Managing Broadband Networks: A Policymaker's Guide

George Ou

The Information Technology and Innovation Foundation

MANAGING BROADBAND NETWORKS: A POLICYMAKER'S GUIDE

George Ou

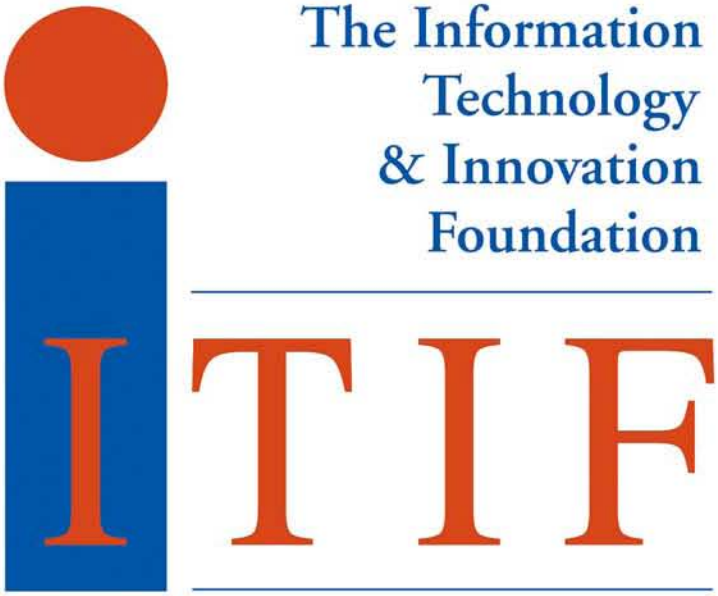

December 2008

Table of Contents

I. Executive Summary 1

II. Managing Broadband Networks: A Policymaker's Guide 6

Introduction 6

The Debate About Net Neutrality 7

The Evolution of Networks and Network Management 8

A. Circuit-Switching Networks Used in the Telephone System 8

B. Packet-Switching Networks such as the Internet 10

C. The need for Quality of Service (QoS) on Packet-Switching Networks 10

Static vs. Dynamic Rationing of Network Capacity 12

Jacobson's Algorithm: A TCP Congestion-Control Mechanism Built into the Internet 13

Network File Distribution Architectures 15

A. Client-Server File Distribution Architecture 15

B. Peer-to-Peer (P2P) File Distribution Architecture 15

C. Content Delivery Network (CDN) File Distribution Architecture 17

Improving Fairness Between Broadband Customers 18

A. Internet Standards 19

B. Protocol-Specific Throttling Systems 19

C. Protocol-Agnostic Network Management Systems 20

Protocol Specific Network Management Solutions 22

Logical Order of Packet Priority for Application Types on the Internet 23

The Problem of Jitter on Packet-Switching Networks 25

A. Queuing Theory's Application to Packet-Switching Networks 25

B. The Misperception that Network Jitter Can Be Solved by More Capacity 27

C. Why Broadband Networks Will Always Have Speed Mismatches and Jitter 27

D. Why Certain Applications Create High Jitter and Others Don't 29

Quality of Service (QoS) and the Internet **31**

A. Solving the Jitter Problem with QoS 33

B. Clearing up Misconceptions About QoS 33

C. QoS for Broadband Networks 34

D. User-Approved and User-Controlled QoS 35

Wireless Networks: The New Frontier of the Internet **35**

A. Why Wireless Networks Require More Management than Wired Networks 35

B. Increasing Spectral Efficiency Through Scheduled Access 36

C. Why Wireless Management Is a Necessity That Enables Innovation 37

Flawed Arguments About Alternatives to Intelligent Network Management **38**

A. Why Increasing the Supply of Bandwidth Will Not Solve the Problem 38

B. Why Metered Pricing and Usage Caps Alone Will Not Solve the Problem 40

C. Why Exclusive QoS on the Internet Is Better Than Exclusive QoS on Private Circuits 41

Conclusion **43**

Endnotes **44**

Appendix A: Networking Glossary **A1**

List of Boxes

Box 1: Net Neutrality Proposals Under Consideration 9

Box 2: Overview of the Internet 11

Box 3: Jacobson's Algorithm for Avoiding Network Congestion at Work 14

Box 4: Operational Modes in P2P Networks 16

Box 5: Basic Concepts of Network Performance 26

Box 6: Debunking the Myth That Network Capacity Is a Substitute for Quality of Service (QoS) 39

List of Tables

Table 1: Network Requirements of the Four Basic Types of Applications That Run on the Internet 24

Table 2: Emerging Network Technologies 37

List of Figures

Figure 1: Exploiting TCP congestion control....13
Figure 2: Client -Server model....15
Figure 3: Peer-to-peer (P2P) model....16
Figure 4: Content Delivery Network (CDN) model....18
Figure 5: Unmanaged versus managed bandwidth allocation....21
Figure 6: Dumb versus Smart network....25
Figure 7: Why there will always be a bottleneck on broadband....27
Figure 8: Effect of low upstream BitTorrent usage on jitter....28
Figure 9: Effect of low upstream and downstream VoIP usage on jitter....28
Figure 10: Effect of high downstream BitTorrent usage on jitter....28
Figure 11: High jitter inducing application....29
Figure 12: Low jitter inducing application....30
Figure 13: How VoIP packets flow....31
Figure 14: VoIP dealing with low jitter still suffers....32
Figure 15: High jitter is much more destructive to VoIP....32
Figure 16: Network device with QoS....33
Figure 17: QoS can completely mitigate jitter damage....34
Figure 18: The need for QoS on both ends of the broadband competition....35
Figure 19: Wireless networks — the new frontier of the Internet....36
Figure 20: Why capacity isn't a substitute for network management: the case of Japan....40

Executive Summary

To make progress to a ubiquitous digital world, bigger pipes are not sufficient. We need not just expanded network capacity but networks that are better and more intelligently managed.

The Internet has changed the face of communications, commerce, and indeed the world. And over time the Internet itself has changed too. Until recently, most Americans at home accessed the Internet using telephone dial-up connections rather than today's faster broadband connections. With slower connections, home users limited themselves to a few basic online activities, such as email and web browsing, which perform passably well even on a slow network. In this environment, the need for Internet service providers (ISPs) to manage their networks to ensure the best possible experience for their customers was limited.

Today most Americans connect to the Internet over broadband connections that are in some cases 400 times faster than the dial-up connections of the late 1990s. But it is precisely because of these new bigger "pipes" that ISPs are finding that they need to more actively manage their networks. Broadband networks have enabled the rise of new applications, including those that need to be managed if they are to work effectively (e.g., voice over Internet Protocol, online gaming, video conferencing, and Internet Protocol-based TV) and those that can cause other applications to fail on an unmanaged network (e.g., many peer-to-peer (P2P) applications).

With this exciting transformation of the Internet into the universal communication platform of the future, network engineers face an array of daunting challenges. Specifically, to provide customers a good Internet service and operate their networks efficiently, ISPs must be able to do two very important things: 1) allocate limited bandwidth fairly among users; and, 2) apply network management tools to shape traffic from multiple applications. ISPs can and should do these things in a fair and nondiscriminatory manner. Thus, they should strive to ensure that customers who pay for the same tier of service get roughly the same bandwidth at a given level of usage, eliminate harmful variations of delay (i.e. jitter), make consumers' broadband service more conducive to using multiple applications simultaneously, while at the same time treating other applications and content fairly.

Unfortunately, network management solutions have come under heavy criticism from many advocates of "net neutrality." The issue of network management came to the fore when Comcast limited the ability of peer-to-peer (P2P) users to operate in upload-only mode whenever P2P traffic exceeded 50 percent of total upstream capacity of the entire

neighborhood. More generally, the issue of network management refers to whether and to what extent ISPs can manage their networks to ensure quality of service for the majority of their customers.

Strong advocates of net neutrality argue that ISPs should have little flexibility to manage their networks and that the solution to any kinds of network congestion or other network performance challenges can and should be solved by simply adding more network capacity—primarily in the form of "bigger pipes." Indeed, they fear that using efficient network management techniques may enable network operators to abuse their power, thereby stifling free speech and civic expression and erecting unfair barriers to other companies seeking to distribute digital content or applications. Moreover, some proponents of net neutrality fear that any improvement in the efficiency of the Internet will eliminate the motivation of ISPs to expand network capacity by "building bigger pipes." As we transition to a ubiquitous digital world, bigger pipes are necessary – and public policy should support their deployment – but they are not a substitute for network management. We need to not just expand network capacity, but also build networks that are better and more intelligently managed.

Many if not most of the fears of the proponents of net neutrality stem from a lack of understanding of the history of the Internet, the economics of the ISP industry, and the science of network engineering. This guide is intended to help policymakers better understand how broadband networks and the applications that run on them work, and calls for a balanced approach to the regulation of broadband network management. A balanced approach should be based on reality: both the economic realities of building broadband networks and the scientific realities of network engineering. In addition, it should provide ISPs the flexibility they need to manage complex networks while also ensuring oversight to insure that network management practices are not being applied in anti-competitive ways.

Effective policy in this area must be based on facts. Unfortunately much of the debate over broadband network management to date has been informed more by rhetoric and emotion than by an actual examination of how advanced networks and the applications that run on them work. By providing policymakers with this guide, ITIF hopes to better inform this debate.

Key Findings and Conclusions:

- **Packet-switched networks, like the Internet, have advantages, but also disadvantages.** Packet-switched networks like the Internet were invented for their flexibility and efficiency, characteristics which are optimum for data applications. But they have two key deficiencies in the absence of network management: 1) inability to equitably allocate bandwidth; and 2) high jitter, which are essentially micro-congestion storms that last tens or hundreds of milliseconds, and which can disrupt real-time applications such as VoIP, online gaming, video conferencing, and IPTV.

- **The Internet and its predecessor ARPANET became the first adopter of packet-switching networks because it was more efficient and flexible than the circuit-switching telephone network.** Unlike telephone networks which only connected a small percentage of users at any given time, packet-switched networks allow everyone to be on the network at the same time and dynamically divide up the resources among the active users. If few users are on the network, then those users get a lot of resources allocated to them. If many users are on the network, then each user gets fewer resources but no user is locked out. This dynamic expansion and contraction of bandwidth makes packet switching networks very efficient but the allocation of bandwidth can become disproportionate whenever applications like P2P resist reallocations of bandwidth. Network management can balance the allocation of bandwidth such that each customer in the same service tier gets an equitable share of the total bandwidth.

- **Network management techniques, such as quality of service (QoS) mechanisms, make a packet-switched network more conducive to simultaneous application usage.** Network management tech-

niques such as QoS essentially carve out virtual circuits within a packet-switched network by providing the necessary resources and performance characteristics that real-time applications need. This gives a packet-switched network the real-time characteristics of a circuit-switched network while maintaining the robustness and flexibility of a packet-switched network.

- **Even since its early days, the Internet has been a managed network.** The Internet has had basic network management mechanisms built into it since its inception, although these mechanisms have undergone and continue to undergo much refinement as usage patterns on the Internet change. Since 1987, for example, computers have used a revised version of the transmission control protocol (TCP) that includes a network congestion control mechanism developed by computer scientist Van Jacobson to slow down endpoints and prevent network meltdown.

- **Peer-to-peer (P2P) applications pose special challenges to broadband networks.** P2P users on unmanaged networks can use a disproportionately high amount of bandwidth and cause network congestion. In Japan, for example, P2P users represent 10 percent of the total broadband population but account for 65 to 90 percent of traffic on the network. By running multiple TCP flows (i.e. connections) per file transfer, P2P applications can effectively circumvent the Jacobson algorithm intended to allocate bandwidth. As a result, P2P applications can maximize the use of available bandwidth, sometimes at the expense of other applications, such as VoIP and video conferencing, which require low latency and jitter.

- **An ISP that dynamically allocates its network capacity can always offer its customers far more unguaranteed bandwidth than its guaranteed minimum level of service.** Because broadband networks are shared, it is more efficient to give consumers access to speeds that can increase when there is less congestion. Since only 1 to 10 percent of network users are active at any point in time, packet switching networks can dynamically allocate 10 to 100 times more bandwidth to each active user. If a network can be built to guarantee 1 megabit per second (Mbps) of performance for each user, for example, it can just as easily offer the customer 1 Mbps of guaranteed performance and up to 20 Mbps of unguaranteed performance. But building a network that provided a guaranteed performance of 20Mbps for example, would be much more expensive and require much higher monthly costs for the consumer.

- **One goal of network management is to fairly allocate bandwidth between paying customers.** Fairness dictates that customers who are paying for the same tier of broadband service from a broadband provider should get roughly the same bandwidth at a given level of usage. Fair bandwidth allocation shouldn't just measure instantaneous bandwidth usage, duration should also be factored in to the equitable distribution of bandwidth. If one application or one customer uses the network hundreds or thousands of times more frequently than another application or customer, it isn't unreasonable to let the short duration application or customer get a short boost in bandwidth over the long duration application or customer.

- **To achieve fair bandwidth allocations, protocol-agnostic schemes are the best solution.** ISPs can use protocol-agnostic network management systems (systems that measure the aggregate bandwidth consumption of each customer and not what protocols they are using) to ensure that bandwidth is shared fairly between customers. Early network management systems that used less accurate protocol-specific schemes to allocate bandwidth between customers worked well most of the time but experienced occasional problems. These protocol agnostic solutions are being evaluated by broadband providers. A key downside of protocol-agnostic network management systems is that they are often too expensive for smaller ISPs to deploy.

- **Another goal of network management is to better share network resources between many different applications.** Different types of applications have different network requirements. Real-time applications

(e.g., VoIP) are most sensitive to network jitter. Video streaming applications (e.g., YouTube) have moderate fixed bandwidth requirements and moderate jitter tolerance. Interactive applications (e.g., web browsing) have brief bursts in bandwidth that could disrupt real-time or streaming applications. Background applications (e.g., P2P applications) are designed to be unattended with no one waiting for an instant response.

- **Packets should be ordered logically with priority given to real-time applications first, streaming applications second, interactive applications third, and background applications last.** In order for all applications efficiently and fairly share an Internet connection, those with higher duration and higher bandwidth consumption (e.g., P2P) are given lower priority than applications with lower duration and lower bandwidth consumption (e.g., VoIP applications). This does not mean P2P applications are being mistreated because they still receive the highest average bandwidth from the network.

- **To better enable multiple applications to share an Internet connection, protocol-specific schemes are necessary.** Application protocols that require low packet delay must be identified and must be protected against high variations in packet delay (e.g., jitter) and Quality of Service network management techniques are the mechanism that provides that protection.

- **Wireless networks require more management than wired networks.** Wireless networks require more network management than wired networks because they have less bandwidth available and it must be shared more frequently. Furthermore, multiple radio transmitters sharing the same wireless frequency in the same geographic location results in a high probability of radio interference which can bring networks to a halt. These unique challenges of wireless networks require the most elaborate network management system of all in the form of a centralized scheduler which coordinates the transmission slots for network users as tightly and efficiently as possible without collision.

- **Wireless network management enables innovation.** Intelligent wireless networks will ultimately spur more adoption and usage of wireless broadband, which facilitates more mobile e-commerce and enables more innovation and generation of wealth.

Responding to Common Misperceptions About Network Management:

- **Network management techniques, such as QoS, do not put low priority applications on a "dirt road."** QoS gives higher prioritization to applications that have lower bandwidth, lower duration, and higher sensitivity to packet delay. In spite of this, applications that are given the least priority still end up receiving the highest average bandwidth from the network. But with this logical prioritization scheme in place, low priority applications like P2P applications interfere less with other applications sharing the same network. This in turn allows P2P applications to operate freely without any artificial constraints on when to use them or how much bandwidth to allocate to them which are commonly used on unmanaged networks.

- **Building more bandwidth, while desirable, does not eliminate the need for network management. A**dvancing the digital economy requires higher speed broadband. However, higher speed networks will not preclude the need for network management. First, as network capacity grows, network demand also grows, as new kinds of applications emerge to take advantage of the capacity. Second, networks with plenty of spare unused capacity on average can still suffer instantaneous shortages at peak times of the day. Third, networks operating at low utilization levels can still suffer packet delay in the form of jitter.

- **Metered pricing and usage caps alone will not solve the problem of network congestion.** Metered pricing and bandwidth usage caps are legitimate tools for ensuring the efficient use of networks, but they cannot control instantaneous bursts in demand nor can they deal with the problem of jitter and the inability of dumb

networks to gracefully support multiple applications. Only advanced network management techniques like quality of service can deal with these challenges.

Policy Implications:

- **Legislation and regulations should not limit efforts by ISPs to fairly use network management to overcome technical challenges and maintain a high quality Internet service for their customers.** As described in this report, ISPs face many technical challenges to manage network congestion and support various online applications. Network management is a necessary and important component of broadband networks, and policymakers should support its use. However, this freedom to manage the network is not a license for ISPs to behave in anti-competitive ways such as blocking legitimate websites or unreasonably degrading services that users have paid to access. Neither should ISPs unreasonably discriminate against any content or service on the open Internet.

- **Policymakers should be cognizant of the effects of certain proposed legislation on the use of network management.** Some proposed net neutrality bills ban differentiated pricing for enhanced QoS and would have undesirable and unintended consequences. One intent of these bills is to facilitate more open Internet bandwidth for broadband consumers, but the result may be just the opposite. Not allowing network operators to prioritize their own IPTV content above other Internet content, for example, will simply push those cable TV-like services onto private circuits that share the same physical network. That would result in less Internet bandwidth being available on a permanent basis for broadband consumers even when they are not using their IPTV service.

- **The federal government has a key role to ensure openness and fair play on the Internet.** However, it should do this with sensible rules. Policies should strive to prevent any potential abuse without eliminating the ability of ISPs to manage their networks in ways that produce the best possible user experience for the largest number of users, and without eliminating incentives to build the next generation broadband network. Toward that end the FCC should oversee broadband providers and ensure that they ISP network management practices are open, transparent and not harmful to competition. And the ISP industry should continue its efforts to develop and abide by industry codes of good conduct regarding network management that include, but are not limited to, fuller and more transparent disclosure to consumers of network management practices.

Conclusion:

The Internet in all its glory has never had a perfect architecture. There have always been conflicts between users and applications competing for scarce network resources. Network management is necessary to fairly allocate bandwidth between customers and seamlessly support multiple applications on shared network connections.

The Internet and broadband technology are continuing to evolve at a fairly rapid rate, and neither shows any signs of maturing. Network engineers continue to find new solutions to improve the Internet experience for all users. This situation makes it very difficult, if not impossible, to predict where the market and technology will evolve. The Internet is so valuable precisely because it is open to anyone, for any use, and for any business model, but participation has always required varying levels of payment for varying levels of service between willing parties. Given this environment, it is best for policymakers not to issue blanket prohibitions on network management technology and existing business models. Instead, policies should focus on creating better transparency for all Internet companies along with FCC oversight to ensure that broadband providers are managing networks in ways that are not unfair or anticompetitive.

Managing Broadband Networks: A Policymaker's Guide

To make progress to a ubiquitous digital world, bigger pipes are not sufficient. We need not just expanded network capacity but networks that are better and more intelligently managed.

The Internet has changed the face of communications, commerce, and indeed the world. What started out as an academic and military network exploded into the commercial and consumer space with the proliferation of e-mail and the World Wide Web in the 1990s. But even through the early years of the current decade, most Americans at home accessed the Internet over telephone lines using dial-up connections. In this world, the need for Internet service providers (ISPs) to manage their networks to ensure the best possible experience for their users was limited. Basic text-based applications such as e-mail and Web browsing perform passably well even on a slow network. And paradoxically because networks were so small, there were few applications that created the kinds of network congestion problems that applications such as many peer-to-peer (P2P) applications cause today.

Today most Americans connect to the Internet over broadband connections that are in some cases 400 times faster than the dial-up connections of the late 1990s. But it is precisely because of these new bigger "pipes" that ISPs are finding that they need to engage in more active steps to manage their networks. Indeed, we are in the midst of a revolution in video distribution, video communications, telemedicine, online gaming, and telephony over the Internet made possible by faster Internet connections. This has resulted not only in an exponential growth in demand for network capacity but also in the increased use of applications that need real-time communication—applications such as Voice over Internet Protocol (VOIP), online gaming, video conferencing, and Internet Protocol-based TV (IPTV) that were never meant to run on the early Internet.

With this exciting transformation of the Internet into the universal communication platform of the future, the engineers face an array of daunting challenges. Specifically, to provide customers a good Internet service and operate their networks efficiently, ISPs must be able to do two very important things. First because network capacity is inherently limited, even on much bigger pipes than exist today, ISPs must be able to allocate bandwidth among users. The most effective way to do this is to dynamically allocate bandwidth (employing "statistical multiplexing") so that each customer is able to get far more bandwidth than the guaranteed minimum.

Second, different applications have different network needs; ISPs need to be able to apply network management tools to shape traffic from multiple applications so that overall all applications work effectively. For example, some applications like VOIP need real-time capabilities, while others, like email do not. A smart and well-managed network will attempt to simultaneously satisfy different types of applications as best as possible. ISPs can and should do these things in a fair and nondiscriminatory manner. Thus, they should strive to ensure that customers of a broadband provider who are paying for the same tier of service get roughly the same bandwidth at a given level of usage, eliminate harmful variations of delay called jitter, and strive to ensure that broadband is more conducive to simultaneous application usage. However, this freedom to manage the network is not a license for the broadband provider to behave in anti-competitive ways such as the blocking of legitimate websites or the unreasonable degradation of services that users have paid to access. Broadband providers should not unreasonably discriminate against any content or services on the open Internet. The FCC should oversee and ensure that the broadband providers remain on their best behavior.

Because of these new bigger "pipes," ISPs are finding that they need to engage in more active steps to manage their networks.

Unfortunately, network management solutions have come under heavy criticism from many advocates of "net neutrality" who long for the idealized golden days of the early "dumb" Internet that, in fact, never was. Proponents of net neutrality fear that using efficient network management techniques may enable network operators to abuse their power, thereby stifling free speech and civic expression and erecting unfair barriers to new market entrants.

Moreover, some proponents of net neutrality fear that any improvement in the efficiency of the Internet will eliminate the motivation of ISPs to expand network capacity by "building bigger pipes."

Many if not most of these fears of the proponents of net neutrality stem from a lack of understanding of the history of the Internet, the economics of the ISP industry, and the science of network engineering. This policymakers' guide to how broadband networks and the applications that run on them actually work report explores these topics and calls for a balanced approach to the regulation of broadband network management. A balanced approach should be based on both economic realities and the realities of network engineering and should provide ISPs the flexibility they need to manage complex networks while also ensuring oversight to prevent any potential abuses. To make progress to a ubiquitous digital world, bigger pipes are not sufficient. We need not just expanded network capacity but networks that are better and more intelligently managed.

THE DEBATE ABOUT NET NEUTRALITY

What is net neutrality and why is there a debate about it? Basically, the proponents of net neutrality legislation argue that ISPs should not be permitted to speed up, slow down, or block Web content on the basis of the content's source, destination, or owner. Opponents of net neutrality legislation argue that ISPs should be free to manage traffic on the networks to provide the best quality service to their customers.

The debate over net neutrality has evolved in at least three main stages. In the first stage, the focus was largely on the ability of ISPs to block or degrade sites or applications that they either didn't like or saw as a commercial threat. A well-known example was the case of Madison River Communications blocking a competing Internet-based telephony service from Vonage, a case the U.S. Federal Communications Commission (FCC) intervened in successfully. As all of the major Internet service providers have agreed to the "Internet Four Freedoms" for broadband consumers (freedom to access legal content of their choice, to use applications of their choice, to attach personal devices of their choice, and to obtain information concerning their service plans) initially laid out by former FCC Chairman Michael Powell, this issue of outright blocking has receded in importance.

In the second stage of the debate, the focus was on the ability of ISPs to use tiered pricing in which they charge content or application providers for giving them priority service, without degrading or slowing other applications. Tiered pricing by ISPs was and continues to be considered controversial. Proponents of net neutrality believe that tiered pricing is unfair and may lead to anticompetitive behavior on the part of ISPs.

Free market proponents and network operators argue, on the other hand, that the market will address any potential for abuse and that no additional oversight is needed. But as the Information Technology and Innovation Foundation noted in "A Third Way on Net Neutrality," neither position adequately describes the nature of the problem nor articulates the kind of balanced approach that is needed.

The third stage of the debate over net neutrality that has emerged more recently pertains to what has been termed broadband network management. The issue of network management came to the fore when Comcast engaged in practices to limit the uploading of certain kinds of peer-to-peer (P2P) files at certain times of the day. More generally, the issue of network management refers to whether and to what extent ISPs can manage their networks to assure quality of service for the majority of their customers. Strong advocates of net neutrality argue that ISPs should have little flexibility to manage their networks and that the solution to any kinds of network congestion or other network performance challenges can and should be solved by simply adding more network capacity—primarily in the form of "bigger pipes." Several net neutrality proposals under consideration in the United States, Canada, and Europe (see box 1).

As this report demonstrates, however, the issue is far more complicated than the advocates of net neutrality suggest. At least for the foreseeable future, adding more network capacity will not solve network congestion or other network performance challenges because demand has a way of soaking up supply. Furthermore, as discussed in this report, certain types of applications on the Internet—including many P2P applications—are intentionally designed to take most of the available capacity on a network. This situation makes the use of other applications, including Internet telephony, which is dependent on low latency, quite difficult. Thus, in the absence of intelligent network management, the quality of many broadband users' Internet experience will be diminished.

The Internet and broadband technology are continuing to evolve at a fairly rapid rate, and neither shows any signs of maturing. This situation makes it very difficult, if not impossible, to predict where the market and technology will evolve. Given this environment, policymakers would be well advised to avoid issuing blanket prohibitions on certain types of behavior by ISPs. A better approach would be to foster concerted efforts by all parties, including the FCC, to find solutions that give ISPs the tools they need to effectively manage their networks but in ways that are clearly in the public interest.

If the United States is to make effective policy in this area, policymakers must base their decisions on an informed view of how advanced broadband networks and the applications running on them actually work. Much of the debate over broadband network management to date has unfortunately been informed more by rhetoric and emotion than by an actual examination of how advanced networks and the applications running on them work. By providing policymakers with a guide to how broadband networks and the applications that run on them work, the authors of this report hope to help change the nature of this debate.

THE EVOLUTION OF NETWORKS AND NETWORK MANAGEMENT

Before the development of the Internet, the telephone system was the most prevalent network in the world. The phone system was built on a common networking technology described below called circuit switching. As described below, the circuit-switching network technology created for real-time voice communications (telephony) inherently lacks the flexibility and efficiency needed for the Internet.

The Internet and its predecessor ARPANET were based on a whole new type of networking technology called packet switching, also described below. As time goes on, circuit switching and packet-switching network architectures are becoming more like each other and adopting each other's strengths. The focus of this paper, however, is predominantly on the management of packet-switching networks such as the Internet.

A. Circuit-Switching Networks Used in the Telephone System

Circuit-switching networks used in the telephone system allocate fixed resources when a connection is initiated and these resources remain in use until the connection is terminated. Thus, if too many people try to call at the same time and all available circuits filled up, subsequent callers are denied access to the

BOX 1: NET NEUTRALITY PROPOSALS UNDER CONSIDERATION AROUND THE WORLD

The governments of the United States, Canada, and Europe are considering three types of net neutrality proposals that seek to regulate prioritization technologies and enhanced Quality of Service (QoS).

Net Neutrality Bills That Completely Ban Prioritization

- S. 2360: Internet Nondiscrimination Act of 2006, Ron Wyden (D-Oregon)[a]
- Bill C-552: Charlie Angus, Canadian House of Commons

Net neutrality bills that completely ban prioritization take a hard-line approach that network prioritization should not be permitted and they make no exception for whether an application is sensitive to packet delay or not. S. 2360 in the United States explicitly bans "allocating bandwidth" and forbids Internet service providers (ISPs) from favoring their own content, which can be interpreted to include Internet Protocol-based TV (IPTV) riding over the last mile of the Internet. Bill-552 in Canada simply bans network prioritization, regardless of the application. Although both bills do make exceptions for prioritization of "emergency communications," there is no practical way for an ISP to detect an emergency phone call in a Voice over Internet Protocol (VoIP) stream, especially if the call-control or the entire VoIP session is encrypted. Even if nothing is encrypted, these bills would effectively mandate that ISPs snoop in on the call-control portion of the VoIP stream to look at every phone number that the consumer dials. Neither of the two aforementioned bills has been enacted. It is important to note that banning prioritization technology on the Internet would force ISPs to privatize more of their network using less-efficient circuit-switching networks and allocate fixed bandwidth to get the QoS they need. This would end up decreasing Internet capacity allocation, a result that is just opposite of the legislators' intention.

Net Neutrality Bills That Ban Multitier Quality of Service (QoS) and Sale of QoS

- H.R. 5273: Net Neutrality Act of 2006, Ed Markey (D-MA)[b]
- S. 215: Internet Freedom and Preservation Act of 2007, Olympia Snowe (R-ME) and Byron Dorgan (D-ND) [c]
- H.R. 5417: Internet Freedom and Nondiscrimination Act of 2006, Jim Sensenbrenner (R-WI) and John Conyers (D-MI). A nearly identical bill was introduced in 2008. [d]

Net neutrality bills that ban multitier QoS and sale of QoS, the most common type, take a more nuanced stance that permits QoS technology only if everyone got the same enhancements regardless of whether they paid for the service. The aforementioned bills, none of which has been enacted, were intended to prevent ISPs from offering content providers premium delivery services at a fee and to prevent broadband providers from favoring their own content such as their IPTV service. But almost all large-scale content providers on the Web pay private CDNs for guaranteed delivery, so broadband providers offering these CDN services simply make the CDN market more competitive. It should also be noted that mandating a single QoS tier and banning exclusive QoS enhancements would force broadband providers to resort to operating video on private networks, a situation that would decrease Internet capacity allocation.

A Proposal for Minimal Quality of Service (QoS) Mandates for All Applications and Services

- The European Parliament Amendment 22(3) put forth a proposal that would have national regulatory authorities issue minimum quality of service requirements on the Internet. The details of the language are still being hammered out in parliament.

Networks are shared by end users, and the performance of a network can vary significantly over time. The Internet is a best effort network, and performance levels are never guaranteed. That is why all large-scale on-demand video distribution services like YouTube bypass much of the Internet using private CDNs. Mandating a minimum quality of service standard could mean that normal variations in Internet performance may now be interpreted as a violation. Some have suggested that the minimal QoS requirement is actually a minimal requirement for universal broadband service in Europe. Amendment 22(3) is not written in the context of universal service, however, because universal service broadband requirements generally define the minimal performance of a broadband connection for the consumers; not minimal performance for the content or application provider. Amendment 22(3) is vague enough that it could be interpreted as minimal quality for content providers. Like the bills proposed in the U.S. Congress that mandate a single tier of QoS, Amendment 22(3) is intended to prevent ISPs from offering content providers premium delivery services at a fee. Premium services would be less relevant if content providers were guaranteed a minimum QoS, assuming that that minimum level of service is high enough.

S. 2360, Internet Nondiscrimination Act of 2006. Available at: <www.publicknowledge.org/pdf/s2360-109.pdf>.
H.R. 5273, Net Neutrality Act of 2006. Available at: <markey.house.gov/docs/telecomm/MARKEY_002_XML.pdf>.
S. 215, Internet Freedom and Preservation Act of 2007. Available at: <www.publicknowledge.org/pdf/s215-110-20070109.pdf>.
H.R. 5417, Internet Freedom and Nondiscrimination Act of 2006 (revived in 2008). Available at: <www.publicknowledge.org/pdf/hr5417-109.pdf>.

system and the user gets an "all circuits busy" message. Additional users can use the phone system only if the circuits are open. If the network is not filled to capacity, circuits remain unused. Unused circuits cannot be dynamically allocated to the existing connections. Thus, for example, a fax machine cannot transmit a fax through the telephone system based on circuit switching any faster even if only 10 percent of the circuits are in use. Another characteristic of circuit switching networks is that communications are generally limited between two devices unless the circuit is terminated and a new circuit created. Simply put, if Bob is on the phone with Alice but he wants to talk to Mary, he has to hang up on Alice and then dial Mary's phone number.

Even though circuit-switching networks offer less than optimal resource utilization and limited flexibility, such networks do offer the advantages of consistency and predictability. Consistency and predictability are ideal for real-time applications such as telephony and video conferencing. Circuit-switching networks are too limiting and inflexible to be used for the Internet. The Internet and its predecessor ARPANET are based on a newer technology called packet switching.

B. Packet-Switching Networks such as the Internet

The Internet and its predecessor ARPANET were the first packet-switching networks in the world. In 1969, when the Defense Advanced Research Projects Agency of the U.S. Department of Defense created a wide area network called ARPANET to connect government research centers and universities. The first ARPANET used Network Control Program (NCP) as its communication protocol (language and syntax).

In 1972, Robert Kahn began working on the next ARPANET communications protocol called the Transmission Control Protocol/Internet Protocol (TCP/IP), and he was joined by Vinton Cerf in 1973. In 1983, ARPANET officially switched over to TCP/IP, the U.S. military split off from ARPANET for security purposes, and the Internet as we know it today was born. As described in Box 2, the Internet is a network of networks.

The packet-switching network used in the Internet is a radical change from the circuit-switching network. A packet-switching network allows everyone to be on the network at the same time and dynamically divides up the resources among the active users on the network. If very few users are on a packet-switching network, then those few users get a lot of resources allocated to them. If many users are on the network, then each user gets fewer resources but at least is not locked out of the system. Unlike a circuit-switching network, a packet-switching network allows multiple devices to communicate simultaneously and to remain connected to the network all the time instead of having to go through a time-consuming dial-up process such as that required on the phone system when a fax machine begins sending a fax.*

C. The need for Quality of Service (QoS) on Packet-Switching Networks

Packet-switching networks such as ARPANET and the Internet were invented for their flexibility and efficiency characteristics, which are optimum for data applications. The disadvantages of a packet-switching network are (1) a lack predictability for real-time applications such as VoIP, online gaming, video conferencing, and IPTV, which stems from the fact that the network is doing multiple things at the same time; and (2) high jitter. High-jitter conditions are essentially micro-congestion storms that last tens or hundreds of milliseconds. High jitter occurs whenever a large number of packets come from a faster network link to a slower network link or where several networks links merge to a single link. When this happens, network devices such as routers and switches get backlogged, and they force packets to wait inside their memory buffers, increasing the time it takes packets to traverse a network.

To prioritize network traffic on packet-switching networks for real-time applications over applications that are relatively insensitive to packet delay, network managers can use traffic-shaping mechanisms called Quality of Service (QoS). There are many common alternative words used to describe QoS mechanisms, including "enhanced Quality of Service," "network intelligence," "prioritization," or "premium service."

*Early Internet users will recall that the first dial-up Internet access accounts relied on phone and modem technology, which also used those slow and noisy handshakes to establish connections to the Internet. This is actually an example where the last-mile access layer of the Internet was a circuit switched phone network where bandwidth was fixed, slow, and dedicated. But once the user got beyond the ISP modem banks, it was a packet switching Internet that allowed users to get to any website they wanted.

BOX 2: OVERVIEW OF THE INTERNET

The Internet, derived from the word internetwork, is a network of networks. Small networks are linked together by a common protocol called Transmission Control Protocol/Internet Protocol (TCP/IP) to create the larger global network called the Internet. What started off as a military and academic project in the 1960s was quickly commercialized in the 1990s, and the Internet became the common protocol and communication medium for every computer in the world.

Today, the Internet primarily consists of private networks that are interconnected via peering arrangements and contractual agreements. A peering arrangement is a contractual agreement between two Internet service providers (ISPs) to interconnect and exchange traffic. This peering arrangement may involve the larger ISP charging the smaller ISP because the smaller ISP will benefit more from the infrastructure built and paid for by the larger ISP or it might involve no money changing hands if each ISP has something of equal value.

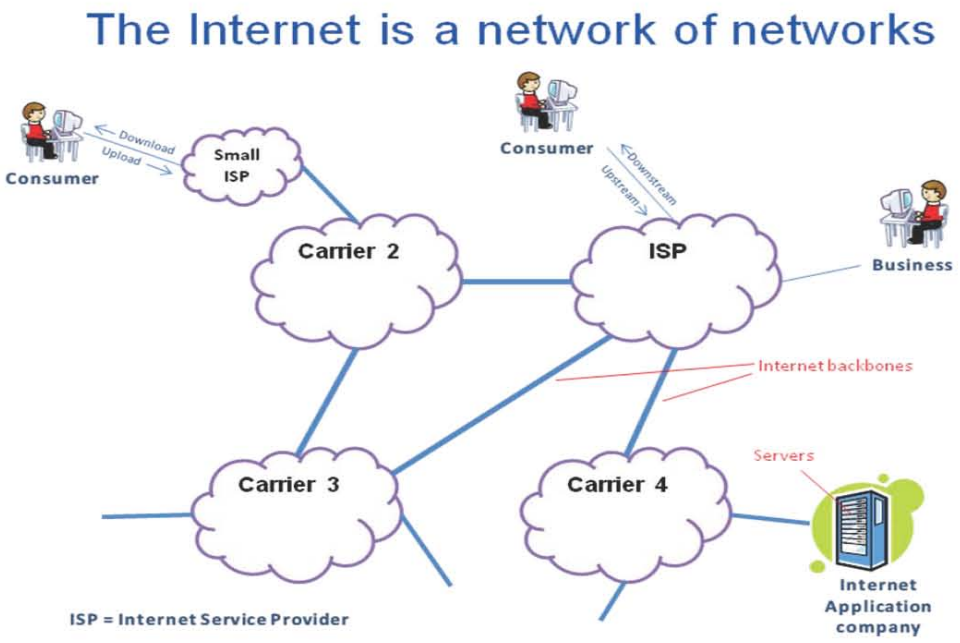

Note: The TCP/IP protocol has been so successful that even networks not connected to the Internet use TCP/IP. It's also common to hear the term "IP Network" refer to a generic network that runs TCP/IP. Many newer phone systems, for example, run on an IP Network, but they're often not connected to the Internet.

QoS essentially carves out virtual circuits within a packet-switching network by providing the necessary resources and performance characteristics that real-time applications need. This gives a packet-switching network the real-time characteristics of a circuit-switching network while maintaining the robustness and flexibility of a packet-switching network. Voice (telephony) and video conferencing have long been handled by circuit-switching networks because of such networks low and predictable packet delays. QoS allows a packet-switching network to become just as reliable as a circuit-switching network.

The Internet is an evolving platform that needs to adapt to support new applications and services such as voice and video conferencing and it has already undergone many changes. Going as far back as 1992, the Internet Engineering Task Force (IETF) has been trying to define and standardize QoS mechanisms. The architecture of the Internet should not remain and has not remained in a frozen state of development. Because the modern Internet is expected to become the "jack of all trades" and handle every type of application, packet-switching networks have to become more predictable and jitter free with more network intelligence.

Effort to make circuit-switching networks more flexible and efficient are also underway, with technologies like switched digital video that allow cable TV companies to make more efficient use of their infrastructure.

STATIC VS. DYNAMIC ALLOCATION OF NETWORK CAPACITY

One of the fundamental questions asked by many people is, "Why do networks even need to be shared and managed and why can't we just build more bandwidth"? The reason is that short of having networks with infinite capacity, network bandwidth will always have to be managed. The fact that infinite capacity networks will never exist means that either (1) the network allocates fixed bandwidth determined by the worst-case minimum bandwidth the network can support; or (2) the network dynamically expands and contracts available bandwidth to each customer depending on the number of active users at any given time.

For example, if a network can be built to guarantee 1 Megabits per second (Mbps) of performance for each user, it can just as easily offer the customer a 1 Mbps of guaranteed performance and up to 20 Mbps of unguaranteed performance. If 5 to 20 percent of people are actively utilizing a network at the same time, the worst-case bandwidth is 5 Mbps and the best case is 20 Mbps if the network divides up the available bandwidth evenly. In other words, even when a managed network is running slower during peak hours of use, it will offer each customer more bandwidth at any given time than the guaranteed minimum bandwidth offered by a network with fixed bandwidth.

The concept of gaining several times more bandwidth through dynamic resource allocation is called "statistical multiplexing,"—and it is the law of efficient networking. Opponents of network management object to the contraction of bandwidth because they feel that building more bandwidth is a better option, but a network that never contracts is by definition a network that can never expand to take advantage of idle capacity. Engineers cannot allocate best-case fixed bandwidth assuming that only 1 percent of the customer base will be active at any time because the minute more than 1 percent of the customers use the network, a network that doesn't permitted bandwidth contraction will melt down. Fixed bandwidth networks by definition have to operate on a worst-case basis and this is precisely the reason designers of the Internet rejected the fixed bandwidth circuit-switching model.

Statistical multiplexing is the law of efficient networks regardless of how fast networks become in the future and using it allows networks to operate 5 to 20 times faster. ISPs use statistical multiplexing to give consumers the most performance at an affordable price. There are guaranteed bandwidth Internet services offered to most commercial sectors, but these services are typically 30 times more expensive per Mbps than the variable bandwidth service offered through statistical multiplexing. Consumers prefer the much faster variable bandwidth service over the guaranteed service because most don't want to pay what it requires to get the higher bandwidth guarantees.

The Internet has had basic network management mechanisms built into it since its inception, although these mechanisms have undergone much refinement over the years and continue to undergo refinements needed because of the changing usage of the Internet.

If statistical multiplexing is such a wonderful and efficient technology, then why does it have such a poor reputation? One reason is the misperception that the advertised peak performance is the same as the guaranteed performance. Perhaps it is because consumers see an advertisement for Internet service offering "up to 16 Mbps" (peak performance) and overlook the "up to" portion of the advertisement. Part of the fault lies with ISPs that have not done a good job educating the public with their marketing campaigns that the service they are selling is shared bandwidth. Consumers expect the advertised bandwidth to be the guaranteed bandwidth, and if they are throttled below the peak advertised performance, they get dissatisfied.

The obvious solution to this problem is clear disclosure, but no ISP wants to come out and advertise the minimum and maximum performance of its services unilaterally while its competitors continue to embellish their offerings by touting peak performance. Government can play an important role in helping to solve this problem by mandating broadband advertising rules that require both the guaranteed and peak

performance figures to ensure an even playing field between ISPs providing broadband.

JACOBSON'S ALGORITHM: A TCP CONGESTION–CONTROL MECHANISM BUILT INTO THE INTERNET

As a shared packet-switching network that dynamically allocates resources to active users, the Internet requires a way to fairly distribute and allocate available bandwidth among active users. It also requires a traffic-policing mechanism that prevents users from overwhelming the network with too much data. As described below, the Internet has had basic network management mechanisms built into it since its inception, although these mechanisms have undergone much refinement over the years and continue to undergo refinements needed because of the changing usage of the Internet.

The Internet is built on various networking protocols, most importantly the Internet suite known as Transmission Control Protocol/Internet Protocol (TCP/IP). TCP/IP is divided into an endpoint (computer) component and a network component. TCP runs on the endpoint devices such as computers, printers, and Web servers, while IP runs on the network infrastructure on devices such as routers. End-point devices attached to the Internet manage things like data transmission speed and error correction, while routers control the proper routing and delivery of packets.

Figure 1: Exploiting TCP Congestion Control

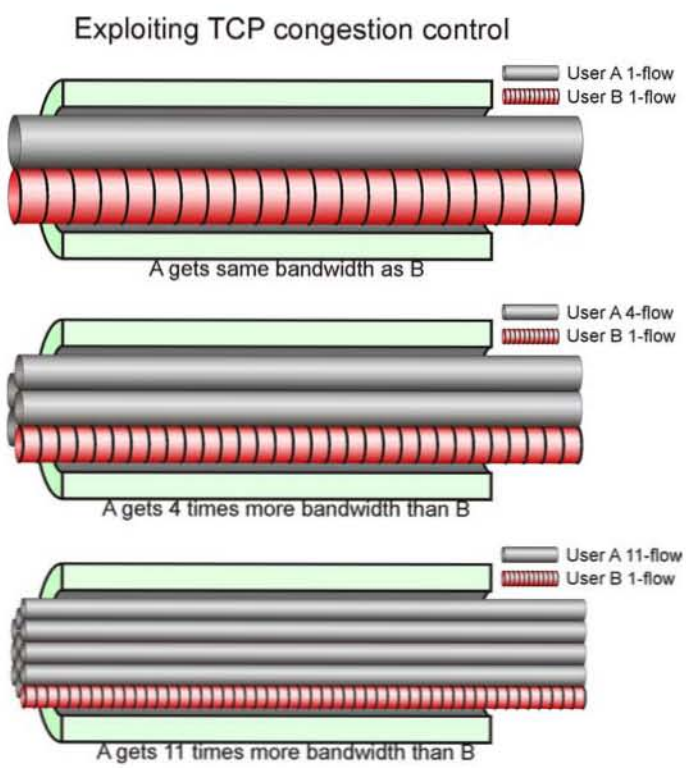

This separation of duties does not mean that computers and routers operate independently of each other. The network and the endpoints must closely interact to make this all work. Since 1987, for example, computers have used a revised version of TCP that includes a network congestion control mechanism developed by computer scientist Van Jacobson to slow down (throttle) the endpoints and prevent meltdown (see box 3).

During the early days of the Internet, Jacobson's TCP algorithm mostly succeeded in fairly allocating bandwidth among users. The allocation was fair because people didn't routinely run multiple applications at the same time and early Internet applications like File Transfer Protocol (FTP) used a single TCP flow to transfer a file. A flow is the data stream between two endpoints that have made a connection to each other. During the 1990s, Web browsers, which use the Hyper Text Transfer Protocol (HTTP), also used single TCP flows to transfer individual files.

This all changed dramatically in 1999 when the first peer-to-peer (P2P) application called Swarmcast was created. Swarmcast actively exploited the fairness loophole in Jacobson's TCP algorithm. By running multiple TCP flows per file transfer, the application could effectively gain "immunity" from Jacobson's algorithm. Under a congested network, P2P could run many times faster than other traditional file transfer protocols such as FTP or HTTP. If the P2P application ran 20 TCP flows, it could run up to 20 times faster than FTP or HTTP under a congested network.

Figure 1 shows small pipes representing data flowing through a cut-away pipe which represent a network. Remember that the TCP/IP packet-switching network is shared and it dynamically allocates and divides bandwidth among the active users. The active users are all competing for shared bandwidth.

Starting from top to bottom, the drawings show User A taking a bigger and bigger share of the pipe as his P2P application increases the number of TCP flows

BOX 3: JACOBSON'S ALGORITHM FOR AVOIDING NETWORK CONGESTION AT WORK

Since 1987, computers have used a revised version of the Transmission Control Protocol (TCP) that includes a network congestion control mechanism developed by computer scientist Van Jacobson.

Jacobson's algorithm works by leveraging the fact that routers in a network becoming overloaded with data traffic will have to drop packets. A computer running Jacobson's algorithm treats the dropping of packets as an implicit notification from the network that it is unable handle any more traffic. When a computer running the algorithm fails to get an acknowledgment from the computer it's sending data to because the data was never delivered, it responds by instantly cutting a computer's transmit rate in half.

Jacobson's algorithm then allows an additive increase in transmit rate with each successful acknowledgment from the computer receiving data. Thus, the computer sending data will continue to speed up the rate of data transmission until the network can no longer support the data flow and begins to drop more packets. At that point, the computer running Jacobson's algorithm will cut the rate of data transmission in half again. Then the whole cycle—the official name of which is "additive increase/multiplicative decrease" (AIMD)—will begin again.

The graph below shows a network in conjunction with Jacobson's algorithm evenly dividing up bandwidth between two TCP flows. As the aggregate bandwidth from the two flows reaches the limit, the router randomly drops packets on the network, which statistically hits the flow with the fastest rate. Every time a TCP flow experiences a dropped packet, its rate is cut in half by a multiplicative decrease. With every successful acknowledgment, it speeds up smoothly with an additive increase.

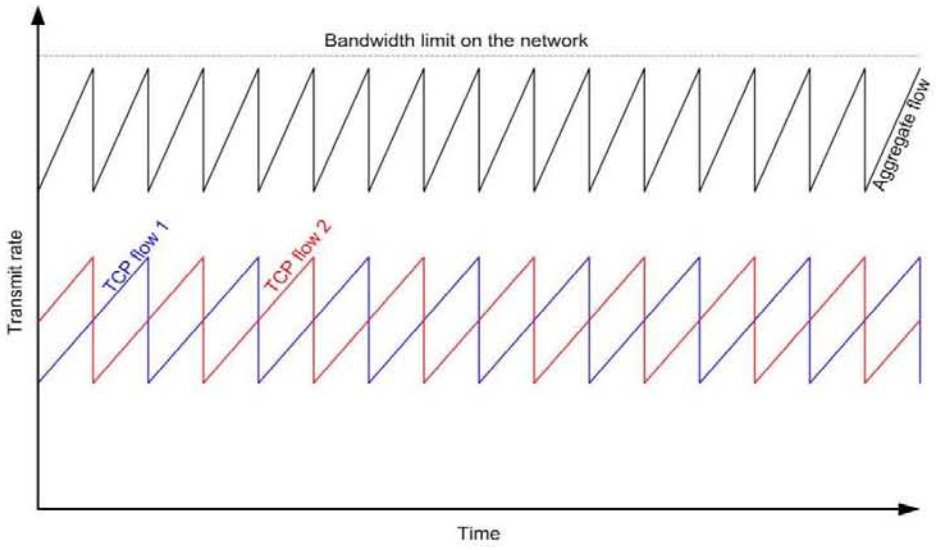

Note: This is an oversimplification of Jacobson's algorithm. There are more complex factors such as the average packet delay, the actual TCP implementation, and other factors in play that determine whether bandwidth is evenly divided up. This mechanism also assumes that the computer on the endpoint "plays nice" and uses a standard TCP implementation. If the computer on the endpoint voluntarily cheats and refuses to back off on transmit rate, it can gain a performance advantage and there isn't much a typical network can do about it. Furthermore, if the endpoint isn't using TCP at all and is instead using User Datagram Protocol (UDP),TCP congestion control doesn't even apply, and the management of UDP is left up to the network equipment.

while User B's traditional application gets squeezed into a smaller and smaller pipe.

All the individual small pipes representing TCP flows are all roughly the same size. The difference for User A is that he has more pipes giving him a higher aggregate speed. User A is effectively taking multiple bites from the apple to gain a larger share and his application is exploiting the system.

This creates a huge degree of unfairness between users where P2P users running applications like BitTorrent are allocated a disproportionately large share of bandwidth at the expense of other users that are vying for the same bandwidth using non-P2P applications.

NETWORK FILE DISTRIBUTION ARCHITECTURES

To understand network management and network congestion, it is important to understand that there are three main file distribution architectures on the Internet: (1) client-server, (2) peer-to-peer (P2P), and Content Delivery Network (CDN). Each of the three major file distribution architectures has different implications for network capacity. These differences are important to consider when implementing network management solutions.

As described below, each of these architectures has pluses and minuses, and each is most suitable to different situations. The client-server architecture has limited scalability and high cost of operation. This led the P2P architecture to very popular within the ad hoc file trading community. But P2P architecture lacks the reliability that the commercial market expects; it also lacks the ability to serve content in order for on-demand video streaming. For that reason, large-scale file distribution services and on-demand video services universally prefer more expensive but more robust CDN architecture.

A. Client-Server File Distribution Architecture

One of the first applications of the Internet was file transfer using File Transfer Protocol (FTP) or Hypertext Transfer Protocol (HTTP) and client-server file distribution architecture. A "file" could be anything such as an electronic document, a digital photo, digital music, or digital video. The client-server method of file distribution was simple to understand because you had a server that served or "uploaded" files and you had clients that wanted the files by "downloading" them from the server (see figure 2).

The problem with the client-server architecture is that the server's scalability depended upon the amount of upstream capacity it had for uploading files. If it had 10 Mbps of upstream capacity, it could either distribute to 10 clients at 1 Mbps each or it could distribute to 31 clients at 0.32 Mbps which is sufficient for low-resolution video streaming. Even if the server capacity is 10,000 Mbps, which requires a huge server and at least $30,000 per month in bandwidth fees as of 2008, the realistic capacity for a video streaming server delivering YouTube level quality at 320 Kbps is only 31,000 simultaneous users. There's also less assurance that the server's bandwidth can overcome congestion at every stage of the Internet to reach every corner of the global network.

Figure 2: Client Server Model

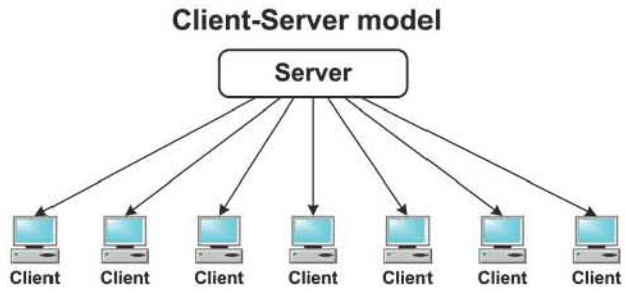

B. Peer-to-Peer (P2P) File Distribution Architecture

The P2P file distribution architecture illustrated to the right solved the scalability and cost issues of the client-server model (see figure 3). In the P2P architecture, the server—now called a "seed"—can leverage the fact that its clients—now called "peers"—can upload any data they receive to other peers, and those peers will in turn upload to other peers and so on. The peers automatically become seeds the instant they accumulate all the necessary pieces of the file. This chain of events can go on indefinitely.

The three operational modes of clients in P2P networks—peers, seeds, and leeches—are defined more fully in Box 4. The higher the ratio of seeds to peers in the P2P architecture, the faster the download speeds for the peers. This is the "secret sauce" that allows

P2P file distribution to scale indefinitely without massive amounts of server bandwidth or paying for distribution services. If no one uploads, the entire P2P architecture breaks down.

P2P file distribution effectively doubles the network traffic load for ISPs because downstream traffic is accompanied by upstream traffic, which is needed to sustain the P2P chain and feed other P2P users. And whereas most other applications use a single TCP flow, P2P applications use multiple flows to gain immunity against Jacobson's algorithm.

The P2P architecture shifts the majority of bandwidth and server costs away from content distributors and to the end users and ISPs. Although this phenomenon imposes some costs on large telecom or cable broadband providers, it is particularly burdensome for smaller ISPs, which pay much higher costs for bandwidth in remote areas than server operators that can purchase bandwidth with bulk discounts in large data centers. As of 2008, bandwidth in large data centers costs as little as $3 per megabit/second whereas small rural ISPs have to pay $100 per megabit/second. Brett Glass, the owner of a small wireless ISP in Wyoming put it best when he said "Most independent ISPs have high bandwidth costs. Any [network management] 'solution' which doesn't recognize that will push us toward duopoly."

Another limitation of P2P architecture is the fact that content is delivered out of order because the content

Figure 3: Peer-to-Peer (P2P) Model

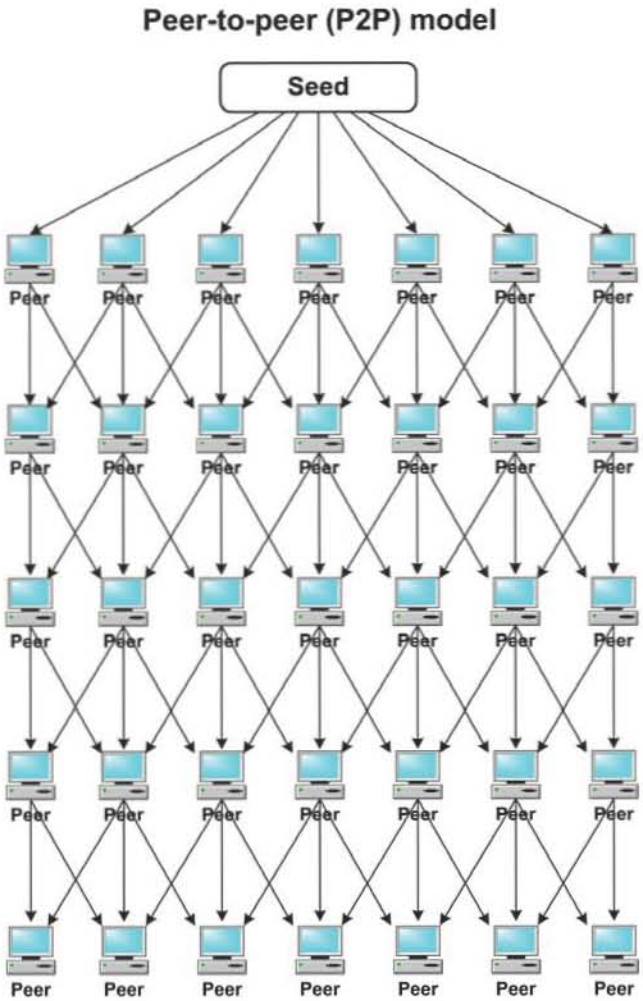

BOX 4: OPERATIONAL MODES IN P2P NETWORKS

Clients on P2P operate in three basic modes:

- **Seed.** A seed is a client on the P2P network that has a complete copy of a file. This could be the original source of the file being distributed or a P2P user who has already finished downloading the file.
- **Peer.** A peer is a client on a P2P network that is uploading and downloading pieces of a file on the network. P2P applications will upload the pieces it already has by default to keep the torrent healthy. P2P peers usually become seeds automatically once they have a complete copy of a file.
- **Leech.** A leech is a client on a P2P network that is downloading pieces of the file but not uploading anything. This is perfectly acceptable if the P2P client doesn't have any pieces that any other client wants, but a P2P user who refuses to upload at any time is heavily frowned upon. Popular P2P applications like BitTorrent will severely penalize a P2P user's download speed if the user refuses to upload.

is coming in from multiple sources. This means the music or video content being downloaded cannot be heard or viewed until the entire file transfer is completed, something that could take several hours or sometimes days. This limitation of P2P architecture makes it unsuitable for on-demand video or music streaming; however, it's a small tradeoff for a free scalable distribution system such as P2P, so P2P is the preferred file distribution architecture for free content. P2P traffic includes considerable illegal content such as pirated music and pirated videos, although it also includes legal content, as described below.

The corporate IT world has been reluctant to embrace P2P technology because it is sometimes slow to adopt new technologies and sometimes mistakenly view all P2P usage as piracy. Nevertheless, the P2P architecture does have some legitimate uses and advantages. P2P enables both small and large companies to distribute files on a large scale within their operations. This is why free software such as Linux is often distributed with P2P technology. P2P can have also significant architecture advantages both in terms of resilience and cost. For one thing, P2P has much better transfer resume capability if a file transfer was interrupted compared to traditional protocols like FTP and HTTP. The other advantage is major advantage is cost. This advantage can be illustrated by the following real-world example. If a national chain store wanted to distribute employee training videos to all of its outlets, it would cost a significant amount of money to buy upstream capacity for a centralized server. Instead, if they leverage their existing upstream capacity in every store and used a P2P application, they could do it at no additional cost.

While many people call the P2P file distribution model "efficient", it is efficient from the point of view of content distributors because it saves them bandwidth and server costs, and as discussed above, this is one reason why some organizations have embraced it for internal file distribution. However, when used to distribute files outside of individual organizations, The P2P architecture is anything but efficient for broadband providers because they have to carry twice the traffic for users to access the same content. The P2P architecture is also inefficient for non-P2P users who end up giving up bandwidth disproportionately because P2P applications don't back down as much under congestion.

To illustrate these points, consider the British Broadcasting Corporation (BBC) software called iPlayer, which is used to distribute high-quality content to BBC viewers on the Internet via a P2P architecture. The iPlayer has received praise from happy viewers pleased with the high-quality downloadable content. But it has been criticized by ISPs because it has caused a significant rise in network traffic. Its P2P file distribution architecture creates twice the traffic load of other architectures because it involves both upload and download traffic. In addition, P2P peers engage in random peering behavior that causes P2P data to traverse more network than otherwise necessary. (However, the P2P protocol has been significantly improved by the P4P standard, which reduces backhaul congestion created by random peering and improves file transfer speed. The P4P working group is a consortium of content providers and network operators who see potential benefits for all parties.)

Another complaint about the BBC's iPlayer is that it installs a hidden KService application, which continues to run in the background using the user's bandwidth even when iPlayer is shut down. The running of this application in the background slows the individual user's broadband connection and could in some cases result in large bandwidth costs for the user. A user who knows how to customize the iPlayer settings can disable this application, but many consumers do not bother to check the default settings. This situation illustrates why it's important to have transparency not just from the ISPs that offer broadband but also from content and application providers.

C. Content Delivery Network (CDN) File Distribution Architecture

The CDN file distribution architecture, an alternative to the client-server architecture and the P2P architecture, uses high-speed cache servers that are distributed across the Internet (see figure 4). These high-speed servers store copies of the files and redistribute them to nearby clients to bypass long-haul Internet connections. This not only speeds up file transfers, it also alleviates congestion on the Internet backbone.

Without CDN, the same content would have to traverse the entire Internet every time someone requested the content. With thousands or millions of people requesting the same popular content, it's crucial to keep

that redundant from choking the core of the Internet. In fact, just a single CDN company Akamai claims to carry 20% of the world's Internet traffic. This is not surprising in light of the tremendous growth in online video streaming and the fact that traffic at private Internet switching centers that host CDN companies like Switch and Data have experienced up to 295 percent growth in 2008.

The CDN file distribution architecture has many advantages. Because the CDN file distribution architecture supports the delivery of on-demand content in order, nearly all large-scale file, video, and music distribution companies have opted to use this architecture. Microsoft, for example, uses CDN to deliver updates to its Windows operating system to hundreds of millions of Windows computers in the world. YouTube uses CDN to deliver on-demand video to millions of users. Apple uses CDN to deliver music and videos to its iTunes customers.

CDN is the dominant architecture because broadband users generally do not want to use their own server and upload capacity for content they've already paid for (either in the form of money or the attention they've paid to commercials); moreover, most consumers want instant gratification by being able to view the content while it's downloading. Because it does not require the clients to upload and download the way P2P architecture does, the CDN architecture puts half the load on broadband networks for a given load. This means that on broadband connections aren't tasked on CDNs with constant uploading of data, a task that they were never designed to handle in the first place. This feature of the CDN architecture has significant implications for broadband networks that have limited upstream capacity, including cable broadband networks, and even bigger implications for wireless broadband networks that share wireless spectrum for both the upstream and downstream.

The downside of CDN architecture is that it's not free to content providers the way the P2P architecture is. But despite the commercial nature of CDN architecture, it is the dominant architecture used by citizens to freely broadcast video. Although it's possible to use P2P architecture to offer higher quality video downloads, people clearly value the instant gratification and accessibility of the CDN video on demand model. This is why wildly popular services like YouTube have used CDN to empower anyone to broadcast video on demand to the entire world. Independent film makers are even utilizing high-quality services like Vimeo to deliver near-DVD-quality video on demand.

IMPROVING FAIRNESS BETWEEN BROADBAND CUSTOMERS

Fairness dictates that customers who are paying for the same tier of broadband service from a broadband provider should get roughly the same bandwidth at a given level of usage. Unfortunately, the Internet's existing mechanisms for controlling congestion does not allocate bandwidth fairly. As a consequence of the multi-flow nature of P2P applications and the exploitation of Jacobson's algorithm, P2P users on dumb

Figure 4: Content Delivery Network (CDN) Model

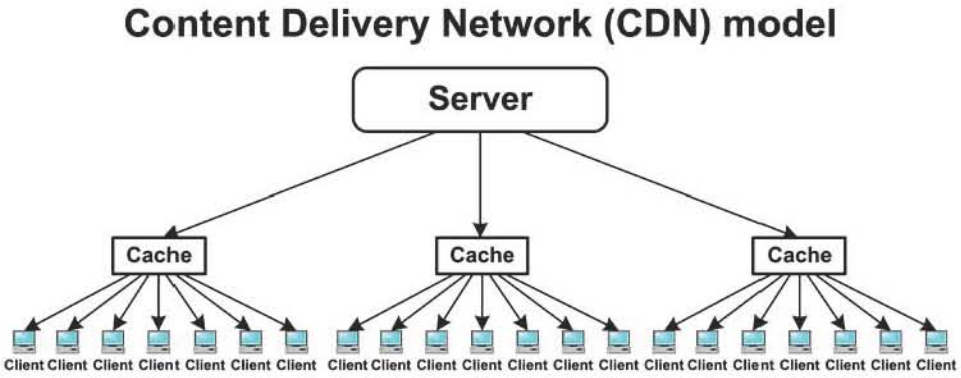

packet-switching networks get a disproportionately high amount of bandwidth.

As discussed below, there have been widespread efforts to address the problem of fairness between broadband customers, including efforts to establish Internet standards to facilitate fairness, protocol-specific throttling mechanisms that target a particular protocol (e.g., P2P), and protocol-agnostic network management solutions. The latest protocol-agnostic solutions, such as those being implemented by Comcast, are an important start but they need to be expanded upon to deal with problems like jitter and making broadband more conducive to simultaneous application usage.

A. Internet Standards

Researchers Frank Kelly and Bob Briscoe have discussed the unfairness of TCP congestion control and the shortcomings of Van Jacobson's algorithm. Briscoe has appeared before the Internet Engineering Task Force (IETF) standards body arguing for changes to the TCP standard that would be designed to facilitate per-user fairness rather than per-flow fairness. He has also released a problem statement summarizing many of the issues. Briscoe wants the TCP algorithm to factor in the number of flows so that a multi-flow TCP application doesn't get an advantage over a traditional single-flow TCP application.

Fixing TCP congestion control fairness at the Internet standards bodies is important for the long term, but years or possibly even decades might elapse before Internet standards are ratified and deployed. To illustrate how long it takes to get new IETF standards adopted, consider Explicit Congestion Notification (ECN). ECN is a superior congestion control mechanism ratified by IETF in 2001. ECN has been implemented in Linux and Microsoft Windows Vista (the latest mainstream operating system in the world), but Microsoft disables ECN capability by default because of the possibility of problems with a small percentage of legacy routers (including home routers) currently in deployment. Thus, in 2008, seven years after ECN became an official IETF standard, we're still nowhere close to widespread deployment of ECN, and everyone is still using Van Jacobson's TCP patch. In 1987, when Jacobson's TCP algorithm was quickly adopted throughout the Internet, there were only around 30,000 computers to deal with. Today, though, there are over a billion devices in the world running TCP, and it will take a long time before we get to implement ECN.

B. Protocol-Specific Network Management Systems

Some institutions, including universities, corporations, and government agencies, deal with applications that consume too much bandwidth by simply blocking high-bandwidth-using P2P applications. In addition to keeping network congestion to a minimum, blocking P2P applications limits students and employees from downloading illegal pirated content.

Corporations that pay for an Internet connection and are paying their employees to work can clearly run their network as they please. For universities, the question of whether such blocking is justifiable is a bit less clear because the students are paying for their Internet connections either directly via their payments for room and board or contributing indirectly via their tuition payments. The question of what constitutes justifiable blocking becomes very complicated when we look at pubic kiosks, libraries, hotspots, schools, and hotels. It's clear that there are valid exceptions where blocking content and applications is justifiable. For example, some airlines have recently begun offering in-flight Internet access with blocked VoIP access because passengers don't want to sit through a flight listening to others talking on the phone.

Unlike corporations that block P2P applications for their employees, broadband providers can't simply block P2P applications because their users pay them for unfettered Internet service and some want to use P2P applications. But a few P2P users can consume so much bandwidth that they make everyone else's Internet experience horrible. In Japan, for example, P2P users representing 10 percent of the total broadband population account for 65 to 90 percent of all traffic on the network, making the network congested for everyone else.

Some broadband providers, in order to neutralize the multiflow advantage of P2P applications that can consume a disproportionately high amount of bandwidth and to ensure that the majority of their customers not using P2P don't suffer, use systems to throttle (slow down) P2P applications. In the United States, for ex-

ample, Comcast uses a protocol-specific throttling system from a company called Sandvine that issues TCP reset commands to disconnect and reduce the number of upstream TCP flows that a P2P seeder can have in a congested network. Comcast's protocol-specific throttling system from Sandvine does not affect P2P downstreams, so P2P peers and leeches are not affected unless they are trying to download from a seeder within Comcast's network that was facing TCP resets. Even then, most BitTorrent transfers have dozens of peers or seeders to download from, so a few reset connections merely slowdown the file transfer.

The dumb packet-switching network is the least fair system because it allows heavy users take resources at the expense of other users. Under a fair share network management scheme, everyone gets an opportunity at high bandwidth.

The protocol-specific throttling system used by Comcast had some unexpected side effects, including the accidental blockage of Lotus Notes, although that problem was quickly fixed. The biggest downside to the protocol-specific throttling system from Sandvine is that it can sometimes overly impact less popular P2P file transfers and frustrate users. If there are no other seeders or peers outside of Comcast's network, Comcast's using TCP resets on seeders inside Comcast's network may temporarily cause the P2P file transfer to completely stop. Comcast's protocol-specific throttling system has little to no effect on the majority of torrents, which typically have one or more seeders outside of Comcast's broadband network. Vuze, a company that distributes content using P2P file distribution, brought an official complaint to the FCC that Comcast was blocking Vuze downloads, but this complaint was not valid. The protocol-throttling system used by Comcast never blocked any Vuze P2P file transfers because Vuze has its own dedicated seed servers outside of Comcast's network.

The reality is that although protocol-specific throttling mechanisms such as the Sandvine system used by Comcast are not a perfect approach to network management, they're better than doing nothing at all for the near term. More accurate and more expensive network management solutions that do not focus on specific protocols do exist, and in fact, Comcast has committed to changing over to a protocol-agnostic approach to network management by the end of 2008. But protocol-agnostic network management approaches (discussed further below) are often too expensive for smaller independent ISPs to deploy. Brett Glass, who operates an independent ISP in Wyoming, put it best when he said, "One of those (advanced protocol-agnostic) boxes costs as much as what we'd pay to deploy new service to 180 square miles of previously unserved countryside."

C. Protocol-Agnostic Network Management Systems

To achieve fair bandwidth allocations, protocol-agnostic schemes are the best solution. ISPs can use protocol-agnostic network management systems (systems that measure the aggregate bandwidth consumption of each customer and not what protocols they are using) to ensure that bandwidth is shared fairly between customers. Early network management systems that used less accurate protocol-specific schemes to allocate bandwidth between customers worked well most of the time but experienced occasional problems. Newer protocol-agnostic solutions are being evaluated by broadband providers.

Jacobson's algorithm achieves per-flow fairness because it does not factor in the number of TCP flows each person uses. This allows P2P applications to use multiple TCP flows to consume far more bandwidth than other single TCP flow applications. Protocol-agnostic approaches to network management are designed to achieve a state of per-user fairness. Figure 5 shows how the per-flow fairness scheme allows some users to get far more bandwidth than other users, whereas the per-user fairness system gives everyone equal bandwidth.

Protocol-agnostic network management devices sit inside the network to reshape the traffic more equitably than dumb networks can. The protocol-agnostic device increases the statistical likelihood that a TCP flow will experience a dropped packet in proportion to the number of active TCP flows belonging to a particular user. Protocol-agnostic network management schemes also factor in User Datagram Protocol (UDP) traffic, which isn't managed by TCP congestion control at all.

This leads to equitable distribution of bandwidth between users of the same service (price) tier regardless of what protocols are being used.

As companies gain more experience with protocol-agnostic solutions, they are implementing them. Comcast initially opted for a protocol-specific bandwidth allocation approach to network management because the solution could be implemented out-of-band and therefore required less dramatic changes to its network architecture. The system did not require Comcast to insert an in-band device in the middle of its network, a measure that requires disconnecting and taking the network down. There was also the fear that in-band devices can also bring down an entire network if they fail during operation.

Comcast's protocol-specific approach would engage in a specific geographic region whenever the total traffic of P2P users exceeded 50 percent of all upstream network capacity within that region. Then the system targeted P2P seeders (customers that were only uploading and not downloading) and limited the number of upstream connections those users could have by sending them TCP resets. The rationale for this was that P2P users that were downloading needed to be able to continue uploading or face slower download performance and P2P users who were only uploading already got the content they desired.

Unfortunately, this meant that the occasional "rare torrent" (P2P BitTorrent files that don't have a lot of peers or seeders) seeding from Comcast's broadband network will have an exceptionally hard time being uploaded. Comcast's TCP reset mechanism made it even more difficult for dowloaders to access those rare torrents. Compounding the situation was the fact that Comcast did not initially disclose its practices and was not fully forthcoming to media inquiries regarding its network management practices. Ultimately, the pressure from the public and the FCC likely played a large role in Comcast's decision to convert entirely to a protocol-agnostic network management system by the end of 2008.

Comcast's new protocol-agnostic network management system called "Fair Share" will identify when the network is in a congested state. If there is no congestion on the Comcast network, then no management takes place. If congestion exists, then the Fair Share system will identify the heaviest users of bandwidth in

Figure 5: Unmanaged versus Managed Bandwidth Allocation

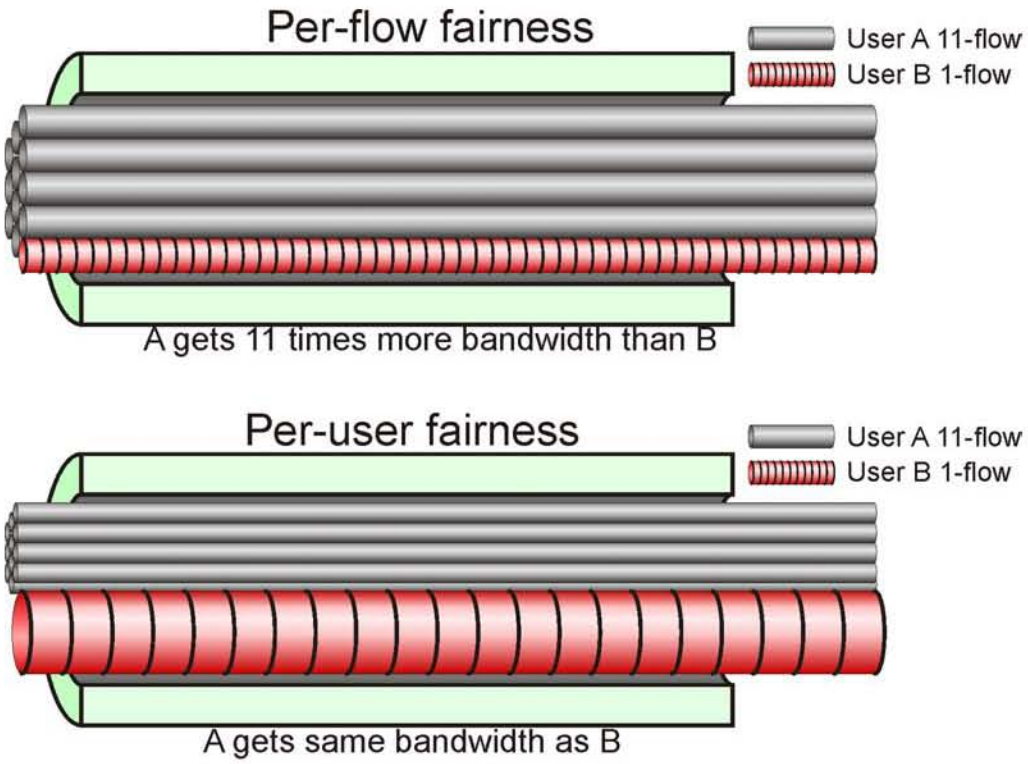

the last 15 minutes and designate their network traffic as "Best Effort." Traffic from all users will normally be designated as "Priority Best Effort." Traffic that has been designated Best Effort will be handled with lower priority in the upstream and downstream. This basically means that the heaviest users on the Comcast network can use the network as much as they like when the network isn't congested. But when many people want to use the network and the network is congested, the heaviest users will be slowed down to give everyone else a fair shot at accessing the shared network. The Comcast network's periodic designations of Best Effort and Priority Best Effort last 15 minutes at most. If a heavy user of the Comcast network stops being a heavy user in the next 15-minute interval and his or her average bandwidth comes in line with most other users of the network, that user's traffic is again designated Priority Best Effort.

To illustrate how Comcast's Fair Share system for ensuring its customers' fair access to bandwidth works, consider two people on a bus, where Person A is taking up four seats and Person B is standing for 15 minutes or longer. If Person B doesn't care to sit down, Person A can take the four seats as long as he wants. But if Person B wants to sit down, he has the option of taking up to three seats for 15 minutes while Person A pulls back to a single seat. If 15 minutes goes by and both people still want the maximum number of seats, each person gets equal priority and each person takes two seats. If one person gives up some or all of his seats, the other person gets the remaining seats.

On a dumb unmanaged packet-switching network that relies solely on Jacobson's TCP control, the heaviest users always get the highest priority because of the fact that the multi-flow applications they are using are immune to the bandwidth allocation scheme built into TCP. The dumb packet-switching network is the least fair system because it allows heavy users take resources at the expense of other users. Under a fair share network management scheme, everyone gets an opportunity at high bandwidth.

PROTOCOL-SPECIFIC NETWORK MANAGEMENT SOLUTIONS

Protocol-specific and protocol-agnostic approaches to the management of packet-switching networks should never be viewed as contradictory because they're meant to solve entirely different problems. Protocol-agnostic systems for network management are ideal for distributing bandwidth equitably among multiple broadband customers, but they can never ensure the equal performance of different applications sharing the same broadband connection and they don't deal with jitter (a measure of the variation in packet delay) on every segment of the Internet. Protocol-specific network management systems such as Quality of Service (QoS) are needed to ensure fairness and harmony between users sharing the same broadband connection.. There are also instances where protocol-specific jitter management is necessary for shared network links between multiple broadband customers. This does not conflict with protocol-agnostic systems so long as equitable sharing of bandwidth between users is maintained.

Many proponents of net neutrality claim that protocol-specific network management such as QoS discriminates against applications designated with lower priority. This argument incorrectly assumes that all applications require the same performance metrics and that all applications have an equal chance for success on a dumb packet-switching network. There are many instances where protocol-specific management schemes are perfectly justifiable and desirable.

Some applications throttle themselves and only ask for very little bandwidth while other applications can take as much bandwidth as possible at the expense of self-throttling low-bandwidth applications. A good network management system will ensure that the self-throttling applications get what little bandwidth they're asking for and will prevent bandwidth-aggressive applications such as P2P from drowning the self-throttling applications out.

In the case of most traditional broadband technologies—cable, Digital Subscriber Line (DSL) over copper phone lines, and fiber—high jitter problems are almost entirely isolated to the customer who created the network congestion in the first place. In other words, the harm mostly comes to the customer whose application initially created the problem. Those customers that are creating jitter within their own homes would want their broadband provider to use protocol-specific network management techniques applied to their broadband connection to give them a better experience when they simultaneously use multiple applications. On wireless networks with much more limited shared bandwidth or networks offered by smaller rural ISPs with limited backhaul connectivity, however,

the jitter created by one customer spills over to other customers. In situations where a broadband network does not have good jitter isolation between customers, protocol-specific jitter management techniques that work across multiple broadband customers are perfectly justifiable. Thus, a blanket rule that prohibits protocol-specific network management techniques between broadband customers would have undesirable consequences.*

At times, some broadband customers may want lower priority for certain high-volume protocols, especially if the ISP is willing to be generous with volume in exchange for lower priority. Lower priority should not be confused with bandwidth throttling because it typically only means a small drop in average bandwidth performance—something that is a great tradeoff for high-volume application if it translates to cheaper volume. Although a broadband pricing model that offers greater volume in exchange for lower priority isn't common yet, it's a fair and attractive pricing model, and public policy should not rule it out with blanket prohibitions against protocol-specific bandwidth management. Such a model is identical in concept to the choice of lower priority FEDEX package shipping if it means an attractively cheaper shipping price. If broadband consumers are voluntarily labeling their high-volume packets with a low priority with BitTorrent applications like Vuze, which already support the P2P standardized scavenger mode, then there's absolutely nothing wrong with the broadband provider honoring that request by giving that packet lower priority.

LOGICAL ORDER OF PACKET PRIORITY FOR APPLICATION TYPES ON THE INTERNET

Four basic types of applications run on the Internet: (1) Platinum—real-time applications such as VoIP, online gaming, video conferencing, and IPTV; (2) Gold (buffered video streaming applications ranging from YouTube to Xbox HD); (3) Silver (interactive applications); and (4) Bronze (background applications such as BitTorrent and Kazaa) (see table 1). Each general type of application has distinctly different needs and impacts on the network. Some applications simply grab as much bandwidth as the network can supply even if they disrupt jitter-sensitive applications that demand very little bandwidth. These bandwidth-bursting applications are not being intentionally malicious; it's just how some applications work. There's nothing wrong with having them so long as the network infrastructure can adjust for them and protect the jitter-sensitive applications.

The problem with the "dumb" packet switching network is that if one application decides to use the network for a prolonged period of time by bursting a large number of packets all at once, other applications can be starved for the duration of that time.

A smart and well-managed network will attempt to simultaneously satisfy all four types of applications as best as possible. For most home broadband customers on current generation broadband networks, a smart network will be able to simultaneously support three of the application types if the bandwidth is low. That means low bandwidth real-time applications like VoIP and online gaming, interactive applications like Web surfing, and even P2P background applications. A dumb network on the other hand can only support interactive and background applications with sluggish interactive performance but real-time applications will suffer badly.

What a smart network can't do is substitute for a next-generation broadband infrastructure because some of the applications like high-definition (HD) video conferencing, IPTV, and many of the other thousand Kbps applications will simply not run well or not run at all. This is why it's crucial that public policy should push for the adoption of next-generation broadband services. Still, it is important for policy makers to recognize that no amount of raw bandwidth will ever be a substitute for intelligent networks and network management and QoS technology. Progress toward a ubiquitous digital world requires both bigger pipes and better managed pipes.

Even with larger pipes, some applications still could have problems. This is why there is a logical order of

*Protocol-specific network management techniques are consistent with FCC Chairman Martin's August 1, 2008, ruling on Comcast, where he stated that prioritizing VoIP, a protocol-specific network management technique, was justifiable. Chairman Martin made no mention or distinction between network management techniques that apply between broadband customers or between the same home and no policymakers or regulatory agencies have gone in to this level of detail.

packet priority that ensures the best compromise for all applications sharing a network to work simultaneously as well as possible. The logical order of packet priority is as follows:

- Real-time applications that are most sensitive to jitter (e.g., VoIP, online gaming, IPTV, Video conferencing)
- Video streaming applications with moderate fixed bandwidth requirements and moderate jitter tolerance (e.g., video streaming applications like YouTube, Xbox live, Hulu, Netflix)
- Interactive applications that have brief bursts in bandwidth that could disrupt real-time or streaming applications if they were given a higher priority (e.g., web browsing, email).
- Background applications, which by design are unattended and for whom no human is looking at the application waiting for an instant response. (e.g., P2P applications, any other bulk file transfer technology)

P2P applications can grab 10 to 40 times more bandwidth using the P2P architecture's multiflow advantage and P2P applications can impose high jitter on the network at the expense of all other application types. This means that unless P2P applications are deprioritized, every other application suffers. Furthermore, the graph below shows that giving a P2P application the lowest priority does not affect the performance of the application one bit.

Although it's true that P2P bandwidth under the smart network was slowed down nine times more to make room for the interactive Web traffic, the fact that the Web browsing traffic runs nine times faster also means that it gets out of the way nine times sooner and the P2P bandwidth resumes full speed sooner. That means over the course of the file download, the time it takes to complete a P2P file transfer is unchanged. Under the smart network, P2P performance stays the same but interactive traffic goes up nine fold in performance (see figure 6). In the context of network management within a home under a

Table 1: Network Requirements of the Four Basic Types of Applications That Run on the Internet

	PLATINUM (REAL-TIME APPLICATIONS)		GOLD (BUFFERED VIDEO STREAMING)		SILVER (INTERACTIVE APPLICATIONS)		BRONZE (BACKGROUND APPLICATIONS)	
	LOW PACKET DELAY LOW TO MEDIUM BANDWIDTH		MEDIUM PACKET DELAY MEDIUM BANDWIDTH		HIGH BANDWIDTH LOW VOLUME		HIGH VOLUME HIGH AVG BANDWIDTH	
Examples	Application	Bandwidth	Application	Bandwidth	Application	BW	Application	BW
	Voice over Internet Protocol (VoIP)	30-90 Kilobits per second (Kbps)	YouTube	320 Kbps	Web browser	Burst	BitTorrent	Peak
	Online gaming	30-90 Kbps	YouTube High Quality	650 Kbps	Email	Burst	LimeWire	Peak
	Video Conferencing	250-8000 Kbps	Vimeo	500 - 1200 Kbps	News reader Network News Transfer Protocol (NNTP)	Burst	Kazaa	Peak
	Internet Protocol-based TV (IPTV) (last mile only)	2000-8000 Kbps	Netflix	4000 Kbps			WinMX+share	Peak
			iTunes "HD"	4000 Kbps				
			Xbox "HD"	6800 Kbps				
							Winny	Peak

single broadband account, this priority scheme needs no justification. In the context of network management between different broadband customers, lower priority for background traffic would be justifiable and even desirable to P2P users if the low-priority traffic got more generous volume caps or cheaper metering rates in return.

When background applications are given lowest priority, their average performance is hardly impacted. Furthermore, such applications are far less likely to get shut down or severely throttled by the P2P user because the user isn't worried about impacting other applications in their home. Without the logical packet priority order shown above, consumers will adopt the least desirable network management solution by shutting down their own P2P application whenever they want to make a VoIP call, play an online game, or do some heavy Web surfing.

THE PROBLEM OF JITTER ON PACKET-SWITCHING NETWORKS

Three basic concepts of network performance are used to determine how well different applications work on the Internet—bandwidth, volume, and delay (which includes latency and jitter) (see box 5).

A. Queuing Theory's Application to Packet-Switching Networks

The price of packet-switching networks' flexibility and efficiency is less predictability in packet delay and higher jitter. Queuing theory is the mathematical study of queues and has applications in many fields, among them as key mathematical background to packet switching, the basic technology behind the Internet. Important work on queueing theory used in modern packet-switching networks was pioneered by Dr. Leonard Kleinrock at the University of California at Los Angeles in the 1960s. Such work was important to the development of ARPANET and the Internet.

Queuing theory is applicable to the study of packet delay on packet-switching networks such as the Internet. It also mathematically explains the cause of jitter on packet-switching networks. The fundamental problem is that packets traveling from a faster network to a slower network or packets merging from multiple network links on to a single network always have the potential of hitting a queuing delay. This is no different from a situation in which cars merging from a five-lane freeway to a three-lane freeway experience a traffic pileup, and a pileup can even happen during off-peak hours if many cars just happen to show up at the same

Figure 6: Dumb versus smart network

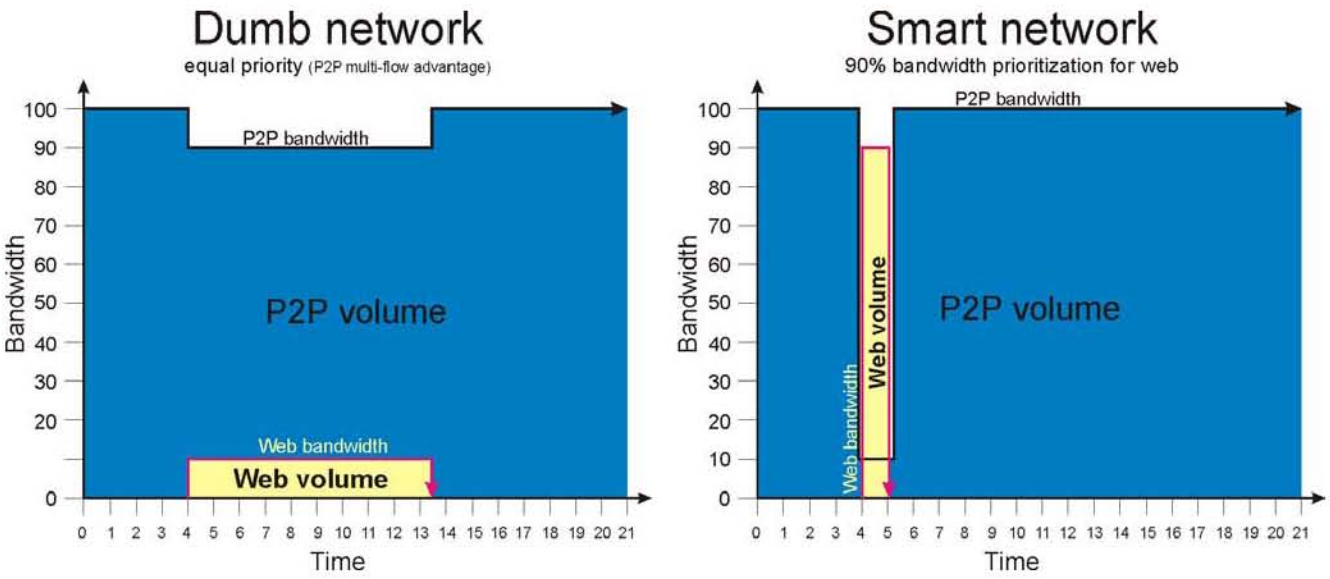

BOX 5: BASIC CONCEPTS OF NETWORK PERFORMANCE

Network management is a very complex topic, and different metrics are used to determine how well different applications work on the Internet. Three fundamental concepts of network performance are critical to understand: bandwidth, volume, and delay (which is further broken down by jitter and latency). These basic concepts of network performance are defined below.

Bandwidth is the rate at which files are transmitted through a network commonly referred to as "speed" but more accurately described as throughput. There's also a tendency to describe higher bandwidth as "fat pipes" when in reality they're not actually fatter; they simply deliver more bits per second. But the concept of a fat pipe is a useful visualization to help describe higher throughput networks. Files (a movie or a song for example) consist of many packets, and the file transfer rate is determined by the size and rate of packets flowing through a network per second. Bandwidth determines the time it takes to transfer a file and how long a person has to wait for a file (e.g., video) to download.

Volume is defined by duration multiplied by the average bandwidth over that duration of time. Simply put, volume is the number of bytes consumed over an arbitrary period of time. Many ISPs, especially outside the United States, limit the volume that a customer can consume per day or per month. Volume is frequently confused and mislabeled as bandwidth. Some applications can be very high bandwidth but very low volume because they're low duration while other applications can be medium bandwidth but high volume because they operate continuously for long periods of time.

Delay is the time it takes an individual packet to travel through a network is also commonly thought of as "speed." The smaller the delay, the better for real-time applications. Packets consist of individual bits, but the packet is considered the basic building block on a packet switching Internet Protocol (IP) network because it contains all the addressing information and IP networks don't deliver individual bits. A single packet might experience more delay and take longer to go from one computer to another because the distance between the computers is great or there are network devices along the way that are backlogged where the packet have to sit and wait. The two types of delay are latency and jitter.

- **Latency** is a simple measurement of delay and the word is commonly misused to describe both latency and jitter. Latency on a computer network is actually the time it takes a bit or packet to traverse a noncongested network before arriving at its intended destination and it's generally measured in milliseconds (ms) where 1,000 ms equals one second. The typical latency from the east coast of the United States to the west coast over the Internet is approximately 40 ms. Not much can be done about this type of latency because it's largely dictated by the distance and speed of light over a fiber optic glass medium. Common latency metrics such as "ping" measure round-trip time of a packet, which is double the one-way latency. Thus, the ping time from the New York to San Francisco is approximately 80 ms on an uncongested network.

- **Jitter** is the measure of the variation in packet delay. High jitter conditions are essentially micro-congestion storms that last tens or hundreds of milliseconds. High jitter occurs whenever a large number of packets come from a faster network link to a slower network link or where several networks links merge to a single link. When this happens, network devices such as routers and switches get backlogged and they force packets to wait inside their memory buffers, thereby increasing the time it takes packets to traverse a network. If a network fluctuates between 80 and 85 ms of delay, then the jitter has a low magnitude of 5 ms. If a network mostly has delays of 20 ms but occasionally spikes to 220 ms, then the magnitude of the jitter is high at 200 ms. Even if the latter example has better average delay, its high jitter makes it less desirable for real-time applications than a network with higher average delay but lower jitter.

It's crucial to understand that bandwidth, volume, and delay are independent metrics that can operate freely. Many net neutrality proponents mistakenly see packet prioritization for real-time applications as a form of discrimination against file transfer applications because they confuse delay with bandwidth. But peer-to-peer (P2P) file transfer is generally immune to packet delay because a network that has very high delay can still achieve high bandwidth. Conversely, a network with low delay might have low bandwidth because it doesn't transmit a lot of packets per second.

time at the place where lanes merge. On a network such as the Internet, a sudden burst of network traffic can cause a backlog of packets waiting to be transmitted in network devices, and this in turn can result in very high jitter.

Networks of all types and speeds are connected by routers and switches on the Internet. Since packets are actually made of electromagnetic signals flowing through wiring or fiber optic glass, only one packet can flow through a network at any given time without collision. Packet collisions result in the destruction of all colliding packets which requires retransmissions of all those packets and too many of these collisions causes a network to perform very poorly. To prevent packet collisions, routers and switches have packet queues (these are memory banks) that temporarily hold packets that can't be transmitted at once. But because some computer applications tend to burst out large number of packets, especially P2P applications, it's possible to have large queuing delay even under relatively light P2P loads. These queuing delays result in huge spikes in packet delay, which is how high levels of jitter are created on packet-switching networks.

B. The Misperception that Network Jitter Can Be Solved by More Capacity

There is a common misperception that QoS is necessary when a network is busy but that QoS is not needed when network capacity is abundant and network utilization levels are low. The idea that abundant capacity and low utilization can substitute for QoS is fundamentally misguided. In fact, it is possible to have good low jitter conditions on networks operating at 90 percent capacity if the data flowing over that network has packets that are evenly spaced. And conversely, it is possible to have bad high jitter conditions on networks operating at 10 percent capacity if the packets are clumped together. This is precisely why IPTV streams that take most of the capacity on fiber to the node (FTTN) broadband network such as AT&T U-verse cause zero measurable increase in latency or jitter. On the other hand, a 20 percent load from multiple P2P TCP flows can cause lots of jitter and very deep queue depths.

Figures 8, 9, and 10 illustrate the amount of round-trip delay induced by BitTorrent (a P2P application) and a VoIP call with similar bandwidth requirements. Despite the fact that the VoIP call is using a little more bandwidth (11 KB/sec) in both directions, it produces negligible jitter. But minimal usage of BitTorrent operating at 10 KB/sec upstream can create high jitter conditions.

C. Why Broadband Networks Will Always Have Speed Mismatches and Jitter

The nature of broadband networks is such that there will always be large mismatches in speed and multiple networks merging. The home network, where distances are measured in meters rather than kilometers, will always be orders of magnitude faster than the broadband connection to the Internet because it's much cheaper to build faster short distance networks. Home

Figure 7: Why there will always be a bottleneck on broadband

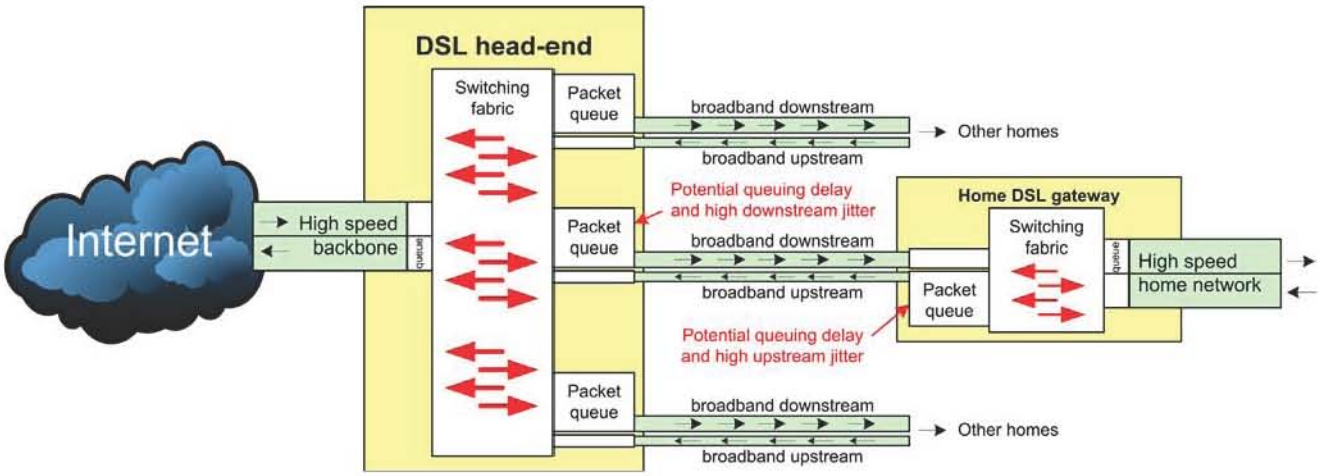

Figure 8: Effect of low upstream BitTorrent usage on jitter

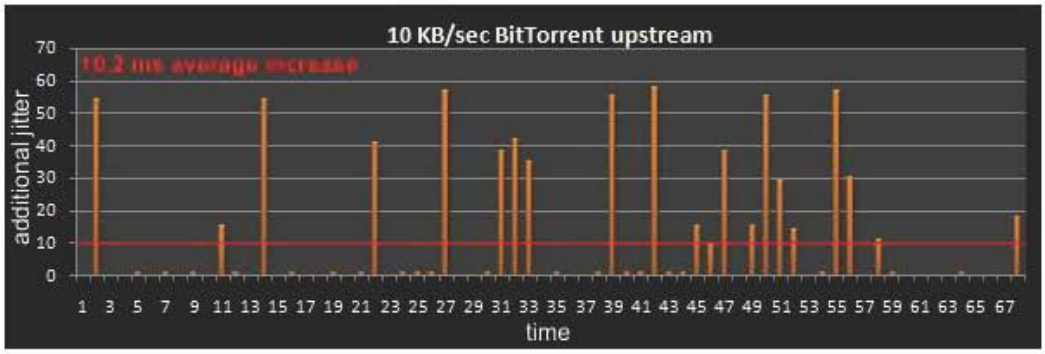

Chart source: Results of tests conducted by George Ou over a residential ADSL broadband connection.

Figure 9: Effect of low upstream and downstream VoIP usage on jitter

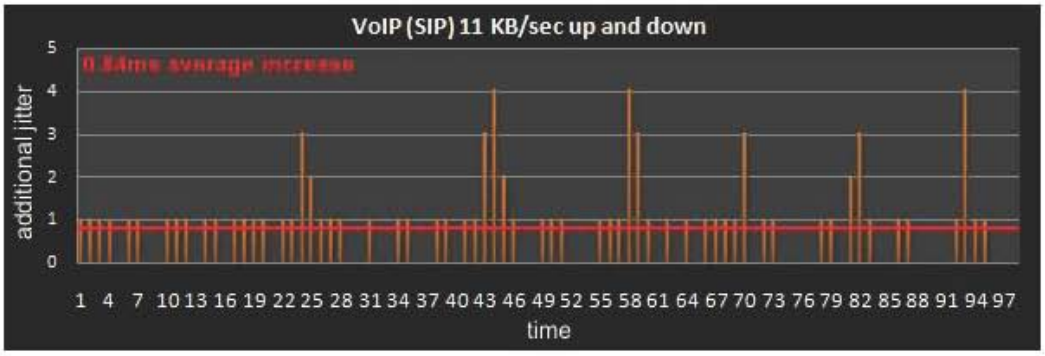

Chart source: Results of tests conducted by George Ou over a residential ADSL broadband connection.

BitTorrent downloading at 260 KB/sec caused even more jitter. The chart below was cropped at 450 milliseconds, but the six spikes shown were actually timeouts, which means the delay exceeded 1000 milliseconds.

Figure 10: Effect of high downstream BitTorrent usage on jitter

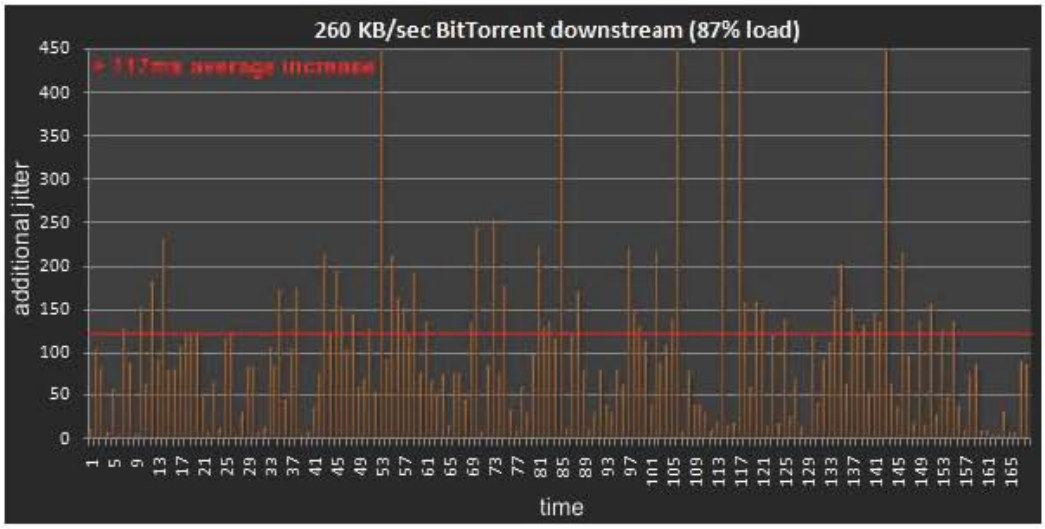

Chart source: Results of tests conducted by George Ou over a residential ADSL broadband connection.

networks, for example, typically operate at 100 Mbps or 1000 Mbps, while broadband networks operate in the single-digit Mbps range or occasionally in the tens of Mbps range. By the time broadband services routinely offer gigabit upstream throughput, home networks will operate at 10, 40, or 100 gigabits, with multiple computers or devices trying to use the Internet at the same time.

This means that regardless of how fast the Internet becomes in the future, there will always be a mismatch in speed coming from the home network to the broadband network and there will always be potential for upstream jitter on the broadband connection whenever applications burst upstream packets.

Downstream jitter is a bigger problem for broadband services. The core interconnects of the Internet must always be orders of magnitudes faster than broadband connections to aggregate traffic from thousands of user. No matter how far technology progresses, the mismatch in speed coming from the Internet backbone on to the broadband downstream will always be present and the potential for downstream jitter will always exist. To eliminate these packet queuing delays and jitter, QoS technology must be employed.

D. Why Certain Applications Create High Jitter and Others Don't

So why is it that some applications create high jitter on a network and others don't despite the fact that they may be operating at around the same speed? Before we can answer this question, we much first debunk a common misconception about networks—namely, that there is such a thing as partial utilization of a network link.

When someone says that a network is experiencing 10 percent utilization, people often imagine a pipe that's filled to 10 percent, with 90 percent of the pipe always available for other uses. Unfortunately, this concept of partial utilization, though a convenient way to visualize a network, is not how packet-switching networks actually work. Network links are either 100 percent utilized and jammed up by one application or they're 0 percent utilized with no traffic at all—and there's never actually an in-between state. Thus, when a network link is being utilized at 10 percent, this means that over time, the average of the 0 percent available states and the 100 percent congested states average out to 10 percent.

Packet switching networks like the Internet by design can only service the packets of one application at any

Figure 11: High jitter inducing application

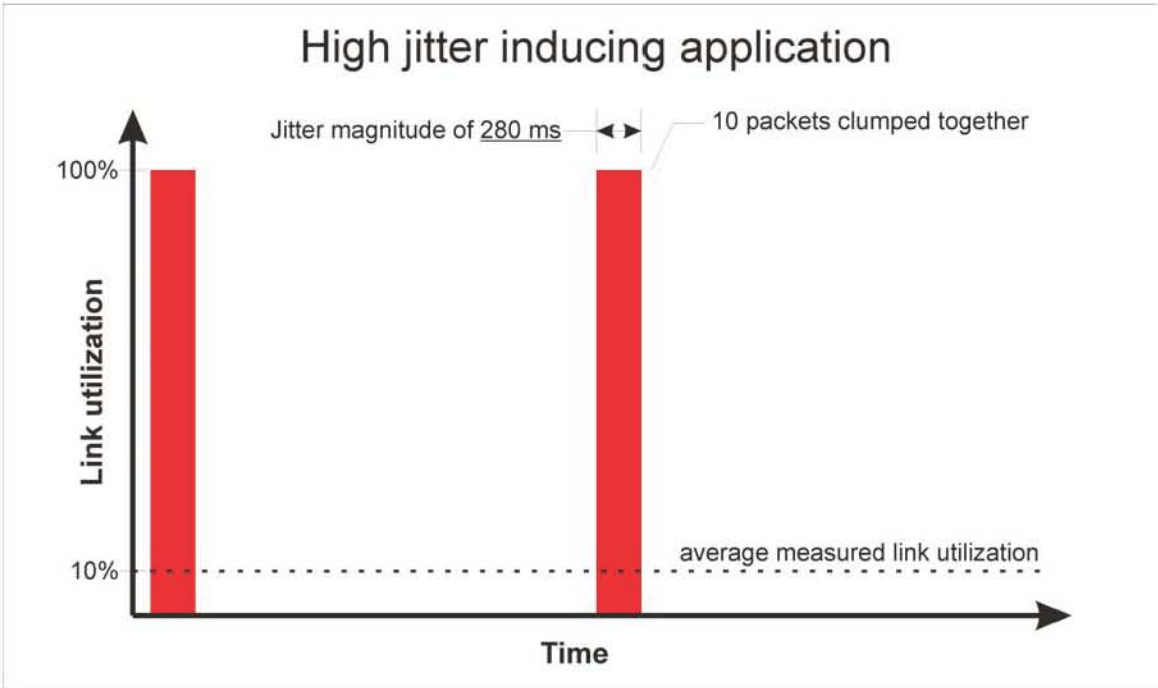

point in time and it supports multiple applications by alternating between the packets. The problem with the "dumb" packet switching network is that if one application decides to use the network for a prolonged period of time by bursting a large number of packets all at once, other applications can be starved for the duration of that time. If applications evenly spread out their packets over time, the jitter remains low and doesn't cause problems for delay-sensitive real-time applications. This high variation of packet delay is called jitter and it is very harmful to real-time applications like VoIP, video conferencing, online gaming, and IPTV. Jitter has minimal impact on Web browsing and almost no impact on P2P file transfer applications.

Imagine a person trying to do two tasks in the home. If the person was cooking food on a hot stove while trying to vacuum the carpet at the same time by switching back and forth between these tasks every hour, the food would probably get ruined even when both tasks are given equal attention time. This is similar to what's happening in a packet-switching network on a millisecond level. The time-sensitive application, like food cooking on a stove, requires consistent and frequent attention. VoIP applications ideally need to be serviced every 1/50th of a second like food on a hot stove may need to be serviced every 5 seconds.

Some applications such as P2P have a tendency to burst out a large number of packets at once, thereby monopolizing a network link for tens or even hundreds of milliseconds, and this is another negative side effect of multiflow applications. Pictured below is an example of an application that sent out a flurry of 10 packets with long periods of rest. This produces very high and undesirable jitter of 280 ms.

Under extreme circumstances, it is possible for P2P applications to cause over 1000 ms latency on the downstream side of a broadband connection. So if a person downloads from 40 P2P peers at the same time and each of those peers send 30 packets at once and they all happen to converge on the last mile over your broadband link at the same time, the downstream link can easily be jammed and unavailable to other applications for more than one second, which is massive jitter.

Other applications such as VoIP, online gaming, and IPTV tend to send out packets that are perfectly spaced one packet at a time. Applications like these can drive utilization up to as high as 90 percent and still keep jitter to an absolute minimum. This is the ideal scenario for minimizing jitter where the packets are perfectly uniform and fine-grain.

Figure 12: Low jitter inducing application

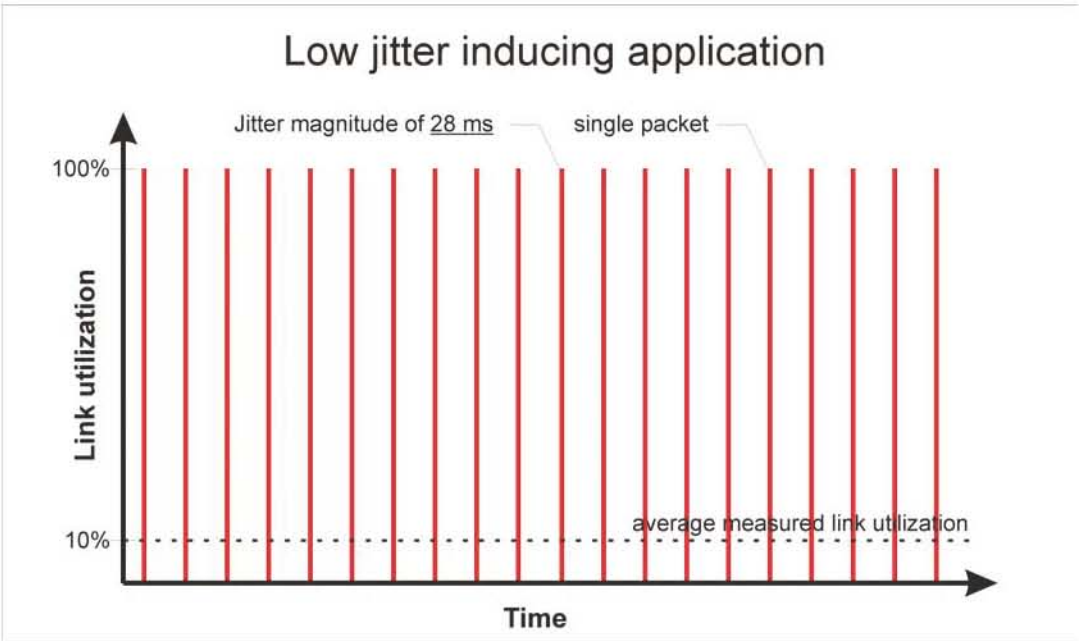

VoIP applications are a classic example of a jitter friendly application. Figure 13 represents VoIP packets utilizing a network at even bursts of single packets with a 20 ms interval, which is 50 packets per second. This is called "isochronous" data transfer, which is a way of transmitting data with periodic small bursts of data. Commercial VoIP services at most use 88 Kilobits per second (Kbps), and this is roughly 20 percent utilization on the upstream side of a residential DSL broadband connection.

Mild P2P usage can cause high jitter when the P2P application bursts out three large packets at a time. Five VoIP packets normally scheduled to be delivered every 20 ms are displaced by 85 ms, which causes three of the packets to be discarded because they arrived too late. That in turn results in 6 percent packet loss for the VoIP application, a loss that represents a small but noticeable decline in quality.

Figure 15 illustrates what can happen if more packets are clumped together. This particular example shows 13 packets being lost, a figure that translates to 26 percent packet loss. This results in a severe decline in phone quality. P2P downloads can generate even higher jitter with spikes going above 1000 ms. There are jitter adaptation techniques where VoIP applications will increase their buffer size to reduce packet loss, but only so much that can be corrected under high jitter conditions and quality still suffers.

QUALITY OF SERVICE (QOS) AND THE INTERNET

QoS in the context of packet-switching networks such as the Internet is a protocol and application specific form of traffic engineering. QoS is not just one technology, it is a complex field of study with dozens of Internet Engineering Task Force (IETF) standards that compete with or complement one another.

The IETF's QoS standards generally fall under two categories of Integrated Services (IntServ) and Differentiated Services (DiffServ). IntServ is a more complex scheduling system that requires resource reservation. Although it offers precise fine-grain control of resources, the added complexity of reservation setup makes IntServ unlikely to scale on large networks much less the Internet. DiffServ is a simpler coarse-grained mechanism that classifies traffic types. Although it does not have the precision of IntServ mechanisms, DiffServ also does not require complex reservation setup and tracking. DiffServ also offers enough control to meet most requirements, making it popular in the marketplace.

DiffServ QoS mechanisms basically create multiple packet queues for different types of applications whereas the dumb packet-switching networks only have one packet queue. By having multiple packet queues, QoS-capable network devices can transmit packets in a more granular way where no single application can

Figure 13: How VoIP packets flow

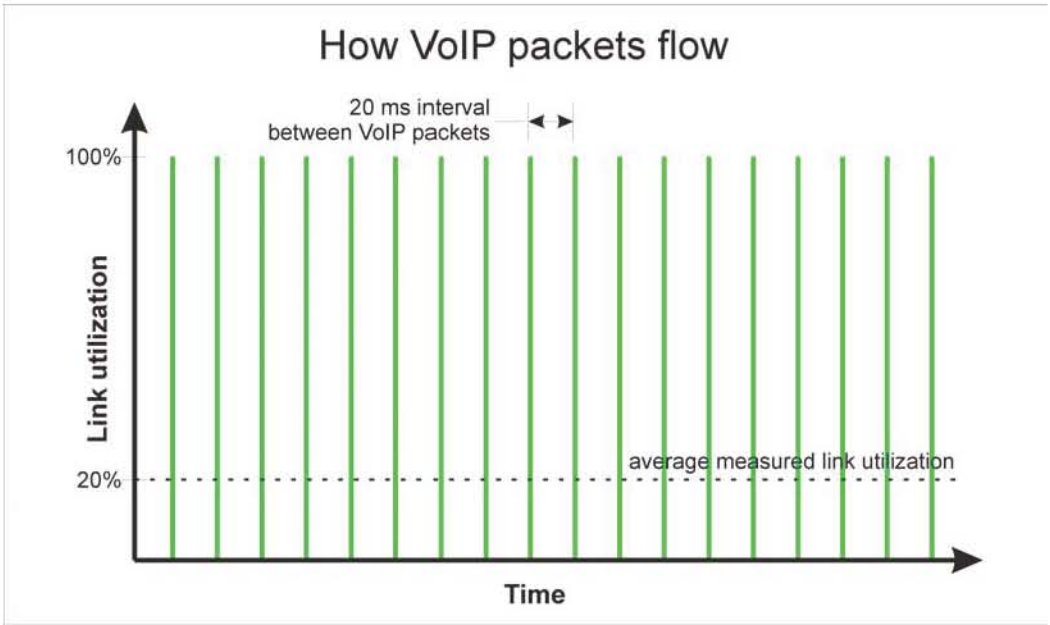

Figure 14: VoIP dealing with low jitter still suffers

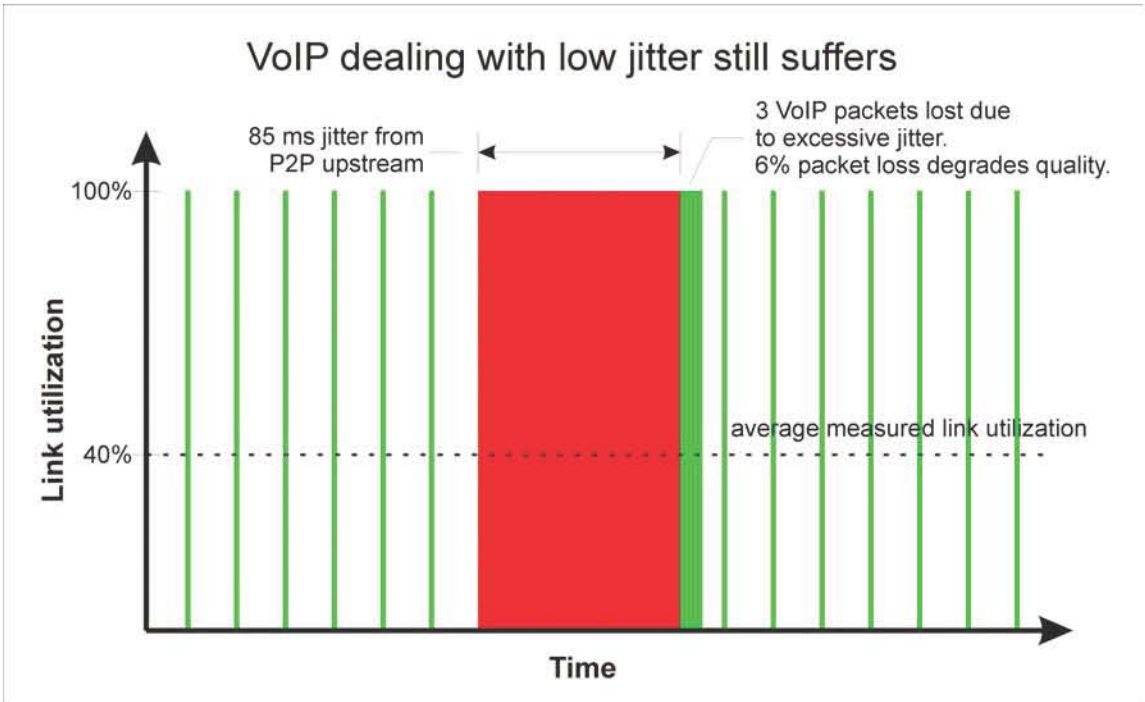

Figure 15: High jitter is much more destructive to VoIP

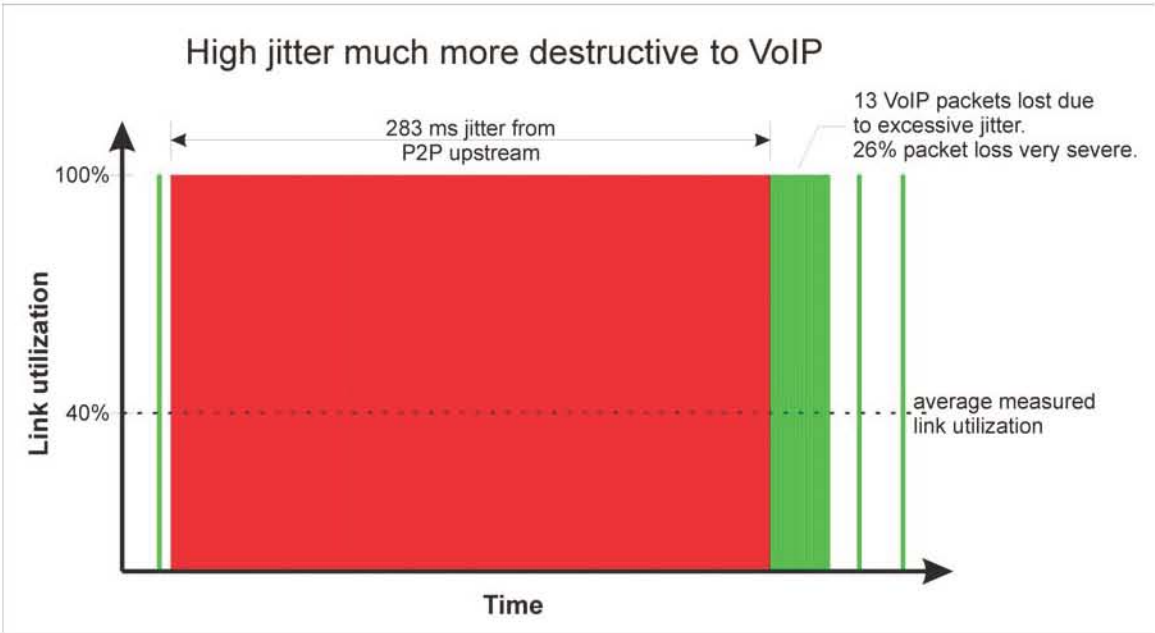

monopolize a network link for too long. Even when the real-time voice or video packet queue isn't given any more priority than the file transfer queue, just the mere act of alternating packet transmissions between different applications can substantially lower jitter. When the priority level of the real-time queue is given a higher priority, jitter for real-time applications can be reduced even further.

A. Solving the Jitter Problem with QoS

Using DiffServ QoS technology, network devices can rearrange and shuffle the packets to minimize jitter. This is not about giving any application a higher bandwidth rate; it's about ensuring that no application gets unattended for a prolonged period of time like the food cooking on the hot stove analogy. A separate queue is created for VoIP such that VoIP packets can be transmitted in-between the larger clumped peer-to-peer packets. This advanced form of QoS can completely eliminate packet loss for real-time applications while allowing file transfer applications to take as much bandwidth as they like.

A cruder QoS technique used inside some consumer devices or software merely limits packet bursting which actually harms file transfer speeds while providing minimal jitter relief. While this provides some minimal relief, it's a poor substitute for advanced QoS technologies that reorder and shuffle packets (see figure 17). Advanced QoS techniques that reorder packets optimally can almost eliminate jitter entirely while providing maximum concurrent file transfer performance.

Another crude way of obtaining QoS is to overprovision a network and maintain a very low utilization level on the network by using pricing mechanisms. This unfortunately means that the network is grossly underutilized, doesn't have minimal jitter, and isn't as cheap and fast as it could be with a QoS enabled network.

B. Clearing up Misconceptions About QoS

There are a number of misconceptions about QoS technology. One misconception about QoS prioritization is that lower priority applications are somehow forced on to a "dirt road" that runs slower. This is false because bandwidth is generally not affected by higher packet delay. Whether someone downloads a file from 50 miles a away with 20 ms latency or 500 miles away with 40 ms latency, a doubling of packet delay does not result in a halving of bandwidth. Only the total delivery time of the file being transferred goes up by 20 ms, which is 1/50th of a second, which is imperceptible.

For more complex reasons, higher latencies do impose lower speed limits for applications that use Transmission Control Protocol (TCP). This can be overcome by using User Datagram Protocol (UDP) in place of

Figure 16: Network device with QoS

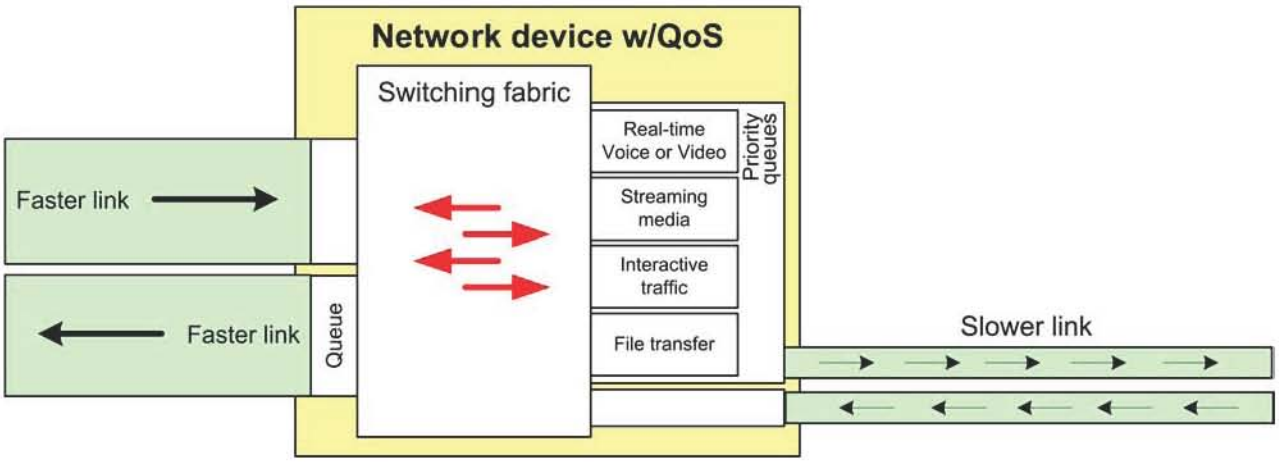

TCP, using a modified TCP implementation, or using multiple TCP flows which is common with P2P applications. Jitter, which is the type of delay that QoS tries to minimize, does not impose a speed limit on TCP. Generally speaking, not much can be done to reduce latency because latency is usually determined by the speed of light and the distance between two endpoints.

Another misconception about QoS is that high-priority applications are given lower jitter delay at the expense of higher jitter for low-priority applications. This idea is based on the misguided assumption that if we take away 40 ms of packet delay for the high-priority application, then we must be adding 40 ms of delay to the low-priority application. QoS doesn't work that way because the network is simply switching between the various applications more frequently so that all priority tiers end up with less packet delay. Every application regardless of whether it is high priority or low priority benefits, but the high-priority applications get the most benefit.

The argument is often made that giving priority to VoIP is somehow unfair to the P2P applications because VoIP packets are transmitted ahead of the P2P packets, and giving priority to VoIP is tantamount to letting VoIP "cut in line." This argument ignores the fact that P2P already gets many orders of magnitude more bandwidth and volume than VoIP. Furthermore, insisting on equal P2P priority ends up being destructive to P2P because lower priority P2P packets are effectively a worry-free license to run P2P applications at full throttle.

Without a good QoS solution in place, consumers today must consciously shut off P2P or severely limit its throughput whenever they want to use VoIP or do online gaming despite the fact that there's plenty of leftover bandwidth. With QoS in place, a consumer can run both applications concurrently with no bandwidth limits on P2P while achieving optimum performance for everyone. By accepting lower packet priority, thereby causing less jitter and less grief for other applications, P2P applications are rewarded with higher bandwidth. This is the reason why P2P standards bodies and P2P vendors are voluntarily standardizing the labeling on their own P2P packets with the "scavenger service" tag, which means lowest packet priority.

C. QoS for Broadband Networks

Implementing a complete QoS network management approach requires participation on both ends of a broadband link because jitter happens in both directions. On a DSL broadband network, the problem is ideally dealt with in the DSL modem in the home

Figure 17: QoS can completely mitigate jitter damage

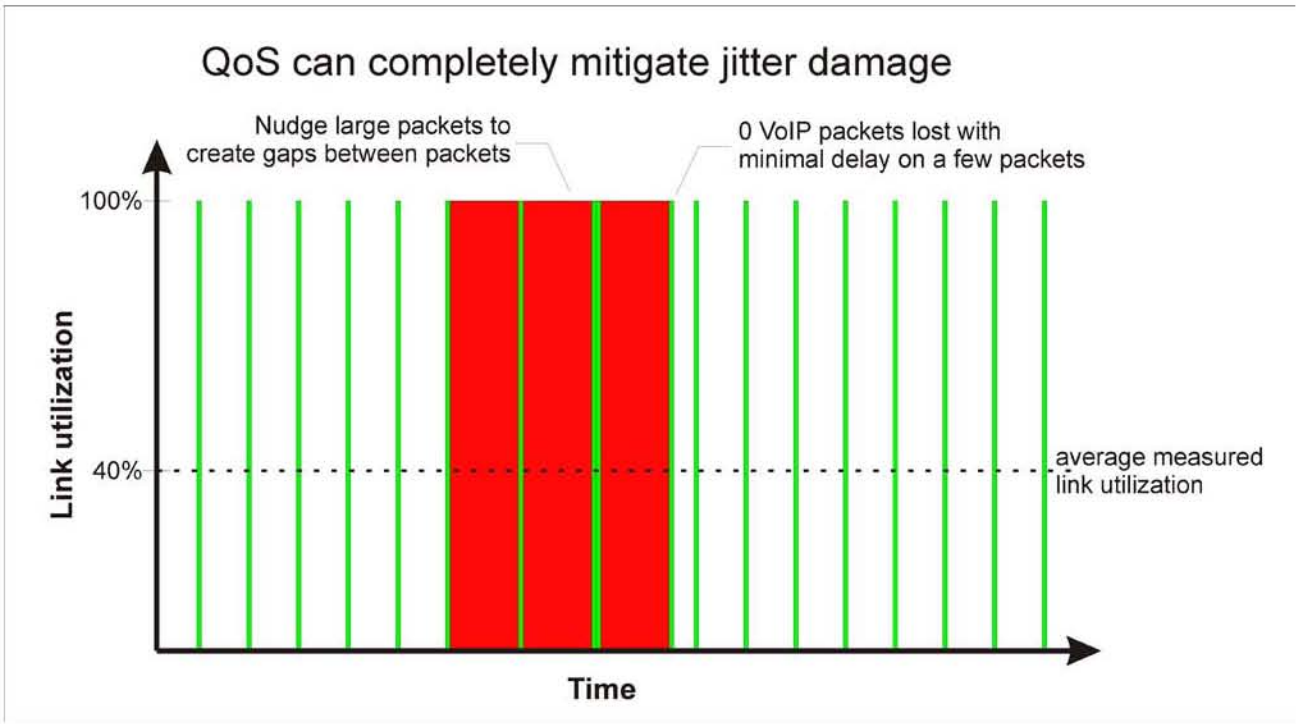

and the Digital Subscriber Line Access Multiplexer (DSLAM) on the ISP's side because queues build up in both of these places. On a cable broadband network, the problem is ideally dealt with on the cable modem in the home and the Cable Modem Termination System (CMTS) on the ISP's side because queues build up in these places.

A partial QoS solution can be implemented by the consumer in the home using high-end routers. However, this solution can only eliminate upstream jitter and not downstream jitter because downstream jitter can only be dealt with at the ISP side in the DSLAM or CMTS where downstream packets queue up. The vast majority of consumers don't even have these expensive high-end home routers with effective QoS technology. Even when they do, those high-end home routers can only partially reduce downstream jitter using the crude method of slowing down the P2P application. This is a horrible tradeoff because jitter isn't entirely eliminated and the P2P application is unnecessarily slowed. If the problem is dealt with by both the ISP and the home, downstream and upstream jitter can be completely eliminated and bandwidth hungry applications get more bandwidth.

D. User-Approved and User-Controlled QoS

Unfortunately, it is unclear whether the FCC's ruling on August 1, 2008, admonishing Comcast has made protocol-specific network management schemes illegal. The unclear nature of the FCC's ruling does make U.S. ISPs worry about another lawsuit or FCC complaint if they try to provide good application multitasking to customers. Some net neutrality advocates have exacerbated the confusion by attempting to paint anyone advocating both protocol-specific and protocol-agnostic solutions for network management as being hypocritical when the reality is that the two technologies are addressing completely different problems.

Comcast at this point has not attempted to perform protocol-specific network management yet. It has solely concentrated on network management solutions that assure equitable bandwidth sharing between customers. It could at some point in the future implement a protocol-specific network management system that addresses the problem of jitter and makes broadband more multi-application friendly with the user's consent or even allow users to configure it themselves. It is unclear if this would be an opt-in system or an opt-out system. Because QoS is such a complex technology that is difficult for many consumers to understand, though, it would benefit the most number of consumers if the system was activated with a default set of logical priority rules that users can opt-out of or adjust.

WIRELESS NETWORKS: THE NEW FRONTIER OF THE INTERNET

In a world where wired broadband such as DSL, cable, and fiber are the last mile of the Internet, wireless technology is becoming more important, both within the home for the "last meter" and outside the home as new high-speed wireless technologies like WiMAX and Long-Term Evolution (LTE) are rolled out. In fact, it's easy to envision a day when wireless broadband access will surpass wired broadband services because the total market for residential broadband service is limited to the number of households whereas the total wireless broadband market is limited by the number of future Internet-enabled mobile phones. To realize this future, wireless network management is more important than ever.

A. Why Wireless Networks Require More Management than Wired Networks

Just as they argue that more capacity on wired networks obviates the need for network management, some net neutrality advocates claim that if the supply of spectrum is increased then we won't need any kind

Figure 18: The need for QoS on both ends of the braodband competition

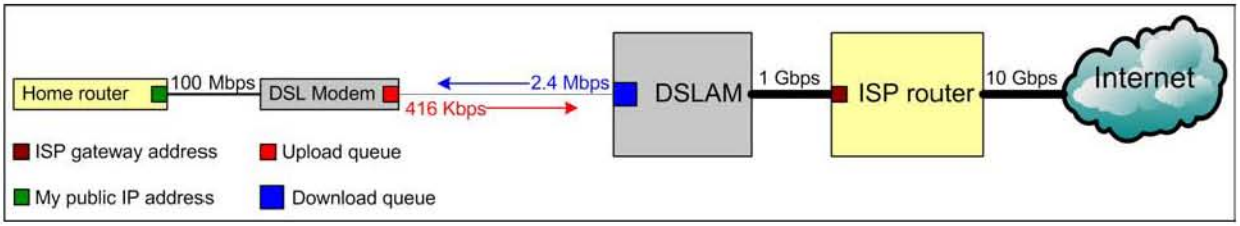

of wireless network management. The reality is that unmanaged networks don't work for wired networks or wireless networks. Not only is the capacity on wireless networks more scarce; wireless networks are far more shared than wired networks, which presents unique engineering challenges not present on copper DSL or fiber networks.

Since wireless devices share their own radios sharing the same air space and frequency, wireless networks require much more sophisticated network management scheme than wired networks. This is because when multiple computers try to send packets out on the same radio frequency at the same time, colliding packets are destroyed and must be retransmitted–an inefficient process. Wired networks have routers and switches that prevent packets from colliding on the same wire by using memory buffers called packet queues but there is no such luxury in the wireless space.

The following table shows the amount of bandwidth that each wireless technology is capable of. It's important to understand that the amount of bandwidth listed is shared bandwidth per radio which typically services 100 to 1000 people. Sharing radios between so many customers is necessary because the average 3G cell tower today costs an average of $650,000. With today's 3G technology, the bandwidth is not only more limited that wired broadband services, it's also shared between more people. By 2009 with the development of next generation wireless data systems like WiMAX and LTE, wireless technologies may close in on current generation cable broadband services in terms of bandwidth but the number of people sharing that bandwidth remains higher so that bandwidth needs to be carefully managed.

B. Increasing Spectral Efficiency Through Scheduled Access

Multiple wireless devices on unmanaged wireless networks have to greatly reduce their rate of transmission such that there are few enough packets flying in the air that the packet collision rate is kept to a minimum. But on intelligent wireless networks where transmissions are coordinated by software that centrally schedules packets, the transmission rates can be kept much higher without the potential for packet collisions. The centralized radio scheduler can be thought of as an air traffic controller for packets just like airplanes need to be centrally scheduled and coordinated to avoid mid-air collisions.

Figure 19: Wireless networks – the new frontier of the Internet

Table 2: Emerging Wireless Technologies

YEAR	TECHNOLOGY	BANDWIDTH (MBPS)		LATENCY
		UP	DOWN	(MS ROUND TRIP)
2007	3GPP R5 – HSDPA	0.375	14.4	150
2007	EVDO Rev A (5 MHz)	7.2	12.4	100
2008	WiMAX (10 MHz)	8	40	60
2009	3GPP R7 – HSPA+	22	42	90
2010	LTE (20 MHz 2xMIMO)	50	150+	20

Unmanaged wireless network technology was used in first-generation 802.11 Wi-Fi technology. The efficiency of the network wasn't great but it was deemed "good enough" for simple data applications even with the packet collision overhead. But as Wi-Fi usage grew in enterprise-class deployments such as hospitals, universities, and corporate campuses and the network started handling telephony applications, good enough for data was no longer good enough for voice. Newer amendments to the 802.11 standard such as 802.11e added two additional network management modes to Wi-Fi to make it more conducive to real-time applications by scheduling around collisions.

With traditional 802.11b Wi-Fi technology, a single Wi-Fi Access Point which uses 20 MHz of radio spectrum can reliably support four Wi-Fi VoIP phones. Any more than four phone sessions on a single Access Point and all the calls begin to rapidly degrade to the point where no one can make a call. The more common and less sophisticated 802.11e mode called Enhanced Distributed Channel Access (EDCA) can triple call capacity. The more advanced form of 802.11e called Hybrid Coordinator Function Controlled Channel Access (HCCA), which uses an advanced centralized scheduling mechanism can potentially increase call capacity tenfold to 40 calls. On the other hand, simply enhancing the signaling speed and quadrupling bandwidth using 802.11g radios without any type of network intelligence may only double call capacity to eight phone calls. The ideal network would increase signaling speed and incorporate centrally scheduled access to maximum call capacity and performance for every user on the network.

The challenges on a small Wi-Fi network are multiplied on a wireless phone and wireless broadband network because there are hundreds of times more devices being served per cell by the network. Voice call capacity and data efficiency on a wireless phone and wireless data network are critical factors in lowering prices for consumers and driving Internet usage. A single wireless broadband access point must simultaneously maximize call capacity and maximize throughput for data applications for hundreds of customers. To increase capacity and performance further, wireless phone and broadband operators are looking to fourth generation (4G) mobile communication standards such as Long Term Evolution (LTE) because of the enhanced speed and advanced network management technologies. Other 4G mobile communication standards such as Ultra Mobile Broadband (UMB) can handle more than 500 calls in just 10 MHz of spectrum which is 250 times more efficient in capacity than unmanaged 802.11b Wi-Fi networks.

C. Why Wireless Management Is a Necessity That Enables Innovation

Many net neutrality advocates claim that if we just had a few unmanaged Wi-Fi Access Points, VoIP applications like Skype would make traditional mobile phone operators obsolete. The reality is that the spectral efficiency of an unmanaged Wi-Fi network and unscheduled VoIP applications are hundreds of times less efficient than the latest managed mobile communication standards used by cell phone providers in terms of call capacity. If we switched to this idealized world of unmanaged wireless networks, the cost per user would be extremely high. While Wi-Fi and VoIP can handle small office needs or even large enterprise deployments when 802.11e is deployed, it does not scale on a metropolitan level.

Many net neutrality advocates also suggest that requiring devices to ask an intelligent network for permission

to transmit data potentially allows network operators to censor speech and stifle innovation. But this type of network intelligence is merely giving the endpoints a way to go faster by avoid data collisions and it has absolutely nothing to do with stifling innovation or censoring the public. The network software grants permission to all paying customers equally to increase spectrum efficiency and everyone benefits with lowers prices, higher call capacity, and higher performance. Intelligent wireless networks will ultimately spurs more adoption and usage of wireless broadband, which facilitates more mobile e-commerce, which enables more innovation and wealth generation. Network intelligence isn't the enemy of innovation because it enables more innovation, but a ban on network intelligence would reduce wireless network quality.

FLAWED ARGUMENTS ABOUT ALTERNATIVES TO INTELLIGENT NETWORK MANAGEMENT

The proponents of net neutrality have suggested several potential alternatives to intelligent network management. As described below, many of the ideas they have proposed, including simply expanding network capacity, are flawed.

A. Why Increasing the Supply of Bandwidth Won't Solve the Problem

One of the most common arguments made by net neutrality advocates is that with sufficient bandwidth, the Internet could simply be a "dumb fat pipe" that didn't require any network management. A key source for this view is a technical paper by Shalunov and Teitelbaum entitled "Why Premium IP Service Has Not Deployed (and Probably Never Will)." As explained in Box 6, the conclusions of this paper are not supported by the facts.

The argument that with sufficient bandwidth, the Internet could simply be a "dumb fat pipe" that didn't require any network management is fundamentally flawed in many ways. First, the argument assumes that bandwidth demand is constant or finite when it isn't. In reality, consumer demand for ever more capacity and bandwidth will outstrip supply for the foreseeable future.

This point is illustrated by the example of Japan with its massive deployment of 100 Mbps fiber to the home (FTTH). The Japanese Ministry of Internal Affairs and Communications conducted a study and published the graph in figure 20 showing massive congestion problems despite their abundant capacity. These congestion problems were largely caused by a very small percentage of Japanese broadband consumers who are using P2P applications. The study by the Japanese Ministry of Internal Affairs showed that P2P users who made up 10 percent of the total broadband customer base accounted for 65 percent to 90 percent of all traffic on the network.

Not only is the capacity on wireless networks more scarce; wireless networks are far more shared than wired networks, which presents unique engineering challenges not present on copper DSL or fiber networks.

The case of Japan illustrates that even the fastest broadband network in the world can be overrun by congestion. The reality is that there is no upper limit in the near to moderate for how much bandwidth consumers can demand because the quality and quantity of video can always go up. Standard high-definition TV (HDTV) resolution 1080p uncompressed video (1920 by 1080 pixel resolution at 60 frames a second) can occupy 3 gigabits per second of bandwidth. Quad-1080p uncompressed video (7680 x 4320 pixel resolution at 60 frames a second), which is already supported in some production HDTVs, can occupy 48 gigabits per second of bandwidth per video stream. No amount of additional bandwidth is ever likely to be enough because people will always want the ability to burst well beyond the guaranteed speed.

The second big problem with the argument that with sufficient bandwidth, the Internet could simply be a "dumb fat pipe" is that even when a network is built to provide more than sufficient bandwidth, it does not eliminate the jitter (the unwanted variations in packet delay), caused by packet bursting and queuing delay. Even networks operating at very low utilization can suffer jitter , and the only way to reduce jitter is by QoS technology.

Clearly, the fact that simply expanding network capacity is not an alternative to intelligent network management does not mean that more capacity isn't needed.

BOX 6: DEBUNKING THE MYTH THAT NETWORK CAPACITY IS A SUBSTITUTE FOR QUALITY OF SERVICE

In recent years, many net neutrality proponents have argued that Quality of Service (QoS) would not be necessary if Internet service providers would just expand network capacity by building "big pipes." A technical paper by Shalunov and Teitelbaum entitled "Why Premium IP Service Has Not Deployed (and Probably Never Will)" is a key source for this view. As explained below, however, the conclusions of this paper are not supported by the facts.

Shalunov and Teitelbaum make sweeping generalizations about QoS being unnecessary on the regular Internet in part by examining Internet 2—the very high-speed networks used predominantly by research universities: "Internet2 networks are generally well-provisioned and almost always lightly loaded. Packet loss and jitter experienced by best-effort traffic on Internet2 paths is almost always zero or is due to noncongestive causes." Yet the authors go on to admit: "This is especially true when there is no per-bit charge for Internet traffic, as is the case within Internet2. Without pricing disincentives, individual users can very significantly and very suddenly affect network utilization."

Shalunov and Teitelbaum clearly admit that the Internet2, which is a well-provisioned and metered network, is a somewhat special case with its lower probability for congestion and jitter than the lesser provisioned and unmetered regular Internet. But even though the Internet2 is relatively congestion- and jitter-free, Shalunov and Teitelbaum go on to explain that QoS may still be necessary under certain circumstances: "Premium service is about guaranteeing service quality. In essence, it is about removing a component of unreliability from the system—the probability that a network transaction fails because of network congestion. Although typical performance may be perfect, there would be considerable value in being able to assure that important sessions receive perfect network performance. Who wants the possibility that their important conference calls are disconnected or suddenly deteriorate in quality? Who wants a surgeon operating through robotic means on a different continent to experience IP packet loss artifacts?"

Shalunov and Teitelbaum go beyond the scope of the Internet2 and offer some unsupported conclusions: "In the U.S. today, the price of network capacity is low and falling (with the notable exception of residential and rural access) and the apparent one-time and recurring costs of Premium are high and rising (with interface speeds). In most bandwidth markets important to network-based research, it is cheaper to buy more capacity and to provide everybody with excellent service than it is to mess with QoS. In those few places where network upgrades are not practical, QoS deployment is usually even harder (due to the high cost of QoS-capable routers and clueful network engineers). "

Residential and rural access constitute quite a large exception to Shalunov and Teitelbaum's assertion that network capacity is cheap. In fact, talking residential and rural access is probably the entire broadband industry. But even Shalunov and Teitelbaum's claim that the price of network capacity is low elsewhere on the Internet is unsupported. The authors claim that QoS-capable routers and "clueful" network engineers are too expensive is also proven false by the existence of widely deployed fiber-to-the-node (FTTN) networks that run QoS to ensure that their Internet Protocol-based TV (IPTV) service has guaranteed bandwidth and low jitter.

In the end, Shalunov and Teitelbaum concede that small-scale ad hoc QoS solutions work well: "In the very few cases where the demand, the money, and the clue are present, but the bandwidth is lacking, ad hoc approaches that prioritize one or two important applications over a single congested link often work well. In this climate, the case for a global, interdomain premium service is dubious."

So a far more reasonable conclusion that should be drawn from the Shalunov and Teitelbaum paper is that large-scale QoS mechanisms across the entire Internet or Internet2 are probably infeasible, but small-scale point solutions that implement simpler DiffServ classification systems work well. The paper's assertion that network capacity is a better and cheaper alternative to QoS is unfounded. Capacity is not cheaper than QoS; nor is capacity an adequate substitute for QoS, particularly for ensuring performance of certain kinds of applications.

BenTeitelbaum and Stanislav Shalunov, "Why Premium IP Service Has Not Deployed (and Probably Never Will)," Internet2 QoS Working Group Informational Document, May 3, 2002 (HTMLized with updated references January 9, 2006) <qos.internet2.edu/wg/documents-informational/20020503-premium-problems-non-architectural.html> (accessed December 4, 2008).

Figure 20: Why capacity isn't a substitute for network management: the case of Japan

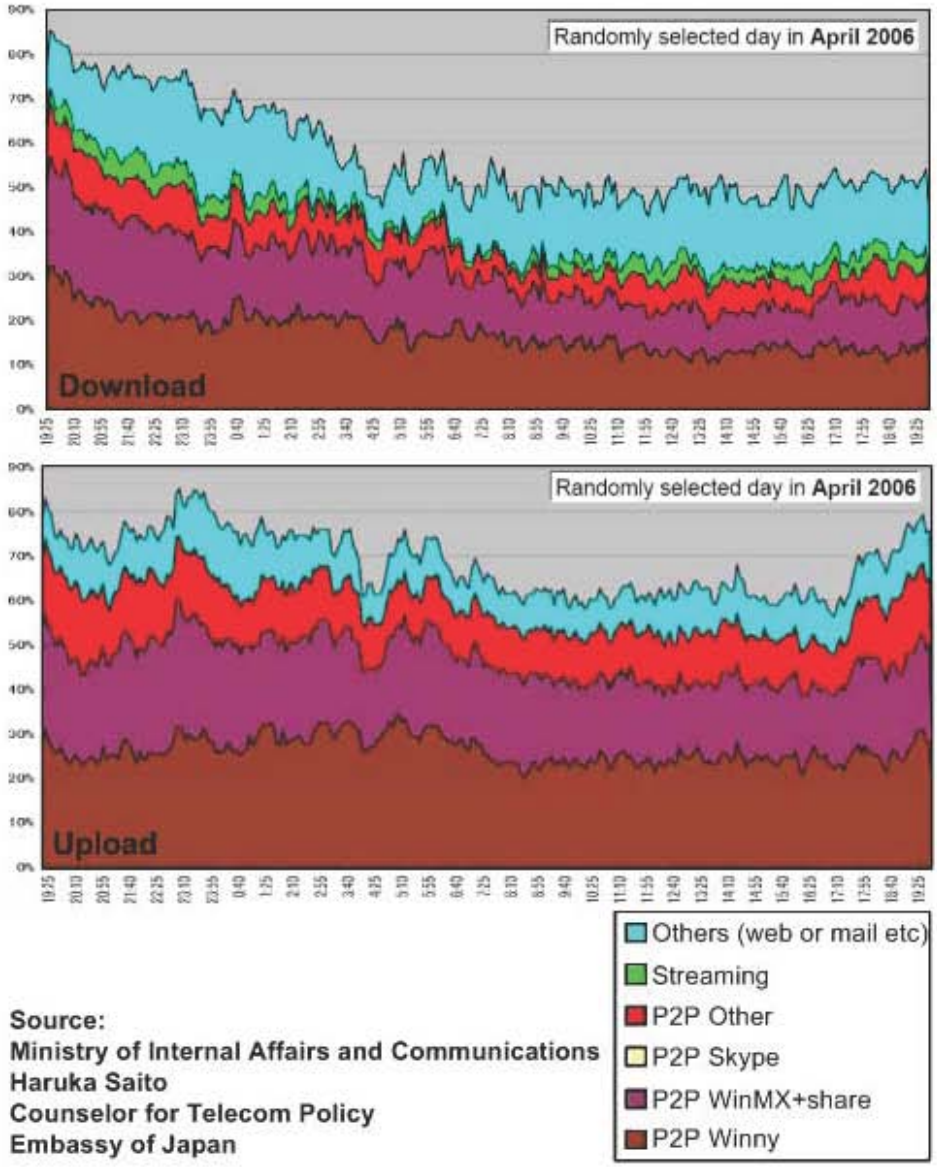

Source:
Ministry of Internal Affairs and Communications
Haruka Saito
Counselor for Telecom Policy
Embassy of Japan

No amount of QoS technology can substitute for capacity, and no amount of capacity can substitute for QoS. Recognizing this, all the large network operators in the United States are spending billions of dollars annually to constantly upgrade their networks.

B. Why Metered Pricing and Usage Caps Alone Will Not Solve the Problem

Some net neutrality proponents have argued that metered Internet or usage caps would be a better alternative to network management because the end-user is in control of what they access rather than letting the ISP control how and what consumers access. While metered Internet services can help, they alone cannot solve the problem anymore than can larger pipes. This is true for several reasons.

First, pricing is a blunt way of managing the network because it can only deal with average utilization on an hourly or daily basis, while congestion happens on a minute-by-minute or second-by-second basis. Pricing merely motivates people to cut back on average bandwidth consumption over the course of a few hours or over the month and there's simply no way for users to control their own consumption on a second-by-second basis.

The second problem with pricing schemes is that they cannot deal with the problem of jitter, which happens on a millisecond-by-millisecond basis. Jitter occurs on a millisecond level, and it can ruin VoIP, online gaming, video conferencing, and IPTV applications.

Instead, a combination of a protocol-agnostic bandwidth throttling system that ensures per-user fairness on a minute-by-minute basis and a protocol-specific QoS scheme that manages jitter with surgical precision at the millisecond level would allow networks to operate fairly and smoothly under existing flat-rate pricing schemes with generous usage caps.

Some people balked at the fact that Comcast was throttling applications for minutes at a time during congested periods. Yet congestion pricing mechanisms such as those in the United Kingdom and Australia would compel users to throttle themselves 12 hours a day during peak periods. Similarly, some usage caps would compel users to limit their use of volume-intensive applications like P2P or high-quality video streaming to a few days of the month.

This is not to say that it is wrong to have any kind of usage-based pricing models or multiple usage tiers. The heaviest users do drive up costs for the ISP to some extent. Heavy usage increases marginal costs but fixed infrastructure costs don't change based on usage. Most broadband providers have to pay usage charges for their backbone connection to the Internet and higher usage from their customers results in higher backbone usage costs. Tier 1 broadband providers like AT&T that don't pay usage fees to use the backbone because they own the backbone are still affected by higher broadband usage because they're forced to do more frequent upgrades to the backbone and other parts of their system.

C. Why Exclusive QoS on the Internet Is Better Than Exclusive QoS on Private Circuits

One of the key concerns of net neutrality advocates is that network operators will favor their own subscription-based video services exclusively and believe that network operators either should not get to prioritize any content or they have to give everyone that same priority. The proposed net neutrality legislation in the U.S. House and Senate follow this line of argument (see Box 1). If any of those bills or bills like them are approved in the United States in the future, network operators will no longer be permitted to pool their network resources and offer television services over the Internet portion of their network because they would no longer be able to guarantee the quality.

The reasoning behind these net neutrality proposals sounds compelling because it mandates equality. The problem is that these proposals are mandating a single service tier when even common carriers are permitted to have tiered services. From an engineering standpoint, this approach would cripple the future of the Internet. Even staunch supporters of net neutrality like Vint Cerf and Tim Berners-Lee acknowledge the need for QoS technology. Berners-Lee even insists that tiered QoS services at different prices are perfectly legitimate. But the biggest problem with enforcing a single QoS tier is that it accomplishes the exact opposite effect of its intended goal. Instead of increasing Internet capacity, it actually decreases Internet capacity.

If legislation bars or severely limits network management, U.S. consumers will either be left with a broken IPTV system or they'll be left with less Internet bandwidth.

When network operators are told that they can no longer favor their own video content on the Internet, they will simply move off the "Internet" on to a private network partition using circuit-switching networks on the same physical cabling. When that happens, they'll use fixed bandwidth allocation to the Internet service and the television service so even if the consumer isn't using the television service, the bandwidth cannot be dynamically shifted to the Internet service and the consumer gets less Internet bandwidth.

A real world example of this is IPTV service from AT&T U-verse in the United States and Deutsche Telekom T-Home in Germany. Both companies use fiber-to-the-node (FTTN) technology, which has a total capacity of 25 Mbps. Deutsche Telekom allows their Internet service to go up to the full 25 Mbps but whenever any high-definition (HD) or standard-definition (SD) channel is in use for the television service, as much as 8 Mbps and 4 Mbps of capacity respectively is set aside for IPTV service. So if two HD channels are being viewed at the same time, the Internet service could drop to somewhere between 9 Mbps and 15 Mbps depending on the amount of motion and complexity in the IPTV HD video streams. AT&T does

something similar in the United States but only offers up to 18 Mbps for the Internet service using their 25 Mbps FTTN broadband link. Like the Deutsche Telekom FTTN service, the service will slow to somewhere between 9 Mbps and 15 Mbps if two HD channels are being viewed on the IPTV service at the same time. In both cases, consumers are told up front about the limitations of the Internet service when IPTV is in use. In both cases when no TV is being watched, the broadband service goes up to the full 18 Mbps or 25 Mbps for AT&T and Deutsche Telekom.

The reaction from many net neutrality proponents is that this is somehow blatant favoritism for the broadband provider's own video service at the expense of all other Internet services. But this fear of exclusive prioritization of video services on Internet-based technology is unfounded in light of the fact that guaranteed video delivery on separate physical cabling or a separate circuit within the same cabling has always been acceptable. Cable companies, for example, dedicate more than 90 percent of their cable infrastructure to their television services while their broadband services are fixed at less than 10 percent. Broadband companies adopting FTTN and IPTV technology are dynamically allocating 0 percent to 65 percent of their cable infrastructure to television while broadband technology gets the remaining 35 percent to 100 percent.

Why should anyone object to the more efficient solution and more generous allocation of resources to broadband and object to a new entrant in the traditional television market? It would certainly not be better to go back to the old fixed allocation model and have one less competitor in the television space. Exclusive QoS prioritization for IPTV service is precisely the behavior that consumers want because this allows IPTV to work properly while other applications operate freely. Without IPTV prioritization, the television service will degrade in quality whenever other jitter-inducing applications are in use because IPTV is a real-time application. Even when, for example, a P2P application is manually restricted to a small percentage of the broadband connection, it still causes occasional "hiccups" in the IPTV service because jitter has little to do with bandwidth consumption and a lot to do with the tendency of certain applications to burst out multiple packets.

With the QoS technology in place, peer-to-peer applications will get to run much faster because the user will be able to lift any speed restrictions on the application without fear of causing problems for the IPTV service. Without QoS technology, a household would have to choose between watching TV or using P2P because non-prioritized IPTV is "allergic" to P2P traffic. This is precisely the reason that P2P companies and P2P Internet standards bodies are voluntarily labeling their own packets as lower priority because they know that a friendly P2P application is less likely to be restricted.

The federal government has a key role to ensure openness and fair play on the Internet. However, it should do this with sensible rules. Policies should strive to prevent any potential abuse without eliminating the ability of ISPs to manage their networks in ways that produce the best possible user experience for the largest number of users, and without eliminating incentives to build the next generation broadband network.

If legislation bars or severely limits network management, U.S. consumers will either be left with a broken IPTV system or they'll be left with less Internet bandwidth. If broadband providers are prohibited from exclusively prioritizing IPTV over broadband, they'll simply only offer no more than 9 Mbps for the Internet service and permanently partition off 16 Mbps of fixed bandwidth on a separate dedicated circuit for their IPTV service. So instead of getting 9 Mbps to 25 Mbps of Internet service, depending on how many HD channels are being watched, consumers will permanently get 9 Mbps of Internet service even when the IPTV service isn't being used at all.

Proponents of net neutrality often ask why telecommunications companies can't simply offer just 25 Mbps of Internet services and charge enough for the service to be profitable and forget about television service. They can't do that for simple reason that, at least currently, not enough consumers will buy the service. Even community or municipally operated broadband services rely on television services to offset the massive investment costs, so why would private companies who can't rely on taxpayer support be any different?

CONCLUSION

The Internet in all its glory was never perfect in architecture. The inability of the ISPs to fairly allocate bandwidth between customers and seamlessly support multiple applications on network connections has long caused conflict between users and applications. More capacity is always welcome and public policy should support it, but for the foreseeable future, it will get used soon after it gets built. Moreover, it will not solve these two fundamental architectural problems of the packet-switching Internet. It will always be necessary to ensure equitable bandwidth distribution between broadband users with protocol-agnostic network management technologies. And, it will always be necessary to manage jitter on shared Internet links or shared broadband connection using QoS technologies to ensure that different types of applications all work effectively.

The federal government has a key role to ensure openness and fair play on the Internet. However, it should do this with sensible rules. Policies should strive to prevent any potential abuse without eliminating the ability of ISPs to manage their networks in ways that produce the best possible user experience for the largest number of users, and without eliminating incentives to build the next generation broadband network. Toward that end the FCC should oversee broadband providers and ensure that ISP network management practices are open, transparent and not anti-competitive. And the ISP industry should continue its efforts to develop and abide by industry codes of good conduct regarding network management that include, but are not limited to, fuller and more transparent disclosure to consumers of network management practices.

ENDNOTES

[1.] Michael K. Powell, "Preserving Internet Freedom: Guiding Principles for the Industry," remarks at the Silicon Flatirons Symposium, "The Digital Broadband Migration: Toward a Regulatory Regime for the Internet Age," University of Colorado School of Law, Boulder, Colo., February 8, 2004 (as prepared for delivery) <hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-243556A1.pdf> (accessed December 4, 2008).

[2.] Robert D. Atkinson and Philip J. Weiser, "A 'Third Way' on Net Neutrality," Information Technology and Innovation Foundation, Washington, D.C., May 30, 2006 <www.itif.org/files/netneutrality.pdf> accessed December 4, 2008).

[3.] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture," informational memo, The Internet Engineering Task Force, June 1994 <tools.ietf.org/html/rfc1633> (accessed December 4, 2008); and P. Almquist, "Type of Service in the Internet Protocol Suite," proposed standard memo, The Internet Engineering Task Force, July 1992 <tools.ietf.org/html/rfc1349> (accessed December 4, 2008).

[4.] A TCP flow is the stream of data flowing between two computers connected via the Internet (or any TCP/IP network). Computers typically use a single flow for a single task but they can open up multiple TCP flows. We can think of a flow as a conversation between two computers but computers can have multiple conversations.

[5.] Todd Enderwood, "Will Work for Bandwidth", CircleID, Nov 20, 2008 <www.circleid.com/posts/20081119_will_work_for_bandwidth/> accessed December 5, 2008.

[6.] Brett Glass, LARIAT.NET, local Internet service provider in Laramie, Wyo., comment from private e-mail exchange dated July 17, 2008 published with permission.

[7.] Copyrighted TV shows, movies, and games dominates the Pirate Bay's top 100 traded P2P files. The Pirate Bay Website <thepiratebay.org/top/all> (accessed December 4, 2008).

[8.] Note that this works for Multi Packet Label Switching (MPLS) and Internet-based Wide Area Network (WAN) architecture but not spoke-hub dedicated circuit architecture.

[9.] Pando Networks, "The P4P Working Group," n.d. <www.pandonetworks.com/p4p> (accessed December 4, 2008).

[10.] Switch and Data, Press release touting the fastest growth in a private peering point <www.switchanddata.com/press.asp?rls_id=143>.

[11.] Ibid.

[12.] Bob Briscoe, T. Moncaster, and L. Burness, "Problem Statement: We Don't Have To Do Fairness Ourselves draft-briscoe-tsvwg-relax-fairness-00," Internet Engineering Task Force working draft, November 12, 2007 <www.cs.ucl.ac.uk/staff/bbriscoe/projects/2020comms/accountability/draft-briscoe-tsvwg-relax-fairness-00.html> (accessed December 4, 2008).

[13.] Explicit Congestion Notification (ECN) is an "explicit" congestion notification system whereas Jacobson's algorithm is an implicit notification system. Pre-ECN routers notify pre-ECN computers to slow down implicitly by dropping packets when the network is close to overflow and the computers respond to dropped packets by slowing down because they assume the network is congested. The problem with Jacobson's solution is that packets are sometimes dropped when there isn't congestion (especially wireless networks experiencing interference) and the clients wrongly assume they need to slow down when they don't need to. The other problem with Jacobson's solution is that the zigzag nature of the throughput and the dropping of packets as an implicit signaling mechanism results in higher overhead which reduces overall network performance.

ECN-capable routers handle congestion more intelligently by explicitly notifying ECN-capable computers how fast they can go rather than relying on dropped packets or zigzag traffic patterns. ECN marks receive acknowledgement packets coming from the receiving computer with information telling the transmitting computer what the network can support in terms of throughput. The ECN client can intelligently speed up or slow down without a drastic halving of throughput and a slow increase in throughput. ECN computers running on ECN networks can perform much more efficiently and faster than a network not running ECN. ECN by itself does not deal with the question of fairness, but it could be part of the solution when it eventually becomes widely adopted.

[14.] Laurence Brett ("Brett") Glass, LARIAT.NET, Laramie, Wyo., Letter to Marlene H. Dortch, Secretary of the Federal Communications Commission, July 29, 2008 <bennett.com/blog/pitchers/20questions.pdf > (accessed December 4, 2008).

[15.] Yasu Taniwaki, Director Telecommunications Policy Division Telecommunications Bureau, Japan's Ministry of Internal Affairs and Communications, "Network Neutrality and Competition Policy in Japan," PowerPoint presentation (charts to the right on slides 15 and 16), December 2007 <www.soumu.go.jp/joho_tsusin/eng/ presentation/pdf/071204_1.pdf> (accessed December 4, 2008).

[16.] Richard Bennett, "FCC's Comcast Ruling Inconsistent and Incoherent," CircleID blog, posted August 1, 2008 <www.circleid.com/posts/88103_fcc_comcast_ruling_inconsistent_incoherent/> (accessed December 4, 2008).

[17.] Iljitsch van Beijnum, "IETF: Find more peer-to-peer bandwidth but use it sparingly," ArsTechnica Website (news), August 3, 2008 <arstechnica.com/news.ars/post/20080803-ietf-find-more-peer-to-peer-bandwidth-but-use-it-sparingly.html> (accessed December 4, 2008).

[18.] Bennett, op. cit.

[19.] Robb Topolski, "George Ou: Protocol Agnostic Doesn't Mean Protocol Agnostic," Public Knowledge Website, July 17, 2008 <www.publicknowledge.org/node/1661> (accessed December 4, 2008).

[20.] CostQuest Associates, "U.S. Ubiquitous Mobility Study", April 17, 2008 <fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6519893737> (Accessed December 6, 2008).

[21.] Robert D. Atkinson, Daniel K. Correa, Julie A. Hedlund, "Explaining International Broadband Leadership," Information Technology and Innovation Foundation, Washington, D.C., May 1, 2008 <www.itif.org/files/ ExplainingBBLeadership.pdf> (accessed December 4, 2008).

[22.] Taniwaki, op. cit.

[23.] Tim Berners-Lee, "Net Neutrality: This is Serious," DIG blog, posted June 21, 2006 <dig.csail.mit.edu/breadcrumbs/node/144> (accessed December 4, 2008).

[24.] Trevor R. Roycroft, Economic Analysis and Network Neutrality: Separating Empirical Facts from Theoretical Fiction," issue brief prepared for the Consumer Federation of America, Consumers Union, and Free Press, June 2006 <www.freepress.net/files/roycroft_study.pdf > (accessed December 4, 2008).

Appendix A: Network Glossary

Bandwidth. The rate at which files are transmitted through a computer network, commonly referred to as "speed" but more accurately described as throughput. This kind of bandwidth is usually expressed in bits (of data) per second (bps). Occasionally, it's expressed as bytes per second (B/s) which is 1/8th the value of bits per second because each byte is 8 bits. Unfortunately, the distinction between Bps and B/s is often mixed up because they sound alike. Bandwidth is also formally expressed in multiples of 1,024 so a Kilobit (Kbps) per second is 1,024 bps, Megabit (Mbps) per second is 1,048,576 bps, and Gigabit (Gbps) is 1,073,741,824. But this can also be a little confusing because people casually use simple multiples of 1,000 to define Kilo, Mega, and Giga.

Bit (b). A binary digit, represented by either a 0 or a 1, which is the smallest basic unit of information storage and communication in the computer world. (A lower case "b" denotes bit, whereas an upper case "B" denotes Byte.)

BitTorrent protocol. A content distribution protocol that enables efficient software distribution and peer-to-peer (P2P) sharing of very large files, such as entire movies and TV shows, by enabling users to serve as network redistribution points.

Broadband. In general, broadband refers to telecommunication in which a wide band of frequencies is available to transmit information. Because a wide band of frequencies is available, information can be multiplexed and sent on many different frequencies or channels within the band concurrently, allowing more information to be transmitted in a given amount of time (much as more lanes on a highway allow more cars to travel on it at the same time).

Broadband technologies. The standard broadband technologies in most areas are digital subscriber line (DSL) and xDSL, which operate over copper telephone lines; cable modems; and optical fiber that can handle TV, voice calls, and Internet access. Newer wireless broadband technologies include Wi-Fi and WiMAX and Long-Term Evolution (LTE).

Buffer. In computing, a buffer is a region of memory used to temporarily hold data while it is being moved from one place to another.

Burst. A short but intense increase in bandwidth.

Byte (B). A unit of measurement of information storage and communication in the computer world that consists of 8 bits. (An upper case "B" denotes Byte, whereas a lower case "b" denotes bit.)

Circuit-switching network. A network built on a networking technology called circuit switching that allocates fixed resources when a connection is initiated and allows additional users to use the system only if the circuits are open. Only a small portion of the total user population can use a circuit switching e network at any given time, and communications on circuit-switching networks are generally limited to two devices unless the circuit is terminated and a new circuit is created. Thus, a circuit-switching network is too limiting and inflexible to be used on the Internet.

Client. The requesting program or user in a client/server relationship. The user of a Web browser, for example, is making client requests for pages from servers all over the Web. The browser itself is a client in its relationship with the computer that is getting and returning the requested HTML file. The computer handling the request and sending back the HTML file is a server.

Client-server file distribution architecture. An architecture for distributing files on the Internet that involves having a single server transmit (upload) files and clients receive (download) files using traditional protocols such as File Transfer Protocol (FTP) or Hypertext Transfer Protocol (HTTP).

Content delivery network (CDN) file distribution architecture. An architecture for distributing files on the Internet that involves globally distributed high-speed cache servers that store copies of files from a central server and redistribute them to nearby clients. The CDN model puts half the load on broadband networks that the peer-to-peer file distribution model does because it does not require clients to upload and download files.

Content provider. An organization or individual that creates information, educational, or entertainment content for the Internet, CD-ROMs, or other software-based products.

Data center. A building that hosts lots of servers, server farms, networking equipment, power, and cooling infrastructure.

Delay. The time it takes an individual packet to travel through a network. Packet delay should not be confused with file transfer time, which correlates to bandwidth. There are two forms of delay that affect packet transfer times: latency (the time it takes a bit or packet to traverse a non-congested network before arriving at its intended destination) and jitter (a measure of the variation in packet delay).

Download. To receive data. In the context of broadband, the term is also a noun that means data flowing towards the home user.

Downstream. The portion of the broadband connection where data travels towards a computer. In the context of broadband, it's the data link traveling towards the home.

Endpoint device/endpoints. In the context of the Internet, devices that attach to the Internet such as computers, servers, digital still or video cameras, music players, printers, or any other devices that uses the Transmission Control Protocol/Internet Protocol (TCP/IP).

File. In the context of computers and computer networks, an electronic document, digital photo, digital music, digital video, a database containing information, or a computer application.

File Transfer Protocol (FTP). One of the earliest file distribution protocols on the Internet. FTP can experience a slowdown in file transfer speeds if it experiences very high latency but it is relatively immune to jitter. FTP is insensitive to packet loss in terms of losing data but it can experience longer and slower download times if there are excessive and continuous packet drops.

Flow. One communication session between two Transmission Control Protocol (TCP) endpoints. Applications traditionally use a single TCP flow, but peer-to-peer (P2P) applications use multiple flows to gain immunity against Jacobson's algorithm.

Hypertext Transfer Protocol (HTTP). A communications protocol used for retrieving interlinked text documents (hypertext) that led to the establishment of the World Wide Web.

Internet Protocol (IP). Part of the set of core communications protocols for the Internet and other similar networks that was originally developed in the late 1960s and early 1970s to meet the data needs of the U.S. Department of Defense. The IP specifies algorithms for interconnecting networks and routing traffic on the Internet. IP addresses are the "phone numbers" of the endpoints connected to the Internet and are used to route data.

Internet. A global network of interconnected computer networks that consists of millions of private and public, academic, business, and government networks that exchange data by packet switching using the standardized Transmission Control Protocol/Internet Protocol (TCP/IP).

Internet service provider (ISP). A company that offers its customers access to the Internet.

IPTV (Internet Protocol Television). Digital television service delivered over Internet Protocol on a network that necessitates a broadband connection.

Jacobson's algorithm. A network congestion control algorithm for Transmission Control Protocol/Internet Protocol (TCP) developed by Van Jacobson. When File Transfer Protocol (FTP) became popular on in the mid 1980s, the Internet became a traffic jam at certain times of the day where nothing could get through. This was known as the first meltdown of the Internet and it required a fundamental change in TCP to fix the Internet. Van Jacobson's elegant revision to TCP was so effective that it was immediately adopted in 1987 on every computer of the Internet.

Jitter. The measure of the variation in packet delay. High-jitter conditions are essentially micro-congestion storms that last tens or hundreds of milliseconds. High jitter occurs whenever a large number of packets come from a faster network link to a slower network link or where several networks links merge to a single link. When this happens, network devices such as routers and switches get backlogged, and they force packets to wait inside their memory buffers, increasing the time it takes packets to traverse a network.

Kilobits per second (Kbps). Thousands of bits per second, a measure of bandwidth (the amount of data that can flow in a given time) on a data transmission medium.

Latency. A simple measure of delay, which is largely dictated by the speed of light in glass and the physical distance between two devices on the Internet.

Long Term Evolution (LTE). A fourth-generation (4G) wireless broadband technology developed by the Third Generation Partnership Project.

Megabits per second (Mbps). A common unit of network data transmission throughput used in the networking industry. One megabit is commonly assumed to be one million bits, but the technically correct definition is 1,048,576 bits or 1024 to the 2nd power. Unfortunately, people use the 1000 base and the 1024 base just as frequently, and sometimes it's hard to tell which one is being used.

Megabytes per second (MB/sec). A common unit reported by computer applications and equal to 8 megabits per second. Unfortunately, people often confuse MB/sec for Mbps when they're different by a factor of 8.

Millisecond (ms or msec). One-thousandth of a second, a measure commonly used in measuring packet travel time on the Internet.

Multiflow. A file transfer technique where multiple TCP flows are simultaneously used to transmit a single file. Multi-flow is commonly used by peer-to-peer (P2P) applications to gain immunity to the bandwidth reallocation mechanism in Jacobson's algorithm for Transmission Control Protocol (TCP).

Network. In the context of computers, a network is a series of points or nodes interconnected by communication paths. Networks can interconnect with other networks and contain subnetworks. Networks come in all sizes and shapes and may operate over copper wires, fiber optics, wireless, or even power lines. Large telephone networks and networks using their infrastructure (such as the Internet) have sharing and exchange arrangements with other companies so that larger networks are created. Networks are commonly referred to as "pipes" (although networks like the Internet are physically made up of loose tube optical fiber).

Network capacity. The number of users that can simultaneously use a network with a certain level of performance.

Online gaming. In the context of this paper, online gaming refers to games where players use the Internet to connect to each other in a virtual world for fun. In order for these games to feel responsive, the delay from either latency or jitter must be low as possible and even lower than Voice over Internet Protocol (VoIP) requirements.

Packet. The unit of data that is routed between an origin and a destination on the Internet or any other packet-switched network. Packets are actually made of electromagnetic signals flowing through wiring or fiber optic glass.

Packet delay. The time it takes individual packets to traverse a network. Packet delay comes from latency and jitter. Depending on the type and amount of delay, it may not be a problem for data applications but it's brutal to real-time applications like Voice over Internet Protocol (VoIP) or online gaming.

Packet-switching network. A network built on a technology called packet switching that splits data traffic (digital representations of text, sound, or video data) into small chunks, called packets, that are then routed over a shared network on the basis of the destination address contained within each packet. Breaking communication down into packets allows the same data path to be shared among many users in the network. A packet-switching network dynamically divides up the resources among the active users on the network. If few users are on the network, then those few users get a lot of resources allocated to them. If many users are on the network, each user gets fewer resources but is not locked out of the system. The Internet's predecessor ARPANET and the Internet were the first packet-switching networks in the world.

Peer. A computer on the Internet that can upload and download small metadata files known as torrents in a peer-to-peer (P2P) Internet file distribution system. Generally, a peer does not have the complete file; if it did, it would be called a "seed."

Peering arrangement. A contractual agreement between two Internet service providers (ISPs) to interconnect and exchange traffic. This peering arrangement may involve the larger ISP charging the smaller ISP because the smaller ISP will benefit more from the infrastructure built and paid for by the larger ISP, or the arrangement may involve no money changing hands if each ISP has something of equal value.

Peer-to-peer (P2P) file distribution architecture. An architecture for distributing files on the Internet invented in 1999 that involves multiple computers called peers that receive and transmit pieces of files to other peers. Once a peer has the complete file, it is called a seed. The P2P method of distribution is free to the content provider but effectively doubles the traffic load for Internet service providers (ISPs), because it involves both downstream traffic and upstream traffic. What makes P2P unique and revolutionary is that it endpoints downloading files are simultaneously uploading to other peers, which in turn upload to other peers. This chain of peers allows P2P file distribution to scale indefinitely at the expense of a lot more upload traffic for the P2P user and his/her ISP. P2P is for the most part immune to latency and jitter because of its multiflow properties. P2P is insensitive to packet loss in terms of losing data but it can experience longer and slower download times if there are excessive and continuous packet drops.

Performance. In the context of networks, performance can have many meanings. It typically refers to throughput performance but it could refer to latency and jitter characteristics or a combination of all three.

Protocol (or communication protocol). In the context of computing, a protocol is the convention or standard that governs the syntax, semantics, and synchronization of communication between computers or devices.

Protocol-agnostic network management system. A network management system that does not look at the protocol to function. Protocol-agnostic systems for network management are ideal for distributing bandwidth equitably among broadband customers but they can never ensure the equal performance of different applications sharing the same broadband connection—and they can never eliminate jitter on bottleneck segments of the Internet. Dealing with these problems requires protocol-specific network management systems.

Protocol-specific network management system. A network management system that takes protocol headers and other factors into account to determine packet priority. Protocol-specific systems for network management systems such as Quality of Service (QoS are needed to ensure the equal performance of different applications sharing the same broadband connection—that is, to make networks (especially broadband) more conducive to simultaneous application usage and to address the problem of jitter on bottleneck segments of the Internet.

Quality of Service (QoS). In the context of packet-switching networks, a form of protocol and application specific traffic engineering, its key benefit being its ability to mitigate jitter and create harmony between applications sharing the same broadband connection. QoS is by definition is a protocol-specific technology. This technology is ideal for creating harmony between applications sharing a single broadband connection. It is controversial if it is used to allocate bandwidth between broadband customers. There are instances where protocol-specific jitter management is necessary for shared network links between multiple broadband customers. This does not conflict with protocol-agnostic systems so long as equitable sharing of bandwidth between users is maintained. There are many common alternative words used to describe QoS mechanisms (e.g., "enhanced Quality of Service," "network intelligence," "prioritization," and "premium service"). QoS is not just one technology; it is a complex field of study with dozens of Internet Engineering Task Force (IETF) standards that compete with or complement one another.

Random peering behavior. The behavior of traditional peer-to-peer (P2P) applications, which connect with any other peer on the Internet regardless of network topology or geographic distance. This results in very inefficient usage of bandwidth at the core of the Internet. Newer P2P applications are beginning to adopt more intelligent ways to distribute content by selecting the nearest peer. This not only accelerates P2P performance but it also alleviates congestion at the core of the network.

Seed. A computer on the Internet that can upload and download small metadata files known as torrents and has the complete file. If it does not have the complete file, it is called a "peer."

Server. In the client/server programming model, a server is a computer or program that awaits and fulfills requests from client programs in the same or other computers.

Server farm. A group of computers acting as servers and housed together in a single location. Server farms are typically owned by a single group or company.

Statistical multiplexing. A system for overbooking bandwidth that is commonly used in all broadband deployments throughout the world. Because most users on a network are idle most of the time and because network bandwidth is divided among users, each user gets several times more bandwidth than the minimum guaranteed bandwidth. Networks that employ statistical multiplexing are several times less expensive per unit bandwidth than dedicated bandwidth networks with guaranteed speeds.

Transmission Control Protocol/Internet Protocol (TCP/IP). The set of core communications protocols for the Internet and other similar networks that was originally developed in the late 1960s and early 1970s to meet the data needs of the U.S. Department of Defense. TCP specifies algorithms for how computers and other endpoint devices communicate. It determines behavior such as the transmission speed between endpoints by responding to network conditions and it provides reliable data transport by handling error correction. IP specifies algorithms for interconnecting networks and routing traffic on the Internet. IP addresses are the "phone numbers" of the endpoints connected to the Internet and are used to route data.

Throttle. Slow down.

Throughput. In networking, a term that commonly refers to bandwidth.

Upload. In the context of computers, the term is a verb that means to transmit data. In the context of broadband, the term is also a noun that means data flowing away from the home user.

Upstream. The portion of the broadband connection where data travels away from a computer. In the context of broadband, it's the data link traveling away from the home.

Voice over Internet Protocol (VOIP). A real-time isochronous technology that implements telephony on packet-switching networks like the Internet or private networks running Internet Protocol (IP). In order for the telephone call to feel responsive, the delay from either latency or jitter must be low and preferably well below 100 milliseconds. VoIP applications are fairly intolerant to packet drops.

Wireless broadband technology. A fairly new technology that provides high-speed wireless Internet and data network access over a wide area. Wireless broadband technologies include Wi-Fi and WiMAX and Long-Term Evolution (LTE). It is easy to envision a day when wireless broadband access will surpass wired broadband services because the total market for residential broadband service is limited to the number of households, whereas the total wireless broadband market is limited by the number of future Internet-enabled mobile phones.

ABOUT THE AUTHOR

George Ou is a Senior Analyst at the Information Technology and Innovation Foundation and works out of Silicon Valley. Outside of ITIF, Mr. Ou is an Information Technology and CISSP Security Consultant who was the founder of ForMortals.com. He recently served two years as Technical Director and Editor at Large for TechRepublic and ZDNet (both property of CNET Networks) doing in-depth coverage on IT and technology topics. Before journalism, he worked as an IT professional who designed and built wired network, wireless network, Internet, storage, security, and server infrastructure for various Fortune 100 companies.

ABOUT THE INFORMATION TECHNOLOGY AND INNOVATION FOUNDATION

The Information Technology and Innovation Foundation (ITIF) is a nonprofit, non-partisan public policy think tank committed to articulating and advancing a pro-productivity, pro-innovation and pro-technology public policy agenda internationally, in Washington and in the states. Through its research, policy proposals, and commentary, ITIF is working to advance and support public policies that boost innovation, e-transformation and productivity.

ACKNOWLEDGMENTS

The author would like to thank Bob Briscoe, Brett Glass, and Richard Bennett for their technical insights. In addition, the author would like to thank Daniel Castro, Kerry Kemp, Priscilla Jang, and Rob Atkinson for editorial efforts.