-effort traffic in the

Internet. For the purposes of this document, we define "simple best-effort traffic" as traffic that does not *rely* on the *differential treatment* of flows either in routers or in policers, enforcers, or other middleboxes along the path, and that does not use admissions control. We define the term "simple best-effort traffic" to avoid unproductive semantic discussions about what the phrase "best-effort traffic" does or does not include. We note that our definition of "simple best-effort traffic" includes traffic that is not necessarily "simple", including mechanisms common in the current Internet such as pairwise agreements between ISPs, volume-based pricing, firewalls, and a wide range of mechanisms in middleboxes.

"Simple best-effort traffic" in the current Internet uses end-to-end transport protocols (e.g., TCP, UDP, or others), with minimal requirements of the network in terms of resource allocation. However, other implementations of simple best-effort service would be possible, including those that would rely on Fair Queueing or some other form of per-flow scheduling in congested routers. Our intention is to define "simple best-effort traffic" to include the dominant traffic class in the current Internet.

In contrast to "simple best-effort traffic", intserv or diffserv-enabled traffic relies on differential scheduling mechanisms at congested routers, with packets from different intserv or diffserv classes receiving different treatment. Similarly, in contrast to "simple best-effort traffic", cost-based fairness [B07] would most likely require the deployment of traffic marking (e.g., ECN) at congested routers, along with policing mechanisms near the two ends of the connection providing differential treatment for packets in different flows or in different traffic classes. Intserv/diffserv, cost-based fairness, and congestion-based pricing could also require more complex pairwise economic relationships among Internet Service Providers (ISPs), and between end-users and ISPs.

This document suggests that it is important to retain the class of "simple best-effort traffic" (though hopefully augmented by a wider deployment of other classes of service). Further, this document suggests that some form of rough flow-rate fairness is an appropriate goal for simple best-effort traffic. We do not argue in this document that flow-rate fairness is the *only possible* or *only desirable* resource allocation goal for simple best-effort traffic. We maintain, however, that it is an appropriate resource allocation goal for simple best-effort traffic in the current Internet, evolving from the Internet's past of end-point congestion control.

This document was motivated by [B07], an internet-draft on "Flow Rate Fairness: Dismantling a Religion" that asserts in the abstract that "Comparing flow rates should never again be used for claims of fairness in production networks." This document does not attempt to be a rebuttal to [B07], or to answer any or all of the issues raised in [B07], or to give the "intellectual heritage" for flow-based fairness in philosophy or social science, or to commit the authors of this document to an extended dialogue with the author of [B07]. This document is simply a separate viewpoint on some related topics.

2. On Simple Best-Effort Traffic

This section makes some observations on the usefulness and limitations of the class of simple best-effort traffic, in comparison with traffic receiving differential treatment.

2.1. The Usefulness of Simple Best-Effort Traffic

We now list some useful aspects of simple best-effort traffic.

Minimal technical demands on the network infrastructure:

Simple best-effort traffic, as implemented in the current Internet, makes minimal technical demands on the infrastructure. There are no technical requirements for scheduling, queue management or enforcement mechanisms in routers.

Minimal demands in terms of economic infrastructure:

Simple best-effort traffic makes minimal demands in terms of economic infrastructure, relying on fairly simple pair-wise economic relationships among ISPs, and between a user and their immediate ISP.

Usefulness in the real world:

Simple best-effort traffic has been shown to work in the Internet for the past 20 years, however imperfectly. As discussed below, simple best-effort traffic is not optimal. However, experience in the Internet has shown that there is value in having a mechanism that generally allows all users to get a portion of the resources, while still preventing congestion collapse.

2.2. The Limitations of Simple Best-Effort Traffic

We now discuss some limitations of simple best-effort traffic.

2.2.1. QoS

Some users would be happy to pay for more bandwidth, less delay, less jitter, or fewer packet drops. It is desirable to accommodate such goals within the Internet architecture while preserving a sufficient amount of bandwidth for best-effort traffic.

One of the obvious dangers of simple differential traffic treatment implementations that do not take steps to protect simple best-effort traffic would be that the users with more money *could* starve users with less money in times of congestion. There seems to be fairly widespread agreement that this would not be a desirable goal.

As a sample of the range of positions, the Internet Society's Internet 2020 Initiative, entitled "The Internet is (still) for Everyone", states that "we remain committed to the openness that ensures equal access and full participation for every user" [Internet2020].

The wide-ranging discussion of "network neutrality" in the United States includes advocates of several positions, including that of "absolute non-discrimination" (with no QoS considerations), "limited discrimination without QoS tiering" (no fees charged for higher-quality service), and "limited discrimination and tiering" (including higher fees allowed for QoS) [NetNeutral]. The proponents of "network neutrality" are opposed to charging based on content (e.g., based on applications, or the content provider).

As the "network neutrality" discussion make clear, there are many voices in the discussion that would disagree with a resource allocation goal of maximizing the combined aggregate utility, particularly where a user's utility is measured by the user's willingness to pay. "You get what you pay for" does not seem to be the consensus goal for resource allocation in the IETF or in the commercial or political realms of the Internet. However, there is a reasonable agreement that higher-priced services, as an adjunct to simple best-effort traffic, can play an important role in helping to finance the Internet infrastructure.

Briscoe argues for cost-fairness [B07], so that senders are made accountable for the congestion they cause. There are, of course, differences of opinion about how well cost-based fairness could be enforced, and how well it fits the commercial reality of the Internet, with [B07] presenting an optimistic view. Another point of view, e.g., from Roberts in a paper on "Internet Traffic, QoS, and Pricing", is that "many proposed schemes are overly concerned with congestion control to the detriment of the primary pricing function of return on investment" [R04].

With *only* simple best-effort traffic, there would be fundamental limitations to the performance that real-time applications could deliver to users. In addition to the obvious needs for high bandwidth, low delay or jitter, or low packet drop rates, some applications would like a fast start-up, or to be able to resume their old high sending rate after a relatively-long idle period, or to be able to rely on a call-setup procedure so that the application is not even started if network resources are not sufficient. There are severe limitations to how effectively these requirements can be accommodated by simple best-effort service in a congested environment.

2.2.2. The Avoidance of Congestion Collapse and the Enforcement of Fairness"

As discussed in Section 3.2 below, there are well-known problems with the enforcement of fairness and the avoidance of congestion collapse [RFC2914] with simple best-effort traffic. In the current Internet, end-to-end congestion control is relied upon to deal with these concerns; this use of end-to-end congestion control essentially requires cooperation from end hosts. '

2.2.3. Control of Traffic Surges

Simple best-effort traffic can suffer from sudden aggregate congestion from traffic surges (e.g., Distributed Denial of Service (DDoS) attacks, flash crowds), resulting in degraded performance for all simple best-effort traffic sharing the path. A wide range of approaches for detecting and responding to sudden aggregate congestion in the network have been proposed and used, including deep packet inspection and rate-limiting traffic aggregates. There are many open questions about both the goals and mechanisms of dealing with aggregates within simple best-effort traffic on congested links.

3. On Flow-Rate Fairness for Simple Best-Effort Traffic

This section argues that rough flow-rate fairness is an acceptable goal for simple best-effort traffic. We do not, however, claim that flow-rate fairness is necessarily an *optimal* fairness goal or resource allocation mechanism for simple best-effort traffic. Simple best-effort traffic and flow-rate fairness are in general not about optimality.

Within simple best-effort traffic, it would be possible to have explicit fairness mechanisms that are implemented by the end-hosts in the network (as in proportional fairness or TCP-fairness), explicit fairness mechanisms enforced by the routers (as in max-min fairness with Fair Queueing), or a traffic class with no explicit fairness mechanisms at all (as in the Internet before TCP congestion control).

This document does *not* address the issues about the implementation of flow-rate fairness. In the current Internet, rough flow-rate fairness is achieved by the fact that *most* of the traffic in the Internet uses TCP, and *most* of the TCP connections in fact use conformant TCP congestion control [MAF05]. However, rough flow-rate fairness could also be achieved by the use of per-flow scheduling at congested routers [DKS89] [LLSZ96], by related router mechanisms [SSZ03], or by congestion-controlled transport protocols other than TCP. This document does not address the pros and cons of TCP-friendly congestion control, equation-based congestion control [FHPW00], or any of the myriad of other issues concerning mechanisms for approximating flow-rate fairness. Le Boudec's tutorial on rate adaption, congestion control, and fairness gives an introduction to some of these issues [B00].

3.1. The Usefulness of Flow-Rate Fairness

We note that the limitations of flow-rate fairness are many, with a long history in the literature. We discuss these limitation in the next section. While the benefits of simple best-effort traffic and rough flow-rate fairness are rarely discussed this does *not* mean that benefits do not exist. In this section we discuss the benefits of flow-rate fairness. For simple best-effort traffic with rough flow-rate fairness, the quote from Winston Churchill about democracy comes to mind: "It has been said that democracy is the worst form of government except all the others that have been tried." Finally, we note that many of the useful aspects of simple best-effort traffic discussed above also qualify as useful aspects of rough flow-rate fairness.

Minimal technical demands on the network infrastructure:

First, the rough flow-rate fairness for best-effort traffic provided by TCP or other transport protocols makes minimal technical demands on the infrastructure, as TCP's congestion control algorithms are wholly implemented in the end-hosts. However, mechanisms for enforcement of the flow-rate fairness *would* require some support from the infrastructure.

Minimal demands in terms of economic infrastructure:

A system based on rough flow-rate fairness for best-effort traffic makes minimal demands in terms of economic relationships among ISPs or between users and ISPs.

Usefulness in the real world:

The current system---based on rough flow-rate fairness and simple best-effort traffic---has shown its usefulness in the real world.

Getting a share of the available bandwidth:

A system based on rough flow-rate fairness and simple best-effort traffic gives all users a reasonable chance of getting a share of the available bandwidth. This seems to be a quality that is much appreciated by today's Internet users (as discussed above).

3.2. The Limitations of Flow-Rate Fairness

This section discusses some of the limitations of flow-rate fairness for simple best-effort traffic.

3.2.1. The Enforcement of Flow-Rate Fairness

One of the limitations of rough flow-rate fairness is the difficulty of enforcement. One possibility for implementing flow-rate fairness would be an infrastructure designed from the start with a requirement for ubiquitous per-flow scheduling in routers. However, when starting with an infrastructure such as the current Internet with best-effort traffic largely served by First-In First-Out (FIFO) scheduling in routers and a design preference for intelligence at the ends, enforcement of flow-rate fairness is difficult at best. Further, a transition to an infrastructure that provides actual flow-rate fairness for best-effort traffic enforced in routers would be difficult.

A second possibility, which is largely how the current Internet is operated, would be simple best-effort traffic where most of the connections, packets, and bytes belong to connections using similar congestion-control mechanisms (in this case, those of TCP congestion control), with few if any enforcement mechanisms. Of course, when this happens, the result is a rough approximation of flow-rate fairness, with no guarantees that the simple best-effort traffic will continue to be dominated by connections using similar congestion-control mechanisms or that users or applications cannot game the system for their benefit. That is our current state of affairs. The good news is that the current Internet continues to successfully carry traffic for many users. In particular, we are not aware of reports of frequent congestion collapse, or of the Internet being dominated by severe congestion or intolerable unfairness.

A third possibility would be simple best-effort traffic with flow-rate fairness provided by the congestion control mechanisms in the transport protocols, with some level of enforcement, either in congested routers, in middleboxes, or by other mechanisms. [Add citations to some of the literature on this.] There seems to us to be considerable promise that incentives among the various players (ISPs, vendors, customers, standards bodies, political entities, etc.) will align somewhat, and that further progress will be made on the deployment of various enforcement mechanisms for flow-rate fairness for simple best-effort traffic. Of course, this is not likely to turn in to a fully-reliable and ubiquitous enforcement of flow-rate fairness, or of any related fairness goals, for simple best-effort traffic, so this is not likely to be satisfactory to purists in this area. However, it may be enough to continue to encourage most systems to use standard congestion control.

3.2.2. The Precise Definition of Flow-based Fairness

A second limitation of flow-based fairness is that there is seemingly no consensus within the research, standards, or technical communities about the precise form of flow-based fairness that should be desired for simple best-effort traffic. This area is very much still in flux, as applications, transport protocols, and the Internet infrastructure evolve.

Some of the areas where there are range of opinions about the desired goals for rough flow-based fairness for simple best-effort traffic include the following:

* Granularity: What is the appropriate fairness granularity? Should fairness be assessed on a per-flow basis? Should fairness take into account multiple flows between a pair of end-hosts (e.g., as suggested by [RFC3124])? If congestion control applies to each individual flow, what controls (if any) should constrain the number of connections opened between a pair of end-hosts? As an example, RFC 2616 specifies that with HTTP 1.1, a single-user client SHOULD NOT maintain more than two persistent connections with any server or proxy [RFC2616] (Section 8.1.4). For peer-to-peer traffic, different operating systems have different limitations on the maximum number of peer-to-peer connections; Windows XP Pro has a limit of ten simultaneous peer-to-peer connections, Windows XP Home (for the client) has a limit of five, and an OS X client has a limit of ten [P2P].

* RTT-fairness: What is the desired relationship between flow bandwidth and round-trip times, for simple best-effort traffic? As shown in Section 3.3 of [FJ92], it would be straightforward to modify TCP's congestion control algorithms so that flows with similar packet drop rates but different round-trip times would receive roughly the same throughput. This question is further studied in [HSMK98]. It remains an open question what would be the desired relationship between throughput and round-trip times for simple best-effort traffic.

* Multiple congested routers: What is the desired relationship between flow bandwidth and the number of congested routers along the path, for simple best-effort traffic? It is well established that for TCP traffic in particular, flows that traverse multiple congested routers receive a higher packet drop rate, and therefore lower throughput, than flows with the same round-trip time that traverse only one congested router [F91]. There is also a long-standing debate between max-min fairness and proportional fairness, and no consensus within the research community on the desired fairness goals in this area.

* Bursty vs. smooth traffic: What is the desired relationship between flow bandwidth and the burstiness in the sending rate of the flow? Is it a goal for a bursty flow to receive the same average or maximum bandwidth as a flow with a smooth sending rate? How does the goal depend on the time scale of the burstiness of the flow [K96]? For instance, a flow that is bursty on time scales of less than a round-trip time has different dynamics than a flow that is bursty on a time scale of seconds or minutes.

* Packets or bytes: Should the rough fairness goals be in terms of packets per second, or in bytes per second? And if the fairness goals are in terms of bytes per second, does this include the bandwidth used by packet headers (e.g., TCP and IP headers)?

* Different transport protocols: Should the transport protocol used (e.g., UDP, TCP, SCTP, DCCP) or the application affect the rough fairness goals for simple best-effort traffic?

* Unicast vs. multicast: What should the fairness goals be between unicast and multicast traffic?

* Precision of fairness: How precise should the fairness goals be? Is the precision that is possible from per-flow scheduling the right benchmark? Or, is a better touchstone the rough fairness over multiple round-trip times achieved by TCP flows over FIFO scheduling? Or, is a goal of even more rough fairness of an order of magnitude or more between flows using different transport protocols right?

There is a range of literature for each of these topics, and we have not attempted to cite it all above. Rough flow-based fairness for simple best-effort traffic could evolve with a range of possibilities for fairness in terms of round-trip times, the number of congested routers, packet size, or the number of receivers per flow. (Further discussion can be found in [METRICS].)

Fairness over time:

One issue raised in [B07] concerns how fairness should be integrated over time. For example, for simple best-effort traffic, should long flows receive less bandwidth in bits per second than short flows? For cost-based fairness or for QoS-based traffic, it seems perfectly viable for there to be some scenarios where the cost is a function of flow or session lifetime. It also seems viable for there to be some scenarios where the cost of QoS-enabled traffic is independent of flow or session lifetime (e.g., for a private Intranet that is measured only by the bandwidth of the access link, but where any traffic sent on that Intranet is guaranteed to receive a certain QoS).

However, for simple best-effort traffic, the current form of rough fairness that is not integrated over time seems acceptable. That is, in the current Internet, a user who opens a single TCP connection for ten hours *might* receive the same average throughput in bits per second, during that TCP connection, as a user who opens a single TCP connection for ten minutes and then goes off-line. Similarly, a user who is on-line for ten hours each day *might* receive the same throughput in bits per second, and pay roughly the same cost, as a user who is on-line for ten minutes each day. That seems acceptable to us. Other pricing mechanisms between users and ISPs seem acceptable also.

4. On the Difficulties of Incremental Deployment

One of the advantages of simple best-effort service is that it is currently operational in the Internet, along with the rough flow-rate fairness that results from dominance of TCP's congestion control.

While additional classes of service would clearly be of use in the Internet, the deployment difficulties of such mechanisms have been non-trivial [B03]. The problem of deploying interlocking changes to the infrastructure do not necessarily have an easy fix as they stem in part from the underlying architecture of the Internet. As explained in RFC 1958 on

"Architectural Principles of the Internet": "Fortunately, nobody owns the Internet, there is no centralized control, and nobody can turn it off [RFC1958]." Some of the difficulties of making changes in the Internet infrastructure, including the difficulties imposed by the political and economic context have been discussed elsewhere (e.g., [CMB07]). The difficulty of making changes to the Internet infrastructure is in contrast to the comparative ease in making changes in Internet applications.

The difficulties of deployment for end-to-end intserv or diffserv mechanisms are well-known, having in part to do with the difficulties of deploying the required economic infrastructure [B03]. It seems likely that cost-based schemes based on re-ECN could also have a difficult deployment path, involving the deployment of ECN-marking at routers, policers at both ends of a connection, and a change in pairwise economic relationships to include a congestion metric [B07].

[Add citations about the deployment difficulties of QoS, IPv6, multicast, ECN, and other mechanisms that have chicken-and-egg deployment problems.]

5. Related Work

5.1. From the IETF

This section discusses IETF documents relating to best-effort service and flow-rate fairness.

RFC 896 on congestion control: Nagle's RFC 896 on "Congestion Control in IP/TCP", from 1984, raises the issue of congestion collapse, and says that "improved handling of congestion is now mandatory" [RFC896]. RFC 896 was written in the context of a heavily-loaded network, the only private TCP/IP long-haul network in existence at the time (that of Ford Motor Company, in 1984). In addition to introducing the Nagle algorithm for minimizing the transmission of small packets in TCP, RFC 896 considers the effectiveness of ICMP Source Quench for congestion control, and comments that future gateways should be capable of defending themselves against obnoxious or malicious hosts. However, RFC 896 does not raise the question of fairness between competing users or flows.

RFC 2914 on congestion control principles: RFC 2914, a Best Current Practice document from 2000 on "Congestion Control Principles", discusses the issues of preventing congestion collapse, maintaining some form of fairness for best-effort traffic, and optimizing a flow's performance in terms of throughput, delay, and loss for the flow in question. In the discussion of fairness, RFC 2914 outlines policy issues concerning the appropriate granularity of a "flow", and acknowledges that end nodes can easily open multiple concurrent flows to the same destination. RFC 2914 also discusses open issues concerning fairness between reliable unicast, unreliable unicast, reliable multicast and unreliable multicast transport protocols.

RFC 3714 on the amorphous problem of fairness: Section 3.3 of RFC 3714, an Informational document from the IAB (Internet Architecture Board) discussing congestion control for best-effort voice traffic, has a discussion of "the amorphous problem of fairness", discussing complicating issues of packet sizes, round-trip times, application-level functionality, and the like [RFC3714].

RFC 2616 on opening multiple connections: RFC 2616, the standards track document for HTTP/1.1, specifies that "clients that use persistent connections SHOULD limit the number of simultaneous connections that they maintain to a given server" [RFC2616] (Section 8.1.4.).

RFC 2309 on unresponsive flows: RFC 2309, an Informational document from the End-to-End Research Group on "Recommendations on Queue Management and Congestion Avoidance in the Internet" in 2000, contains the following recommendation: "It is urgent to begin or continue research, engineering, and measurement efforts contributing to the design of mechanisms to deal with flows that are unresponsive to congestion notification or are responsive but more aggressive than TCP." [RFC2309]

RFCs on QoS: There is a long history in the IETF of the development of QoS mechanisms for integrated and differentiated services [RFC2212, RFC2475].

5.2. From Elsewhere

This section briefly mentions some of the many papers in the literature on best-effort traffic or on fairness for competing flows or users. [B07] also has a section on some of the literature regarding fairness in the Internet.

Fairness with AIMD: Fairness with AIMD (Additive Increase Multiplicative Decrease) congestion control was studied by Chiu and Jain in 1987, where fairness is maximized when each user or flow gets equal allocations of the bottleneck bandwidth [CJ89]. Van Jacobson's 1988 paper on "Congestion Avoidance and Control" defined TCP's AIMD-based congestion control mechanisms.

Fair Queueing: The 1989 paper of Fair Queueing by Demers et al. promoted Fair Queueing scheduling at routers as providing fair allocation of bandwidth, lower delay for low-bandwidth traffic, and protection from ill-behaved sources [DKS89].

Congestion-based pricing: One of the early papers on congestion-based pricing in networks is the 1993 paper on "Pricing the Internet" by MacKie-Mason and Varian [MV93]. This paper proposed a "Smart Market" to price congestion in real time, with a per-packet-charge reflecting marginal congestion costs. Frank Kelly's web page at [Proportional] has citations to papers on proportional fairness, including [K97] on "Charging and Rate Control for Elastic Traffic".

Other papers on pricing in computer networks include [SCEH96], which is in part a critique of some of the pricing proposals in the literature at the time. [SCEH96] argues that usage charges must remain at significant levels even if congestion is extremely low.

6. Security Considerations

This document does not propose any new mechanisms for the Internet, and so does not require any security considerations.

7. IANA Considerations

There are no IANA considerations in this document.

8. Conclusions

9. Acknowledgements

Informative References

[B00] J.-Y. Le Boudec, Rate adaptation, Congestion Control and Fairness: A Tutorial, 2000. URL "http://citeseer.ist.psu.edu/boudec00rate.html".

[B03] G. Bell, Failure to Thrive: QoS and the Culture of Operational Networking, Lawrence Berkeley National Laboratory, September 2003.

[B07] B. Briscoe, Flow Rate Fairness: Dismantling a Religion, internet-draft draft-briscoe-tsvarea-fair-02.txt, work in progress, July 2007.

[CJ89] Chiu, D.-M., and Jain, R., Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks, Computer Networks and ISDN Systems, V. 17, pp. 1-14, 1989. [The DEC Technical Report DEC-TR-509 was in 1987.]

[CMB07] kc claffy, Sascha D. Meinrath, and Scott O. Bradner, The (un)Economic Internet?, Internet Economics Track, 2007.

[DKS89] A. Demers, S. Keshav, and S. Shenker, Analysis and Simulation of a Fair Queueing Algorithm, SIGCOMM, 1989.

[F91] Floyd, S., Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1: One-way Traffic, Computer Communication Review, Vol.21, No.5, October 1991.

[FHPW00] Floyd, S., Handley, M., Padhye, J., and Widmer, J, Equation-Based Congestion Control for Unicast Applications, SIGCOMM, August 2000.

[FJ92] On Traffic Phase Effects in Packet-Switched Gateways, Floyd, S. and Jacobson, V., Internetworking: Research and Experience, V.3 N.3, September 1992.

[HSMK98] Henderson, T.R., E. Sahouria, S. McCanne, and R.H. Katz, "On Improving the Fairness of TCP Congestion Avoidance," Globecom, November 1998.

[Internet2020] Internet Society, "An Internet 2020 Initiative: "The Internet is (still) for Everyone", 2007. URL "http://www.isoc.org/orgs/ac/cms/uploads/docs/2020_vision.pdf".

[K96] F. Kelly, Charging and Accounting for Bursty Connections, In L. W. McKnight and J. P. Bailey, editors, Internet Economics. MIT Press, 1997.

[K97] F. Kelly, Charging and Rate Control for Elastic Traffic, European Transactions on Telecommunications, 8:33--37, 1997.

[LLSZ96] C. Lefelhocz, B. Lyles, S. Shenker, and L. Zhang, Congestion Control for Best-effort Service: Why We Need a New Paradigm, IEEE Network, vol. 10, pp. 10-19, Jan. 1996.

[MAF05] A. Medina, M. Allman, and S. Floyd, Measuring the Evolution of Transport Protocols in the Internet, Computer Communications Review, April 2005.

[METRICS] S. Floyd, Metrics for the Evaluation of Congestion Control Mechanisms, internet-draft draft-irtf-tmrg-metrics-10.txt, work in progress, July 2007.

[MV93] J. K. Mackie-Mason and H. Varian, "Pricing the Internet', in the conference on Public Access to the Internet, JFK School of Government, May 1993.

[NetNeutral] Network Neutrality, Wikipedia. URL "http://en.wikipedia.org/wiki/Net_neutrality". [Added for its citations to the literature.]

[P2P] "Maximum Number of Peer-to-Peer Connections", MAC OS X Hints web site, February 2007, URL "http://forums.macosxhints.com/showthread.php?t=67237".

[Proportional] Kelly, F., papers on Proportional Fairness.

[R04] J. Roberts, Internet Traffic, QoS and Pricing, Proceedings of the IEEE, V.92 N.9, September 2004.

[RFC896] Nagle, J., Congestion Control in IP/TCP, RFC 896, January 1984.

[RFC1958] B. Carpenter, Architectural Principles of the Internet, RFC 1958, June 1996.

[RFC2212] Shenker, S., Partridge, C. and R. Guerin, Specification of Guaranteed Quality of Service, RFC 2212, September 1997.

[RFC2309] B. Braden at al, Recommendations on Queue Management and Congestion Avoidance in the Internet, RFC 2309, April 1998.

[RFC2475] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z. and W. Weiss, An Architecture for Differentiated Services, RFC 2475, December 1998.

[RFC2616] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P. and T. Berners-Lee, Hypertext Transfer Protocol -- HTTP/1.1, RFC 2616, June 1999.

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

[RFC3124] H. Balakrishnan and S. Seshan, The Congestion Manager, RFC 3124, June 2001.

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

[SCEH96] Shenker, D. D. Clark, D. Estrin, and S. Herzog, Pricing in Computer Networks: Reshaping the Research Agenda, ACM Computer Communication Review, vol. 26, April 1996.

[SSZ03] I. Stoica, S. Shenker, and H. Zhang, Core-Stateless Fair Queueing: a Scalable Architecture to Approximate Fair Bandwidth Allocations in High-speed Networks, IEEE/ACM Transactions on Networking 11(1): 33-46, 2003.

Authors' Addresses

Sally Floyd
ICSI Center for Internet Research
1947 Center Street, Suite 600
Berkeley, CA 94704
USA
EMail: floyd <at> icir <dot> 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 at icir.org
URL: http://www.icir.org/mallman/

Full Copyright Statement

Intellectual Property

ipr@ietf.org.