Engineering a Reliable and Efficient Internet: Priority Scheduling, Congestion, and Security

Dr. Charles L. Jackson
21 July 2009

Table of Contents

Introduction and Overview

The Internet is one of the great technological successes of the last several decades. It has gone from being a research project used by a few specialists to an essential tool for most businesses and for a large fraction of consumers. Policies that facilitate the wider availability and adoption of broadband access to the Internet would promote a wide variety of public interest objectives, including jobs, safety-of-life, and quality-of-life. At the same time, restrictive regulations that tie the hands of network engineers and managers and prevent continued innovation would have the result of making broadband networks less robust, less useful, and less secure than they could otherwise be, while also denying consumers the choice of certain services that may be effectively precluded in the absence of particular forms of network management The successful operation of a broadband network requires considerable attention by network operators to many significant background details, such as protecting against security threats, controlling congestion, and making sure that delay-sensitive applications like VoIP and interactive games perform well. Allowing providers the flexibility to employ the tools and practices that most effectively address these concerns benefits all broadband consumers and users.[1]

Some commenters seem to oppose any form of congestion control or priority routing (which gives one class of traffic priority over another class) in the public Internet. Some go so far as to assert that priority routing could provide no benefits—that it is a "zero sum game" because speeding one packet slows another.[2] By the same logic, it would make no sense for commuters to pull off the highway to allow an ambulance to pass—an accident victim may get to the hospital 10 minutes sooner but a hundred commuters will get home 6 seconds later. But effective congestion control, including priority routing can make otherwise useless resources useful. In modern wireless systems, for example, priority routing of packets creates new useful capacity and allows services that otherwise would

[1] The author thanks Verizon for support and assistance in developing this report. The opinions expressed here are my own and are not necessarily the views of any other entity.

[2] Comments of Free Press at pages 145 and 164.

not be available to consumers.[3]

Communications engineering has long focused on the issues of priority routing and congestion control. The discipline of queueing theory was born in communications engineering but has now moved to much wider applications.

Congestion has long been a real problem for the Internet. In the 1980s, there were several congestion collapses of the Internet—which led to improved and widely accepted congestion control mechanisms.[4] Congestion remains a problem, as evidenced by more recent losses of network capacity and severe network congestion. Moreover, as evidenced by cyberattacks by some malicious actors, congestion can be used as a powerful weapon to harm particular users, networks, or the Internet as a whole.

Priority routing can, among other things, be an effective tool at controlling and minimizing the harms of congestion. Giving one class of traffic priority over another can substantially reduce the harms from congestion by enabling latency-sensitive applications that would fail in the absence of network management. Moreover, in the wireless world, giving some traffic priority over others permits expanding capacity without imposing significant costs.

Below I review these related issues from an engineering perspective. That is, I look at how priority routing and congestion control can improve system efficiency. There is a large economic literature regarding the benefits of congestion control through pricing or other priority routing technologies. My comments necessarily touch on this topic, but it is not a focus of this paper. However, it is important to note that economists have

---

[3] Technically speaking, in most cases priority scheduling does not really create new capacity. Rather, it allows the use of capacity that cannot be exploited under a single-priority regime or it allows the more efficient use of capacity.

[4] See "Congestion Avoidance and Control," V. Jacobson, *Proceedings of SIGCOMM '88,* ACM. This is frequently referred to as authored by Van Jacobson and Michael Karels because an updated version was prepared by both authors. The updated version is available at http://www-nrg.ee.lbl.gov/papers/congavoid.pdf.

repeatedly shown that various forms of priority routing, often called *congestion pricing,* can create enormous benefits.[5]

I begin by examining modern wireless to show how giving more-urgent packets priority over less-urgent packets can create a system that delivers more capacity to users than does a system without such priority routing. Or, stating it another way, imposing a rule that all packets must be treated equally would reduce the capacity of such a system without improving the performance of the high-priority service. I then discuss how the same principle applies to both satellite and electrical power capacity.

I then turn to the issue of congestion and congestion control. I begin with congestion control in the Internet as it has been practiced in the past and as it is practiced today. I also describe recent incidents of system collapse and how blocking low-priority traffic was a key factor in recovering from such collapses. I conclude that congestion controls within the network—congestion controls that do not treat each packet equally—offer substantial benefits for consumer welfare and public safety.

The third topic I address, cross-layer design, also offers many potential benefits to the users of broadband networks but would also be threatened by some proposals seeking to prohibit any differentiation between bits. *Cross-layer design* refers to the development of networking technologies that gain efficiency at the expense of complexity by simultaneously operating at more than one layer of the Internet's protocol stack.[6] Cross-layer design has proven to be particularly valuable in wireless networks, although it creates benefits in other networks as well. The basic concept is simple. A lower-layer protocol, such as the link level on a radio or Ethernet connection, can deliver more value to users if it takes into account the higher-level protocol being carried over the

---

[5] For an exposition of congestion pricing in transportation, see *Congestion Pricing: A Primer,* U.S. Department of Transportation, Federal Highway Administration, FHWA-HOP-08-039, October 2008. Available at http://ops.fhwa.dot.gov/publications/fhwahop08039/fhwahop08039.pdf. The Obama administration has strongly supported the development of a smart grid for the electrical power network that would permit the use of peak-load or time-of-day pricing (congestion pricing by another name) for electricity.

[6] This layering is also referred to as the ISO model. See "OSI Reference Model — The ISO Model of Architecture for Open Systems Interconnection," H. Zimmerman, IEEE Transactions on Communications, 28(4):425-432, April 1980.

connection. I review some circumstances in which cross-layer design can deliver significant efficiencies. Prohibiting particular approaches to network management, such as packet inspection technologies, would forfeit such efficiencies.

Fourth, I describe how certain tools, technologies, and congestion control techniques criticized by some parties to this proceeding – including packet inspection technologies – can provide highly effective defenses against attacks against networks—in particular against denial-of-service attacks.

As this discussion will show, imposing any form of a rule that prohibits any differential treatment or handling of different packets would create substantial efficiency losses by prohibiting the use of technologies that expand capacity, protect against congestion, and enable services or applications that would otherwise not function effectively. Such a rule would also make broadband networks less robust and less secure than they would otherwise be. Finally, such a rule would handicap innovation in both network services and network-hosted applications.

Consumer Benefits from Priority Routing

Despite the best efforts of their designers and operators, nature makes the capacity of systems variable—best modeled as a random quantity. Consider the capacity of the airways between Washington, DC, and New York. Although there is an upper limit set by the capacity of the airports at each end, weather often reduces capacity well below that upper limit. The supply of electricity also fluctuates. Generators fail, river flows and winds vary, and transmission lines fail. Nature imposes a similar random fluctuation on the capacity of communications networks. For example, the capacity of the wireless channels used for cellular and PCS fluctuates over fractions of a second. Dividing wireless traffic into high-priority and low-priority streams increases the benefits that can be delivered by wireless channels to consumers. Similarly, the capacity available for particular users of wireline and wireless networks can be dramatically affected when many users (or many applications initiating multiple session flows) attempt to use a shared link at the same time.

Scheduling Wireless Transmissions

Modern wireless voice networks transmit signals to and from user handsets over radio channels that carry many conversations simultaneously. The quality of the radio signal received by each user varies rapidly—received signal strength can change by a factor of 10 within as little as a hundredth of a second. If the radio signal received by user A becomes weaker, say because he or she has just stepped away from the window in a building, the base station in the wireless system must increase the power it uses to transmit to user A or the telephone call will be lost. Most of the time, another user's radio channel, say user B's channel, improves at the same time and the power used to transmit to user B can be lowered—consequently the increases and the decreases cancel and total power from the base stays even.

However, sometimes the increases and decreases do not cancel out—and many users need extra power. If a user needs more power on the downlink but the power cannot be increased, the call will be lost. Wireless systems protect against the threat of such failures by keeping some power in reserve—they restrict the number of calls served on a single radio link so that there will be such a power reserve. Consequently, on those occasions when substantially more than average power is needed, the system can draw on the reserve and avoid dropping any calls.

At times when the reserve power is not needed for the voice service, the reserve power can be put to effective use for data services, thus making better use of the finite capacity available in the system. In order to keep the voice service working acceptably, this data service must necessarily be lower priority than the voice service. At times, the voice service would demand all the downlink power and the data service would have to be suspended for as long as several hundred milliseconds. Nevertheless, a data service with substantial capacity—about 50% of the throughput on the voice channels in some circumstances—can be created this way if the system is able to schedule voice packets for transmission ahead of packets for the data service.

This is not a hypothetical analysis. Multiple studies have shown this to be the case for both cdma2000 and WCDMA.[7] Mehmet Yavuz and his coworkers at QUALCOMM report that

> DO-Rev A can provide VoIP capacity comparable to circuit-switched cellular CDMA technologies (e.g., IS-2000) and *simultaneously* carry significant amount of other types of traffic such as non-delay sensitive applications and downlink multicast.[8]

Ozcan Ozturk and his coauthors, also at QUALCOMM, state,

> Simulations also show that a significant amount of best effort traffic can still be served on the downlink at the VoIP capacity operating point.[9]

Imposing a rule on wireless systems that prohibits any differential treatment of packets would present a system operator with a choice between (1) running the system but restricting traffic to the level consistent with high-quality voice or (2) running the system with more traffic but delivering a service with delay and jitter that would make voice service unacceptable. If the operator chooses to offer voice—the all-time killer application—then the traffic capacity offered by the reserve power would be wasted.

The heart of this issue in wireless arises from the fact that the capacity of the wireless link varies randomly over times that are short compared with a phone call but that can be long compared with the duration of a word. Humans find it hard to deal with telephone services in which occasional words are missing—there is a big difference in meaning between "Don't call me after 11:00 PM" and "Call me after 11:00 PM." Because people cannot tolerate such dropouts, the wireless system must have enough reserve power to cope with the variations in the radio channel. Similarly, people dislike phone service that often drops calls. In contrast, an email transfer that sometimes is blocked from accessing the radio channel for a second or two works just fine for most people. Consistent with widely accepted practices throughout the industry, priority routing is the tool that lets these differing demands of voice and data customers be satisfied. In this case, priority

---

[7] For example, see "VoIP over cdma2000 1xEV-DO Revision A," M. Yavuz et al, *IEEE Communications Magazine,* February 2006, at p. 88; "Class-Based Quality of Service over Air Interfaces in 4G Mobile Networks," Y. Guo and H. Chaskar, *IEEE Communications Magazine,* March 2002, at p. 132; and "Performance of VoIP Services over 3GPP WCDMA Networks," O. Ozturk et al., *IEEE 19th International Symposium on Personal, Indoor and Mobile Radio Communications,* September 2008, at p. 1.

[8] Op. cit. at p. 88, emphasis in original.

[9] Op. cit. at p. 5.

routing is clearly not a zero-sum game. Priority routing permits use of resources that would otherwise sit idle.

Clearwire's wireless data network illustrates the benefits of priority scheduling from a different starting point. There is a well-known tradeoff between latency and capacity in mobile data networks.[10] The basic idea is simple. As described above, the quality of the radio channel to each user varies over time. If the system waits to transmit data to a specific user until the channel to that user is particularly good, the system capacity is increased. The longer the system is willing to wait for a good channel, the more capacity becomes available.[11] Today, Clearwire is selling a basic Internet access service for mobile users. They can deliver more megabits per second to their customers if they run their network with relatively high latency—perhaps a several hundred milliseconds—during periods of heavy usage.[12] But, such high latency would make it virtually impossible to place a useful VoIP call over their network. If Clearwire were able to offer a priority service option—an option that would allow users to shorten the latency on some packets—users would be able to buy the right to use VoIP. Prohibiting Clearwire from offering to their users a priority service option would essentially prevent Vonage or Skype from serving those Clearwire users that need to make phone calls when the load on the Clearwire network is high. More generally, prohibiting ISPs from offering priority services handicaps all application providers whose applications require connections capable of minimizing jitter or latency.

Scheduling and Priority in Satellites and Electricity

Nature has imposed similar random fluctuations on the capacity of other types of important services. The capacity of some geostationary communications satellites comes in physical units called transponders. A satellite might have 24 transponders. Satellite

---

[10] The formal name for the technology that exploits this tradeoff is *multi-user diversity*.

[11] Xin Liu and other researchers at Sprint used measurements on an EV/DO system to determine that such scheduling of transmissions can add about 20% to system capacity. See "Experiences in a 3G Network: Interplay between the Wireless Channel and Applications," Xin Liu et al., ACM MobiCom'08, 2008. Available at https://research.sprintlabs.com/publications/uploads/com0765-liu.pdf.

[12] I subscribed to the Xohm 2.5 GHz WiMax service that has now become part of Clearwire and tested its performance. I measured download rates between 3 and 4 Mbps and upload rates between 1.5 and 2.5 Mbps. Vonage calls worked quite poorly with gaps and delays. I believe that most people would not be willing to accept the call quality that I experienced. It appears that the Xohm network was optimized for capacity at the expense of latency.

providers often sell the capacity of an entire transponder to a customer. Unfortunately, transponders are like computers or refrigerators—they can work fine for months or years and then unexpectedly fail. Satellite carriers and satellite users have a good idea of the probability of these failures. Thus, at the time that a 24-transponder satellite is launched, a planner might expect that 5 years later there would be a 100% chance that the satellite would have 20 or more working transponders, a 50% chance of having 22 or more working transponders, and a 10% chance of having all 24 transponders working.

As is the case for the wireless channels described above, the capacity of a satellite varies randomly. The satellite industry deals with this uncertainty by offering three types of transponder services—protected, unprotected, and preemptible. *Protected service* provides the highest reliability. If a protected transponder fails, the user's traffic is transferred to a different transponder that is still working. *Unprotected service* provides less reliability but costs less. If an unprotected transponder fails, the user is out of luck—the user loses the satellite link through that transponder. *Preemptible service* provides the least reliability. When a protected transponder fails, a user of a preemptible transponder may see service terminated in order to free up a transponder for the user with protected service. If there were a rule that all satellite transponders had to be offered on the same terms, then either (1) a user who needed highly reliable service, say a TV programming service, would need to rent multiple transponders in order to ensure access to backup capacity or (2) the satellite operator would need to keep the backup transponders idle. Giving some transponder users priority over others increases the total value delivered by the satellite system. Moreover, it makes available to users several price/service quality options.

Electrical power systems also have uncertain capacity because generators fail, transmission lines fail, river flows vary, and the wind is stronger at some times than at other times. Naturally enough, and for much the same reasons as described above, wholesale electric power producers sell products such as firm power and interruptible power.[13] Interruptible power would be unacceptable for most homes and businesses.

---

[13] See the Bonneville Power Administration glossary at http://www.bpa.gov/power/pl/subscription/prodglos.htm. The power industry also faces variations in

However some commercial uses of electricity, such as refining aluminum or pumping water for irrigation, can be operated efficiently on interruptible power.

Summing Up

These systems—wireless, satellite, and electrical power—are all subject to random fluctuations in system capacity that arise from the nature of the physical world. These fluctuations create the opportunity to define new products and services that make the best use of a given system, consistent with its capacity and limitations. If all satellite users needed exactly the same service, then there would be no separate market for preemptible transponders. If all electrical power users were like hospitals or office buildings, then there would be no market for interruptible power.

A wireless system engineered to support human conversation may have no more capacity for telephone calls but may still have capacity to carry delay-tolerant packets. Because some Internet applications are far more tolerant of delay than are human conversations, this additional capacity can be used to deliver useful service to consumers. Similar considerations apply to other broadband networks—with "smarter" networks better able to make the most out of the finite capacity that exists in any given system. A rule prohibiting any differential treatment of packets, i.e., that no priority be afforded one class of packets over another, would block consumer access to this additional capacity and prevent the efficient use of the radio spectrum and of the base stations and radios used to communicate across that spectrum.

Although the above discussion focuses on wireless networks because the underlying physical system varies over time, demand variations create essentially identical concerns in the wireline world. For example, it is well known that when multiple users go online at the same time—such as when kids leave school in the afternoon—the resulting congestion can affect the latency and jitter experienced by cable modem users competing for the finite and shared resource. In that context as well, approaches that differentiate between latency-sensitive traffic and other traffic could yield substantial consumer

---

demand and offers a variety of user pricing mechanisms designed to limit peak demand or to move demand from peak to off-peak times. The application of congestion pricing to energy through Advanced Metering Infrastructure is a key part of the Department of Energy's Smart Grid policy. See the Department of Energy publication The Smart Grid: An Introduction, available at http://www.oe.energy.gov/1165.htm.

benefits and enable services that otherwise might not function well or at all at times of congestion.

Congestion and Congestion Control

Congestion in the Internet is not merely a theoretical concern—it has long presented a real-world challenge for network engineers. A famous paper by Van Jacobson and Michael Karels describes several congestion collapses of the Internet.[14] The development of effective congestion control mechanisms was a key step in developing the modern Internet. Unfortunately, the primary congestion control mechanisms in today's Internet depend on the honor system for their effective operation. Incompetent or malicious programmers may subvert the honor system and set the stage for congestion failures. Happenstance, malicious acts, or equipment failure may also lead to congestion failures. Congestion is not just a problem of the 1980s, as evidenced by more recent system collapses.

The early Internet suffered a series of congestion collapses in the mid-1980s.[15] The collapses arose from a simple cause—users were transmitting more data on some paths than the paths could handle. Router queues would fill up, and subsequently arriving packets would be discarded. User machines would retransmit the lost packets, and congestion would continue. The Internet congestion was like the Beltway in Prince Georges County after a Redskins home game—except for the retransmissions.[16]

Internet Congestion Control on the Honor System

In 1993, researcher Van Jacobson of Laurence Berkeley Laboratory described the congestion problem and the solution that he and his coworkers developed:

> "If too many people try to communicate at once," explains Jacobson, "the network can't deal with that and rejects the packets, sending them back. When a workstation retransmits immediately, this aggravates the situation. What we did

---

[14] Op. Cit.

[15] Van Jacobson and Karels state, "In October of '86, the Internet had the first of what became a series of 'congestion collapses'. During this period, the data throughput from LBL to UC Berkeley (sites separated by 400 yards and two IMP hops) dropped from 32 Kbps to 40 bps. We were fascinated by this sudden factor-of-thousand drop in bandwidth and embarked on an investigation of why things had gotten so bad."

[16] Skin's fans stuck in a traffic jam are not magically cloned in the parking lot to start out again and add even more to the congestion.

> was write polite protocols that require a slight wait before a packet is retransmitted. **Everybody has to use these polite protocols or the Internet doesn't work for anybody**."[17]

Substantial thought and research went into developing congestion control mechanisms that have been embedded in TCP implementations. Although these methods are complex and subtle, the basic idea is simple—if a server or user terminal senses that the network seems to be losing packets, the server or user terminal should cut back sharply the rate at which it is transmitting data. Putting congestion control in the user devices at the edge of the network made sense for many reasons, and over the next few years, TCP implementations included congestion control features and such congestion failures became far rarer and more localized.[18]

But it is widely recognized that the fundamental problem still remains. There is finite capacity at many points in a network. If the traffic delivered to such a point exceeds the traffic that the point can forward on, some packets must be discarded. Furthermore, today's Internet congestion control works mostly on the honor system. Windows, Linux, and the Apple operating systems all come with TCP congestion control built in, but users can install software that violates the honor system.

Claiming that congestion control on the Internet works on the honor system is not a metaphorical usage for emphasis—it is a statement of fact. Users' systems must act altruistically, sacrificing their network service for the greater good, in order for these congestion control approaches to be effective. A recent (May 2009) publication by the Internet standards body, the Internet Engineering Task Force (IETF), makes this point, saying,

> In the current Internet architecture, **congestion control depends on parties acting against their own interests.** It is not in a receiver's interest to honestly return feedback about congestion on the path, effectively requesting a slower transfer. It is not in the sender's interest to reduce its rate in response to

---

[17] "Building and Rescuing the Information Superhighway," Jeffery Kahn, *LBL Research Review,* Summer 1993. Available at http://www.lbl.gov/Science-Articles/Archive/information-superhighway.html. Emphasis added.

[18] The reasons that deploying congestion control at the edges was appropriate included the facts that deploying changes to user and server software can be easier than changing routers, that user and server computers have more computing capacity available for managing such congestion, and that a key part of congestion control is a change in the behavior of devices connected to the network.

> congestion if it can rely on others to do so. Additionally, networks may have strategic reasons to make other networks appear congested.[19]

A recent textbook makes much the same point:

> It is possible for an ill-behaved source (flow) to capture an arbitrarily large fraction of the network capacity. . . . Such an application is able to flood the Internet's routers with its own packets, thereby causing other applications' packets to be discarded.[20]

Despite the success of TCP congestion control mechanisms developed in the 1980s and 1990s, researchers have remained concerned about the threat of congestion caused by software that violates the honor code. In 1998, for example, a group of prominent computer scientists authored RFC 2309, titled *Recommendations on Queue Management and Congestion Avoidance in the Internet,* setting forth some of their concerns.[21] The 15 authors of this RFC include many of the best-known researchers on congestion control in the Internet. The authors repeatedly express concern about **"**the potential for future congestion collapse of the Internet" and describe scenarios in which "the Internet is chronically congested."[22] In particular, they address congestion from applications which "can grab an unfair share of the network bandwidth."[23] As the authors recognize, software with the capability to do exactly was available a decade ago. Such software is far more widespread today.

In the web-services context, persistent connections are TCP connections that are kept alive over time in order to speed web-server response by avoiding connection set up delays. Persistent connections speed up web downloading but they can impose higher traffic bursts than newly-established connections. If a user kept a large number of persistent connections open to a web server, he could download multiple files quickly—but at the risk of creating congestion problems on the route between the web server and the user's computer. Consequently, Internet standards recommend that web browsers not

---

[19] IETF Working Group Draft, "Open Research Issues in Internet Congestion Control," Scharf and Briscoe, May 2009, available at http://tools.ietf.org/html/draft-irtf-iccrg-welzl-congestion-control-open-research-04. (expires November 2009)

[20] *Computer Networks: A Systems Approach,* L. Peterson and B. Davie, 4th ed., Morgan Kauffman, 2007, at p. 470.

[21] RFCs are the standardization documents for the Internet and are published by the IETF. They are available at http://www.ietf.org/rfc.html.

[22] RFC 2309

[23] Ibid.

have more than two persistent connections to a single web site.[24] But, not all web browsers follow this recommendation. The extensively used Firefox web browser, for example, allows the user to edit some of the networking settings. Figure 1 shows the control panel of an add-in that simplifies that editing with the number persistent connections per server on my web browser set to 16 and the maximum connections per server set to 64.[25] These settings improve performance for me, but they clearly violate the honor system of the Internet and have the potential (particularly if widely used) to hinder the overall performance of the network and to degrade the service of other users.

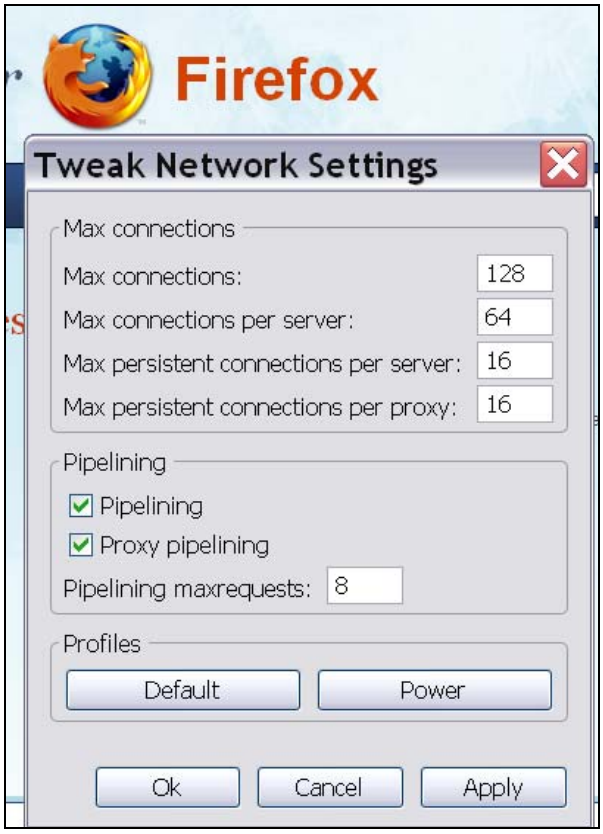

Figure 1. Firefox Networking Control Panel Showing a Maximum of 16 Persistent Connections Rather than the Maximum of 2 of RFC 2616.

[24] RFC 2914 states, "The specific issue of a browser opening multiple connections to the same destination has been addressed by RFC 2616 [RFC2616], which states in Section 8.1.4 that 'Clients that use persistent connections SHOULD limit the number of simultaneous connections that they maintain to a given server. A single-user client SHOULD NOT maintain more than 2 connections with any server or proxy.'" (emphasis in original). Figure 1 shows my Firefox browser configured to maintain 16 connections to a server or proxy—that is 8 times more than the number in the standard. This set up is illustrative. I run my browser with the default settings, not these greedy settings. Of course, the default setting is 6—triple the recommended number.

[25]

The Internet community is well aware of the congestion risk created by nonconformant applications such as the Firefox browser. For example, an Agilent white paper states,

> *Mischievous Applications* - In spite of efforts to modify TCP or queue management to improve fairness, achieve better link utilization, and so on, an important consideration is that applications themselves are evolving to exploit the nature of networks and take an unfair share of bandwidth. For example, the open-source browser Firefox opens multiple TCP connections in attempt to manipulate the network. More widespread and problematic are peer-to-peer applications such as BitTorrent that make multiple small requests over different TCP connections, ultimately defeating the principle of fairness that TCP and queue management researchers seek to uphold. Properly managing such mischievous applications requires going beyond dealing with individual flows or connections.[26]

Sophisticated users and developers of applications are also well aware of both the potential individual benefits and collective harms of violating the congestion-control honor code. Here is a statement from a blog describing how to improve Firefox performance.

> Bear in mind however that the more connections you are tying up, the less that will be available to others wishing to connect to the same server - so don't set this excessively high just because you can.[27]

Web browsers are not the only software that may violate the honor code of the Internet and contribute disproportionately to network congestion and increased delay. Some peer-to-peer software also does. The Agilent white paper notes that BitTorrent can open dozens of TCP connections to download a file—thus greatly speeding downloading but risking congestion and possibly taking an unfair share of network resources.[28] Agilent's reference to taking an unfair share of network resources reflects the fact that if two users are sharing a communications link—one using a web browser to view a video feed from

---

[26] *TCP and Queue Management,* Agilent Technologies, 2008. Available at http://cp.literature.agilent.com/litweb/pdf/5989-7873EN.pdf

[27] See http://pinguy.infogami.com/blog/39l5. Other blogs also suggest tuning Firefox to increase performance but do not explain the negative consequences for others. See http://www.blogsdna.com/372/21-aboutconfig-hackstweaks-for-firefox-3.htm, http://www.computerworld.com/action/article.do?command=viewArticleBasic&taxonomyName=Networking+and+Internet&articleId=9020880&taxonomyId=16&pageNumber=5, and http://maketecheasier.com/28-coolest-firefox-aboutconfig-tricks/2008/08/21. The help page for the Opera browser states, "It is recommended to keep the default setting of 8 [maximum connections to a server], but you can try changing the maximum number of connections to a single server if you are experiencing problems with browsing speed."

[28] BitTorrent opens multiple TCP connections that together are less responsive to congestion than a single TCP connection. See the discussion of BitTorrent below.

Hulu and the other using BitTorrent to download a movie, the BitTorrent user might receive 50 times as much of the link's capacity than would the viewer of the video. This unfair sharing would not create a problem if the link had 100 times more capacity than needed to view the video stream. But, if the link had only 10 times as much capacity as needed to view the video stream, the Hulu user would get about one-fifth of a video channel and the BitTorrent user would get about 9.8 video channels of capacity.[29] The Hulu user would get to watch the clip. But, he or she would either have to wait half an hour to watch a six-minute clip with interruptions or have to accept pauses in viewing while the programming trickled into the buffer. Applications such as BitTorrent can also fill network buffers and thereby delay other applications and other users.

BitTorrent does not dispute this latter fact. About a year ago, a BitTorrent position paper explained,

> When a user starts a typical implementation of BitTorrent today, multiple uploading TCP connections entirely saturate the uplink and fill the buffer in the bottleneck device, typically cable or DSL modem. This imposes an additional delay on all traffic, equal to the size of this buffer divided by the uplink bitrate. In typical home usage cases, this additional delay can range from a second to four seconds or so. An increase in RTT of this magnitude not only starves out other TCP connections, **it quickly makes real-time communication, such as VoIP and games, entirely impossible**.[30]

BitTorrent is aware of the problems created by its protocol and is working to develop, deploy, and standardize a protocol that can coexist more peacefully with VoIP and interactive gaming.[31] But, even if BitTorrent does fix its protocol to be more friendly to other applications, ISPs will always have to deal with new software and new problems. Denying ISPs tools to deal with disruptive or unfair software will harm consumers.

One of the factors that permits the public Internet to work is that most software follows

---

[29] On June 30, 2009, I used packet capture tools to verify that Hulu uses a single TCP connection to transfer a video clip.

[30] "Users want P2P, we make it work," S. Shalunov, BitTorrent position paper IETF P2P Infrastructure Workshop, Boston, May 28, 2008. Emphasis added.

[31] See the charter of the Low Extra Delay Background Transport (LEDBAT) working group of the IETF's Transport Area. The group is cochaired by a BitTorrent employee, and BitTorrent has contributed in other ways to the working group's operation. http://tools.ietf.org/wg/ledbat/.

the honor system for congestion control. However, if ISPs lack the ability to manage traffic that is not obeying the honor system and to use approaches that make their networks "smarter," then they may be unable in the future to keep their network running—at least at a level that satisfies consumers' expectations and needs—if widespread violations of the honor system proliferate.

More Recent System Collapses

Concern about congestion collapse in today's Internet is not theoretical. On December 26, 2006, a large earthquake took down 12 of the 18 cables between Taiwan and the Philippines. Internet service in much of Asia was seriously impaired. Bob Briscoe reports that an ISP in Singapore, SingNet, restored service before the cables were repaired by blocking video downloads and gaming traffic.[32] That is, by the simple expedient of giving email, VoIP, and normal web-browsing priority over video downloads and gaming, SingNet was able to restore Internet service to most users.

In this case, network overload was precipitated by a massive hardware failure. But, network overload can arise from many other factors. Flawed hardware can create overloads as can malicious or faulty software. Automated access to Network Time Protocol (NTP) servers has been the source of several localized network overloads. The NTP provides the Internet's equivalent of a clock-on-the-wall. Any computer on the net can query an NTP server and find out the current time. Operating systems and network hardware often have NTP clients built in. These built-in clients permit the equipment to set the time automatically without any operator intervention. For example, once a week the time-of-day clock on my computer asks the NTP server at *time.windows.com* to provide the correct time.

There have been several incidents in which such NTP client software went awry and overloaded some facilities. Perhaps the most well-known occurred in May 2003 when the University of Wisconsin NTP server was flooded with hundreds of megabits per

---

[32] Internet—Draft, draft-briscoe-tsvwg-relax-fairness-00, B. Briscoe et al., November 12, 2007. Cable failures in the Mediterranean in January 2008 also precipitated Internet failures. See "Disaster's Impact on Internet Performance – Case Study," Tomasz Bilski, in A. Kwiecień, P. Gaj, and P. Stera (Eds.): CN 2009, CCIS 39, pp. 210–217, 2009, Springer-Verlag.

second of NTP traffic.[33] The cause of this traffic was a router manufactured by Netgear that was hard-coded to query the university's NTP server. That code in the router queried the NTP server once per second until it received an answer. If the Netgear router was located behind a firewall that blocked incoming UDP packets, then the router would send one query per second continuously. Plonka reports that Netgear had manufactured about 700,000 of the affected products. If all of these were operating in the defective mode, they would send about 426 megabits per second of traffic towards the University of Wisconsin.[34]

Perhaps a greater threat is posed by widely used software that automatically downloads and installs software updates. Consider what would happen to the Internet if faulty code in a Microsoft Windows update made available in early July and downloaded millions of times since than had the property that, beginning in the month of August, it would query the timeserver once a second. At midnight July 31, there would be a sudden flood of queries to the time server—a flood that would grow as midnight rolled across the globe. If we assume, conservatively, that only 10 million Windows machines would have installed the software update and would be connected to the Internet, they would generate a flow of about 6 gigabytes per second toward the time.windows.com time server.[35] This sudden flow might disrupt parts of the network.[36] And, if many more copies of the software had been installed before the error surfaced, say it was installed on 100 million machines, then the disruption might be widespread.

---

[33] See the account by Dave Plonka at http://pages.cs.wisc.edu/~plonka/netgear-sntp/. See also the notice on the Netgear website. http://kb.netgear.com/app/answers/detail/a_id/1112

[34] Netgear was not the only firm to make such defective equipment. The Wikipedia article on NTP server misuse details several others. For an account of one of these see Richard Clayton's article, "When firmware attacks! (DDoS by D-Link)," April 7, 2006, available at http://www.lightbluetouchpaper.org/2006/04/07/when-firmware-attacks-ddos-by-d-link/

[35] Microsoft has its own large network that is interconnected with that of many ISPs at various locations. Consequently, the attack I describe might only cause problems on Microsoft's internal network rather than on the public Internet. I chose Microsoft Windows to illustrate this threat because most people are aware of how pervasive Windows is in the computing environment. However, many other software packages automatically download and install updates and impose similar risks.

[36] It may seem unreasonable to posit such a programming error. However, the list of programming errors that caused massive losses is extensive. For example, CNN reported that in 2007 a flight of Air Force F-22's lost their navigation and communication systems as they flew across the International Date Line. See http://transcripts.cnn.com/TRANSCRIPTS/0702/24/tww.01.html. Navigation and communications systems support safety-of-life and are critical to the mission of these fighters so one would expect that the software in these systems is subject to substantial testing and quality verification. Yet this critical software failed as the aircraft passed across the International Date Line.

Brett Glass operates a wireless ISP named Lariat in Laramie, Wyoming. In May of this year, his network was brought to its knees by his customer's Windows machines. The customer machines were all automatically downloading a large security update to Windows. He restored normal service by managing the traffic triggered by the Microsoft update in order to ensure that it did not overwhelm the network.[37]

In addition to incompetent software, there is also the threat of malicious code. *Botnets,* networks of user computers that have been infected with software that permits the use of those computers by operators of the network, are often used to create distributed-denial-of-service attacks.[38] In April 2007, there was what appeared to be an attack on the Internet in Estonia resulting in substantial disruption of Internet service in Estonia.[39]

Most recently, on July 4, 2009, a wave of denial-of-service attacks hit federal government computer facilities and a few commercial computers in the United States. Some computers in South Korea were also attacked. The web server for the Department of Transportation appears to have been out of service for four days.[40] One can also imagine malicious code being embedded in widely used software and being used in a similar fashion to flood networks.

Summing Up

As the above discussion illustrates, the threat of a congestion failure on the Internet is real. Congestion failures of various magnitudes occur in parts of the Internet today, as the Estonia, SingNet, Lariat and recent attacks of U.S. government computers all demonstrate. Congestion failure can be caused by hardware failures, software that fails to follow the Internet honor system, incompetently designed hardware and software, and

---

[37] See http://www.interesting-people.org/archives/interesting-people/200905/msg00021.html. Notice that he restored service by throttling legitimate Internet traffic. The Windows security update was valuable and having user machines automatically download and install it served efficiency. However, having them all download it at the same time over Lariat's relatively small middle-mile connection to the larger Internet did not serve efficiency.

[38] The term botnet is derived from ro*bot net*work. See the wikipedia entry on botnets. In 2007, Google's Vint Cerf estimated that one-sixth to one-quarter of the computers on the Internet had been subverted by botnet operators. See "Criminals 'may overwhelm the web'", Tim Weber, BBC, 25 January 2007. Available at http://news.bbc.co.uk/2/hi/business/6298641.stm.

[39] See "Hackers Take Down the Most Wired Country in Europe," *Wired Magazine,* August, 2007.

[40] See "Federal Web sites knocked out by cyber attack," L. C. Baldor, *Washington Post,* July 8, 2009. Several articles indicated that the attacks were triggered by the government of North Korea.

malicious actors.

A well-accepted and essential tool in fighting these failures is the ability of ISPs to differentiate among different types of traffic, including by directly managing the threat caused by particular harmful traffic. If SingNet had been unable to block file sharing applications, it would have taken days or weeks before basic Internet services were functioning properly again. If Bret Glass had been unable to address the Microsoft downloads that were causing the problems, the users on his network would have had to endure poor service. A technology called *deep packet inspection* is one of the tools that ISPs can use to identify and manage the traffic that is disrupting network performance. Priority routing, tools such as deep packet inspection, and ISPs that are permitted to be flexible and agile are important factors well-accepted by network engineers for their role in averting and resolving congestion failures.

Use of Established Congestion-Avoidance Technologies

The concept of priority traffic is not new to the 21st century. Networking researchers experimented with voice over packet networks as early as the mid-1970s.[41] It was immediately clear to these researchers that it would make sense in many situations to give voice priority over applications such as file transfer. And, from the very first days of TCP/IP the Internet standards community adopted standards supporting such priority routing. To date, there have been multiple Internet standards established that can be used to provide priority routing of packets. These include type-of-service, DiffServ, IntServ/RSVP, and MPLS.[42] For a variety of reasons, the first three of these approaches have not been extensively adopted in the Internet. However, the fourth approach, MPLS,

---

[41] I clearly recall attending a demonstration of voice over the ARPANET in the seventies done by, as I recall, Bob Kahn and others. (The voice did not sound very good.)

[42] Type-of-service was an option in the original IP standard (RFC 760, January 1980) which had a 3-bit field for priority. This was modified slightly by RFC 791. Later RFCs (e.g., RFC 1349, RFC 2474, RFC 3260) provided substantial modifications to the priority mechanism creating a new approach to priority that was called differentiated services or Diffserv. RFC 2205 defined the Resource ReSerVation Protocol (RSVP). RSVP permits the reservation of resources, such as bandwidth and queue capacity in routers, along the path between two computers on the Internet. RSVP permits reserving capacity for a communications process, such as a VoIP connection, before the process begins. Such a reservation assures that the communications process will not suffer from congestion when it is active. MPLS, described in RFC 3031, can be regarded as a cross between ATM and TCP/IP—a hybrid that has advantages over either of its parents. MPLS permits network operators employ a wide range of quality-of-service and traffic engineering techniques. RFC 4094 offers a survey of some of these quality-of-service technologies.

is widely used. For example, Level 3 operates a converged MPLS core network. Level 3's public Internet and private virtual network traffic travels on the same core network, with private network traffic being given assured performance levels.[43] Any rule that requires all packets to be treated the same would probably outlaw the use of long-established approaches like Diffserv, IntServ, and RSVP. It might also threaten the efficient and beneficial separation of traffic into various priority classes on MPLS networks—a common and efficient practice benefitting consumers today.

Technology does not stand still. There are multiple research efforts to find better ways to provide priority service or assured quality-of-service over the Internet. A December 2008 presentation by Dr. Tim Gibson of the Defense Advanced Projects Research Agency (DARPA) described the performance of a new router developed by HP and Anagran with funding from DARPA. Energy efficiency was improved by a factor of 4 and throughput under conditions unfavorable to TCP was improved by a factor of 40.[44] Intimately tied to the efficiency gains of the new router are priority mechanisms that give some flows priority over others or can completely exclude flows that would overload the network. The IETF's NSIS working group is also working on improved quality-of-service over the Internet.[45]

Priority-enforcing technologies offer the opportunity to combine all communications on a single broadband link to the Internet.[46] In contrast, any prohibition on priority routing would steer traffic away from smaller service providers that operate only one network. For example, a hospital cannot use the Internet for latency-sensitive traffic, perhaps a medical monitoring service, if it must live with the threat that another user's rogue

---

[43] See http://www.level3.com/index.cfm?pageID=54.

[44] See Building Authenticated and Responsive Networks that are Faster and More Efficient, T. Gibson, 18 December 2008. Available at http://www.darpa.mil/STO/solicitations/baa09-11/pdf/Proposers_Day.pdf A more detailed description of this research is given in "The CHART System: A High-Performance, Fair Transport Architecture Based on Explicit-Rate Signaling," J. Bassil et al., HP Labs, undated. Available at http://www.hpl.hp.com/news/2009/jan-mar/pdf/brassil_osr_crc_21.pdf.

[45] See http://www.ietf.org/html.charters/nsis-charter.html.

[46] Larry Roberts, one of the true pioneers of the Internet, described the benefits from improved routing in a seminar at Stanford last year saying, "Recent improvement in flow technology which maintains information for each active flow, insures quality voice/video, allows utilization in the 95% region, and maintains unprecedented fairness." See http://netseminar.stanford.edu/seminars/01_17_08.html and http://netseminar.stanford.edu/seminars/01_17_08.ppt.

application can seriously degrade or cut off service.[47] Rather, a hospital would need to purchase dedicated connections from a provider able to provide such service on a network separate from the public Internet. Any form of network regulation that prohibits priority routing or other approaches to assuring service quality would make it necessary for our nation to have multiple networks for voice, high-priority data, and general Internet data. The requirement to connect to and use multiple networks may not be a significant burden for a large corporation in an office building in Manhattan. Fiber runs to the basement of the building, and the organization has sufficient scale to operate three networks efficiently. Smaller organizations, however, would face proportionately larger costs to manage the multiple networks and pay the various fixed costs. The development of applications that require high-quality network service would be handicapped as such applications would perform better on dedicated networks than over the public Internet.[48] Aggressive but delay-tolerant applications will thrive and latency-sensitive applications will stumble along. In such cases, regulation and the physics of networks rather than consumer preferences will determine which firms and applications succeed in the market.

Cross-Layer Design and Cross-Layer Quality-of-Service

Cross-layer design refers to the design of network elements, such as wireless access links, that take into account information from other layers to optimize performance. Cross-layer design gets its benefits at the cost of avoiding the simplifications created by the layering principal. Often this results in explicitly distinguishing between packets – something that some network regulation proposals would prohibit.

An example illustrates how cross-layer design can aid efficiency. Consider a radio link carrying two streams of traffic to and from the Internet. One stream is VoIP; the other is a TCP transfer of a web page. VoIP traffic can tolerate little delay, but an occasional packet can be lost without significant harm to the conversation.[49] The web page transfer

---

[47] Recall that the BitTorrent whitepaper said that BitTorrent software does exactly this at times. See footnote 30.

[48] Recall the discussion of Clearwire above in which I showed that if their system were engineered to maximize capacity for applications such as web browsing, it would be unlikely to support adequate VoIP service during busy periods.

[49] Typically, about 1/50 of a second of voice is encoded in a single packet; a packet carries only part of a single syllable.

is more tolerant of delay, but if a packet is lost, the TCP software will retransmit it until proper reception occurs.

Because radio links have much higher error rates than wired LANs, it is common for radio links to include error-detecting and error-correcting capabilities at the link level. Suppose a packet is transmitted over the radio link and is found at the receiver to have arrived in error. The receiver can request partial retransmission of that packet using a technology called Hybrid-ARQ. In Hybrid-ARQ retransmission, the transmitter sends information, such as additional error-correcting coding, that supplements the original transmission rather than retransmitting the entire packet.

In this situation, if the receiving system detects that a packet has become corrupted on the radio link, the efficient action for the receiving system depends on the type of packet that was received in error. If the packet is part of the TCP stream, then the receiving system should request link-level retransmission. A Hybrid-ARQ retransmission uses significantly less of the resources of the radio system than does a retransmission at the TCP level. In contrast, the receiving system should probably discard the VoIP packet that was received in error. Retransmitting the VoIP packet would add delay to the voice stream without any corresponding increase in the quality of the voice connection. Such a "nonneutral" link increases efficiency and improves customer's Internet experience without any harmful effects.[50] Consumers get more for their dollars.

Somewhat related to cross-layer design is the use of cross-layer processing to improve service quality. Several manufacturers offer Ethernet switches that inspect Ethernet frames and route those frames taking into account level 3 or level 4 protocol information. Cisco touts the capability of their ESW 500 series of switches for small business to give VoIP priority saying, "QoS level assures that voice-over-IP (VoIP) traffic takes precedence."[51]

An analogous service could be provided in the public Internet. For example, with deep

---

[50] This example is illustrative. Wireless networks contain a subsystem, called the *scheduler* that manages transmissions. The exact algorithms used by the schedulers in various systems are proprietary to the manufacturers.

[51] See http://cisco.com/en/US/prod/collateral/switches/ps5718/ps10143/data_sheet_c78-521740.html

packet inspection, a carrier could examine packets to see if they represented an attempt to set up a voice call to 911 and give that call-setup attempt priority in the network. A sufficiently smart network would also be able to give priority to voice traffic to and from 911.[52]

Proposals that ISPs only provide "dumb pipes"—pipes that are not smart enough to choose the most efficient retransmission and routing policies—would eliminate such potentially useful practices. Worse yet, they would stifle innovation in the development and use of such practices.

Security

Adoption of the proposals mandating undifferentiated treatment of packets could also make broadband networks and services less secure and less able to defend against a variety of threats.[53] Above I described the denial-of-service attacks against Estonia and against U.S. and Korean computers. Brett Glass humorously describes the congestion caused by the automatic Windows updates as "An unusual denial-of-service attack."[54] The same tools that can limit inadvertent causes of congestion can be used to prevent and address malicious congestion.

Packet inspection or deep packet inspection provides one potentially significant tool for increasing security. Cisco sells a pair of products, the Traffic Anomaly Detector and the Anomaly Guard Module, that are designed to detect distributed denial-of-service attacks and to mitigate their harms.[55] Cisco describes the functioning of the system saying,

> When the Cisco Traffic Anomaly Detector Module identifies a potential attack, it alerts the Cisco Anomaly Guard Module to begin dynamic diversion, which redirects traffic destined for the targeted resources-and only that traffic-for inspection and scrubbing. All other traffic continues to flow directly to its

---

[52] For example, the network could note the SIP messages from a user attempting to set up a call to 911 and could give priority to all telephony traffic from that user.

[53] Many of the various proposals for network neutrality have language that appears to exempt security practices. However, if a policy reduces the incentive to invest in equipment that controls congestion and that can also be used to provide security capabilities, networks will have less investment in security capabilities. Also, the definition of security is unclear.

[54] Glass, op. cit.

[55] See http://www.cisco.com/en/US/prod/collateral/modules/ps2706/ps6235/product_data_sheet0900aecd80220a7c.html

> intended destination, delivering a low-impact, highly reliable, and economical solution that offers easy installation.
>
> Diverted traffic is rerouted through the Cisco Anomaly Guard Module, where it is subjected to multiple layers of scrutiny to identify and separate "bad" flows from legitimate transactions. Specific attack packets are identified and removed, while "good" traffic is forwarded to its original destination, helping to ensure that real users and real transactions always get through, and providing maximum availability.[56]

Some denial-of-service traffic could be detected by deep packet inspection but not by inspection of just the headers. The ability to inspect packets also would provide an effective tool to detect and divert spam and emails that carry computer viruses and other malware. Packet inspection could also detect some malware that is attempting to propagate itself over the Internet.

The threat from malware is real. The National Science Foundation and the U.S. Army funded an analysis of the Conficker virus by SRI International. SRI made clear the magnitude of the threat:

> Perhaps the most obvious frightening aspect of Conficker C is its clear potential to do harm. Among the long history of malware epidemics, very few can claim sustained worldwide infiltration of multiple millions of infected drones. Perhaps in the best case, Conficker may be used as a sustained and profitable platform for massive Internet fraud and theft. **In the worst case, Conficker could be turned into a powerful offensive weapon for performing concerted information warfare attacks that could disrupt not just countries, but the Internet itself**.[57]

Blocking some packets – those that are harmful to users or to broadband networks – serves security. A test of my Comcast cable modem service reveals that Comcast blocks incoming traffic to TCP ports 135, 139, and 445. Each of these ports is commonly used for a service on the local network—not on the larger Internet.[58] The United States-Computer Emergency Response Team (US-CERT), an activity of the Department of Homeland Security, recommends blocking traffic to and from these ports in order to

---

[56] Data Sheet for the Cisco Traffic Anomaly Detector XT 5600. Available at http://www.cisco.com/en/US/prod/collateral/vpndevc/ps5879/ps6264/ps5887/product_data_sheet0900aecd800fa552.html

[57] "An Analysis of Conficker's Logic And Rendezvous Points," Phillip Porras et al., *SRI International Technical Report,* 19 March 2009 in the Addendum "Conficker C Analysis" dated 4 April 2009. See also "Computer Experts Unite to Hunt Worm" John Markoff, *New York Times,* March 18, 2009.

[58] The services are RPC, NetBIOS, and SMB.

protect against various attacks.[59] Many home computer users lack the knowledge and skills to do such blocking. Consequently, consumers benefit from Comcast's decision to block traffic to these ports and also benefit from Comcast's ability to block traffic to any other port should that port become a security vulnerability. Many ISPs block TCP access to port 25 as compromised user machines send email spam using connections to port 25.[60]

Conclusions

Although some commenters opposed any form of congestion control or priority routing mechanism that would favor one class of packets over another or otherwise differentiate between packets, they failed to identify or discuss the many costs that would flow from adopting such a policy.

There is no simple rule that can identify when priority routing should be applied or which flows it should be applied to. Above in my discussions of priority in wireless and of cross-layer design, I gave examples of well-accepted practices that give preferential processing to one category of packet over another, effectively expanding capacity and improving efficiency in the use of a limited resource. A discussed above, a careful analysis of the nature of the application and of the higher-level protocols permits doing more with the limited resources of broadband networks.

Likewise, consistent with widely accepted practices, differentiation among packets can combat the real problem of congestion. Congestion was a severe problem in the Internet in the mid-1980s. The solution to that congestion was the adoption of improved versions of TCP that incorporated congestion control. Unfortunately, this is congestion control on the honor system. Some current web browsers and peer-to-peer applications bend or break the honor system—permitting them to deliver better service to their user but at the

---

[59] Several CERT Vulnerability Notes recommend blocking some or all of these ports. See, for example, US-CERT Vulnerability Note VU#827627, October, 2008.

[60] In May, 2005 the Industry Canada Task Force on Spam's report recommended best practices for ISPs to fight spam. These "best practices" include blocking port 25. They explained, "Port 25 has been widely abused by spammers running zombie networks (or "botnets"). By monitoring and limiting the use of port 25, ISPs and other network operators can close off a major avenue for spamming. Canadian ISPs that have already implemented port 25 blocking have seen very significant declines in the amounts of spam originating on their networks." See their report at http://www.ic.gc.ca/eic/site/ecic-ceac.nsf/eng/h_gv00317.html.

expense of more congestion for other users. No simple rule regarding priority for one class of packets can encompass this complexity.

Congestion can also arise from network equipment failures, software features, and malicious software. I described four recent incidents of such congestion failures. I suspect that there were many more that went unpublicized.[61] In three of these examples, the ability of networks to manage congestion-causing traffic permitted most uses of the network to continue in a close to normal fashion.[62] Consumers benefit if networks have these capacities during times of congestion whether that congestion is caused by normal patterns of use, hardware failures, software failures, or malicious software.

Although this paper has focused on technical issues—such as how priority scheduling expands wireless capacity or how packet inspection limits denial of service attacks—one should remember that there is also an economic argument for priority. Just as it makes sense to give an ambulance priority over commuters' cars, it makes sense to give packets carrying VoIP calls to 911 priority over packets carrying music downloads.

Although some commenters urged the Commission to prohibit service providers from distinguishing between packets, or ever favoring one packet over another, their analysis was silent on the many costs and unintended consequences that this policy would impose. Indeed, some essentially argued that it would impose no costs. But, as the above discussion shows, it is difficult to conceive that an informed engineer or economist would consider priority scheduling of packets to be a zero-sum game. Today, ISPs, wireless carriers, and private networks use a variety of technologies to defend networks against malicious traffic and to give priority to traffic that is sensitive to delay or jitter. Prohibiting or restricting such technologies would harm consumers and pose risks to the economy and to public safety. Perhaps worst of all it would hamper innovation and create artificial incentives to have multiple, fragmented networks.

---

[61] See the anomaly case studies list at SLAC for a few examples. https://confluence.slac.stanford.edu/display/IEPM/Anomaly+Case+Studies

[62] I have not seen any account of the countermeasures used for the recent Fourth-of-July attacks.

About the Author

Dr. Charles L. Jackson is an electrical engineer who has worked extensively in communications and wireless. He has been both a digital designer and a system programmer. He works as a consultant and as an adjunct professor at George Washington University, where he has taught graduate courses on computer security, networking and the Internet, mobile communications, and wireless networks. Dr. Jackson consults on technology issues—primarily wireless and telecommunications. Dr. Jackson served three terms on the FCC's Technological Advisory Council. He previously worked at both the FCC and the House Commerce Committee. He holds two U.S. patents. Dr. Jackson received his PhD from MIT.

Although Dr. Jackson was not in any way a contributor to the design and development of the Internet, he has had a next-bench view of much of its development and growth. He designed and built some of the hardware on the SRI computer that was part of the first ARPANET connection. He has been a user of the ARPANET/Internet since the early 1970s. He did participate, along with several people who were real networking pioneers, in three projects in the 1970s and 1980s involving the use of email for the deaf and hearing impaired.