2007 STEP White Paper Competition, UC Berkeley

BRINGING NEUTRALITY TO NETWORK NEUTRALITY

Kai Zhu

EXECUTIVE SUMMARY

Network neutrality (NN) is becoming a fierily debated issue within Internet-related industries, academia, the public, and Congress. Internet service providers such as AT&T and Verizon want to prioritize Internet traffic and profit on the prioritization. Internet content providers such as Google, Yahoo, and Microsoft want to ban such prioritization via NN legislation.[1] Five bills for or against NN were introduced in Congress in 2006.[2] The issue has sharply divided scholars. Both proponents and opponents of NN claim that Internet innovation is at risk.

The debate is complicated, and it immediately puts the interests of many industry giants at stake. Any solution for the debate will likely have profound and enduring social impacts. In searching for a solution, it is crucial to distinguish between the real problems that NN seeks to solve and the interests of the debating parties. The many dimensions of the debate render this exercise challenging because it is hard for anyone to grasp all the technical, economic, legal, and social complexities as well as subtleties of something as big as the Internet. The fundamental policy challenge in the NN debate is to strike a balance between incentives that are essential for innovation, and monopolies that are necessary for incentives but harmful for innovation.

This paper argues that (1) the Internet has never been neutral and has never been designed to be neutral; (2) Internet traffic prioritization can both coexist with and encourage Internet innovation; and (3) some minimal regulation is needed to prevent market power abuse and usage discrimination in the Internet service market. The paper proposes a technically feasible middle ground solution to the debate.

I. THE REALITIES OF THE INTERNET

For a serious network neutrality (NN) debate, it is critical to understand, from a technical standpoint, that the modern Internet is very different from what it was thirty years ago. The Internet started its life in 1969 as a research network funded by the U.S. government, but was commercialized in the early 1990s.[3] Since then it has grown rapidly and steadily—in its size, its territorial scope, the number of its users, the number and types of applications running over it, and, most importantly, the sophistication of the many technologies underlying it. The architects of the original Internet did not and could not envision the many new technologies and applications that are now common for the Internet. The evolution of the Internet shows that the original Internet architecture cannot serve well current or future Internet applications.

A. A Quick Technical Overview

For an Internet communication, the source computer splits digitized data into small pieces called packets and submits those packets into the network; the network delivers the packets to the destination computer.[3] Multiple intermediate hops, called routers, exist between the source and the destination.[3] Along this path, each router receives a packet from an upstream router and then forwards it to a downstream router.[3] Thus the packet is "routed" hop-by-hop to its destination. Each packet contains some basic information such as its source and destination Internet Protocol (IP) addresses.

Routers run distributed routing protocols to learn the topology of the Internet.[3] A router knows how to forward a packet by looking at both its routing table(s) and the destination IP address of the packet.[3] When packets arrive, a router may need to queue them first and then forward them.[3] This is because those packets may arrive from different upstream routers around the same time but need to go to the same downstream router, while the instant router has fixed bandwidth toward that downstream router. Thus a contest for limited resources may exist in a router, and a packet may experience unpredictable queueing delay at each router. The technical essence of the NN debate is whether routers can reduce the queueing delays of some packets by increasing the delays of other packets.

B. Real-time Applications and Quality of Service

The arrival of real-time applications distinguishes the modern Internet from the original one.[4] A real-time application such as streaming video or online gaming is time-sensitive; its data can only tolerate a limited end-to-end delay.[3] A real-time packet not meeting its delay limit is useless; a small fraction of such packets may lower the quality of the communication to the point of rendering the whole application useless.[3] Thus, a real-time application cannot work well without reasonable quality of service (QoS), for which an end-to-end delay bound is the major metric. Interactive real-time applications such as online gaming or tele-surgery—real-time in both directions—have demanding QoS requirements.[3] QoS is at the center along the technical dimension of the NN debate. NN proponents tend to downplay or misinterpret its importance because they are not familiar with the technical challenges faced by QoS provision.[18]

A router controls queueing delays primarily via link schedulers, which are functional modules within the router that control the sending orders of queued packets at outgoing links.[3] Because a link has fixed bandwidth, a scheduler cannot limit the delay of every waiting packet if the queue is long; it can only limit the delays of some packets by giving them higher priorities for transmission. The simplest link scheduler is First-In-First-Out (FIFO), which sends packets out in their arrival order.[3] This order-preservation nature of FIFO is tightly coupled with the "neutrality" concept in the NN debate. Although FIFO is trivial and does not control queueing delays, it can be easily implemented; thus, FIFO became ubiquitous in older routers and is still dominant in modern routers. A link scheduler, however, can implement very complex algorithms by taking as its inputs many parameters, including the QoS requirements and traffic characteristics of each application passing the queue.[3]

Some have argued that over-provision, which means building a network with significantly more bandwidth than what the normal level of network traffic load demands, will solve the QoS problem.[3,18] However, the idea has not become mainstream.[3] Two related problems challenge this idea. First, the "normal level" of network traffic load is a moving target because whenever the network has "extra resources" due to over-provision, those "extra resources" will induce newer applications with heavier traffic to appear; such heavier traffic tends to exhaust the "extra resources."[3] Second, this idea can at best "almost" solve the QoS problem; it cannot guarantee QoS for mission-critical applications such as tele-surgery.[3]

C. Internet Architecture Is Alive and Growing

The Internet Engineering Task Force (IETF), an international organization in charge of the technical development of the Internet, sets up de facto technical standards for the Internet by publishing a series of documents called "Request for Comments" ("RFCs").[5] Currently, the actual technical work of IETF is split among more than 100 Working Groups (WGs) for specific technical areas of the Internet.[6] Those WGs keep producing new RFCs and shaping the overall Internet architecture at an ever-increasing speed. Beyond the sheer number of new RFCs produced each year, the Internet continues to evolve even in its most fundamental areas. For its two key protocols, Transport Control Protocol (TCP) and IP, many changes and updates were made throughout the last twenty-five years. Routing is an extremely critical function for any data network, and Internet routing keeps evolving.[7] Many of the IETF WGs are completely new to the original Internet. The first RFC dedicated to real-time applications appeared in as late as 1990, while today 14 WGs work in the Real-time Applications and Infrastructure Area.[6] No major RFCs on Internet security appeared before 1990, but today 17 WGs work in the Security Area alone.[6] All the examples and data above show that the Internet architecture is alive and growing.

II. THE GREAT NETWORK NEUTRALITY DEBATE

At its core, the NN debate focuses on a technical question that has great economic significance, although the exact meaning of the term NN has received different and confusing interpretations. Interestingly, the legal community originated and popularized the debate, which has since fallen victim to political and ideological polarities.

A. Controversies about Internet Usages and Services

In 2004, Madison River Communications, a small Internet service provider (ISP) in North Carolina, blocked their customers from using the market-leading Vonage Voice over IP (VoIP) service. After the Federal Communications Commission (FCC) intervened, Madison River restored the service, entered into a consent decree, and paid a $15,000 fine.[8] In 2005, Canadian telephone giant Telus blocked the access to a website supporting the company's labor union during a labor dispute; the blocking lasted for about sixteen hours.[9] In April 2006, Time Warner's AOL blocked all emails mentioning an advocacy campaign opposing AOL's pay-to-send email scheme, but the company said the incident was a "software glitch."[10]

ISPs now complain that major Internet content providers (ICPs) have generated too much traffic that burdens the network and worsens the experiences of general Internet users. ISPs have proposed ways to charge ICPs higher fees for better QoS, a scheme termed "a two-tiered Internet."[11] Debates over such proposals turned bitter. Some ISP executives have used hyperbolic language to threaten ICPs. Those threats resulted in equally hyperbolic responses from ICPs such as Google, counter-threatening to pursue NN legislation and antitrust lawsuits.[12] Not surprisingly, telecommunications equipment vendors such as Cisco and Motorola joined their customers to oppose NN because they expected to sell ISPs equipment that enable QoS.[13]

Numerous advocacy groups are taking strong positions on the NN debate. Among many others, SavetheInternet.com is an NN-supporting coalition, which includes the largest consumer advocacy groups in the nation;[14] Hands Off The Internet is an anti-regulation coalition funded by major telecommunications companies.[15] SavetheInternet.com was created in April 2006, and it collected more than one million signatures in just two months in 2006 to support NN; it has set up an ambitious agenda to push Congress to pass NN laws in 2007.[14] Collectively, such advocacy groups have generated much publicity for the NN debate.

B. Views of Academics on Network Neutrality

Some legal scholars hatched and advocated the NN concept, but they tend to disagree on its scope. Other legal scholars and economists do not believe in NN as a solid public policy. In 2003, law professors Tim Wu and Lawrence Lessig sent the FCC an *ex parte* letter proposing a set of NN rules.[16] They proposed to grant Internet users a general right to use their broadband connections, and prohibit carriers from restricting this right. Separately, Wu studied NN and discriminatory behaviors in the broadband access market.[17] Wu acknowledged that some Internet applications required special QoS guarantees.[17] Wu also implicitly endorsed price discrimination in Internet services if such discrimination would not be based on application types.[17] More generally, he agreed that there existed "both justified and suspect bases of discrimination."[17]

To Lessig, however, the NN proposal was just a narrow case of his general belief that the Internet is a platform for innovation and should remain an "innovation commons."[18] Lessig based his belief partially on an end-to-end principle promoted by some early Internet architects, where the principle says that most of the Internet intelligence should exist at the edge of the network and within applications, rather than inside the network.[19] Although this principle was articulated in a purely technical context, Lessig asserted that it was a good public policy because it made the Internet neutral to applications and thus encouraged innovation.[18] Among other things, Lessig argued that the principle created a neutral platform "because the network owner can't discriminate against some packets while favoring others."[18] This is the most general statement of NN and it bans ISPs from prioritizing packets in any way; in particular, ISPs cannot use product differentiation to serve different markets, especially an emerging QoS market.

Numerous scholars have taken opposing or cautious views on NN, either from a policy perspective or from an economic perspective. Among them, professor Chris Yoo is the leading academic opposing NN. He has argued that NN will reduce ISPs' incentives to invest and innovate, that NN will defeat the QoS requirements from newer Internet applications, and that the end-to-end principle has been misread into the NN debate.[20] Professor James Speta has argued that regulations such as NN are unnecessary because ISPs have no incentives to discriminate against independent applications.[21] Adam Thierer has argued that the "dumb pipe" approach toward the Internet's architecture, as mandated by the end-to-end principle, is not a good public policy because it both oversimplifies the need of emerging Internet applications and discourages the development of competing infrastructures.[22] Professor Philip Weiser has proposed to let the FCC take an antitrust-like, *ex post* approach to ensure competition and prevent discrimination in the Internet service market.[23]

C. The Dilemma of ISPs

The chief problem that ISPs face is a pricing model crisis—they cannot capture the consumer surplus in an emerging QoS market. Due to various difficulties, ISPs were forced to adopt a flat-rate pricing model to sell their bandwidth in the early days. Commercial and individual customers have become used to this model for many years and now take it for granted; ISPs just have found that they are locked into this flat-rate model.

D. The Rights and Obligations of ICPs

Because of the ubiquitous flat-rate ISP service contracts and the rivalrous nature of bandwidth, the ever-increasing traffic volume generated by ICPs has indeed caused a "tragedy of the commons" problem for bandwidth consumption.[24] Despite this problem, the ISPs do not have the right to block the contents of those ICPs for two coupled reasons. First, this tragedy-of-the-commons problem does not result in a breach of the flat-rate contracts. Second, the ISPs have an implied-warranty duty under contract law not to block. On the other hand, the flat-rate contracts provide no guarantee that the ISPs cannot downgrade the services to a certain degree.

Concerns about ISPs' blocking behaviors, however, have been largely historical. The FCC set a precedent to ban such discriminatory behaviors in the *Madison River* case.[8] Nevertheless, ICPs tend to use those narrow and obsolete examples of blocking to launch a wholesale attack on QoS, and cloak their bandwidth tragedy-of-the-commons behaviors.

E. The Irrationalities of Consumers

Because the ISP market—in particular the broadband access market—has limited competition, consumers in general need some protection to deter market power abuse. But that is the traditional regime of antitrust laws. Many consumers are both the initiators and the victims of the bandwidth tragedy-of-the-commons problem, which is the essence of the term "tragedy": on one hand, they generate heavy traffic when they retrieve contents from ICPs; on the other hand, they downgrade the services of each other by their own bandwidth-hogging behaviors. For consumers, NN as a law is probably an over-kill because it is much stronger than antitrust laws. At the same time, NN is also a double-edged sword because it may prevent many useful QoS-based applications from taking off. Nevertheless, consumer advocacy groups, exemplified by SavetheInternet.com, decided to fight together with the ICPs to advocate NN.

III. WHAT IS NETWORK NEUTRALITY AND WHAT IS IT FOR?

For the NN debate, nothing is more paramount than agreeing on what "neutrality" means. Consider two packets at a router: a packet from an e-mail arriving slightly earlier than a packet from a tele-surgery application. Should the router send out the e-mail packet first? An e-mail message can wait for a short while, but a patient under a surgery cannot. So is this FIFO order neutral? Such a question inevitably asks for value judgment, but the example used here illustrates a point: NN cannot be debated in the abstract without considering the underlying engineering realities.

A. The Internet Has Never Been Neutral and Has Never Been Designed to Be Neutral

Contrary to what many NN proponents have asserted,[25] the Internet has never been neutral and has never been designed to be neutral. Simplifying the technically complex and elegant TCP/IP into a "dumb pipe" is both technically inaccurate and conceptually misleading for the NN debate. Examples of this non-neutrality abound, but this paper focuses on those that are most fundamental to the Internet.

In RFC 791, the RFC that was published in 1981 to define an IP packet, a Type of Service (TOS) field was defined for every IP packet.[26] This TOS field was designed to convey QoS information, such as "precedence," "delay," and "throughput." The field is mandatory and it takes one full byte, which is significant in protocol design. As explicitly indicated in RFC 791, the early Internet architects had QoS and packet prioritization in their minds seriously. IPv6, the newer version of IP published in 1995, emphasized QoS even more.[27] As FIFO scheduling does not need information in a TOS field, the dominance of FIFO on the Internet shows that TOS has not been utilized much. Such a result, however, was due to the facts that (1) FIFO is very simple; (2) it is very difficult to implement complex link schedulers; and (3) real-time applications emerged only recently. The result was not because of a "neutrality" principle.

Border Gateway Protocol (BGP) is the most important routing protocol for the Internet.[28] RFC 1105, the first RFC on BGP that was published in 1989, specified policy routing as a fundamental design goal.[29] In non-technical terms, this means AT&T routers can make discriminatory routing decisions such as treating traffic from Sprint more favorably than traffic from Verizon, or even rejecting Verizon traffic altogether. In practice, almost all routers from Cisco and Juniper, the two dominating router vendors that have consistently captured more than eighty percent of the world's Internet router market in the past, provide rich functions for ISPs to implement such routing policies on a daily basis.

TCP is not neutral either. TCP implements an elegant network congestion control algorithm.[3] The algorithm voluntarily reduces the data rate of an application upon detection of network congestion, even if that application is not an actual contributor to the congestion.[3] From that application's perspective, however, TCP is a lower layer and is part of the network.[3] Thus, the network indirectly discriminates the well-behaving application via TCP.

B. Neutrality as a Public Policy for the Internet

The social and economic dimensions of the NN debate center around Internet innovation. Proponents argue that application innovations, especially those from individuals or "garage" innovators, need NN protection because traffic prioritization may deny their access to the Internet completely. Opponents argue that NN will deter network innovation because it will discourage ISPs from investing in the network infrastructure.

A general dilemma about the relationship between a network and its applications can shed some light on the NN debate: do new applications drive the development of a better network or does a good network drive the development of newer applications? To support newer applications such as tele-surgery that have heavy traffic and strict real-time requirements, the network needs faster hardware, faster physical links, better algorithms, more sophisticated and more stable software, and possibly even a better architecture. Enhancing such network capacities requires significant investments in scientific research, engineering development, and large-scale network upgrades. Such investments can be justified only if newer applications are emerging to predictably make the investments profitable. On the other hand, developing major new applications usually takes considerable time and institutional resources; motivation to develop such applications will be seriously dampened if the network stops its evolution and does not technically support those applications. This is a classic chicken-and-egg problem.

The past evolution of the Internet, however, had a simple answer for this problem: both the network development and the applications development were incremental, and they drove each other in a positive feedback loop. More specifically, a few new applications such as the World Wide Web generated more traffic and greater demand for a faster network, which stimulated ISPs to build a somewhat but not revolutionarily better network. This marginally improved network gave the birth of a few even newer applications such as online stock trading and Internet chatting, which in turn stimulated the building of an even better network. This positive feedback loop continued to drive Internet evolution forward.

The NN debate is partially a chicken-and-egg dilemma in the following sense. The development of major QoS-oriented applications needs QoS support from the network, but uncontrolled QoS provision may, as those NN proponents have worried, stifle garage innovation if the innovators cannot receive meaningful bandwidth under the product differentiation regimes that would be in place. The evolution of the Internet has witnessed both institutional and garage-based innovations. From a technical perspective, the incremental nature of Internet evolution makes those garage innovations, which are typically smaller application innovations such as Wikipedia, particularly significant. From an economic perspective, the many application innovations discussed by Lessig[18] have not only directly driven up the demand side of the networking market, but also naturally generated the *network effect* that is invaluable to an information economy.[30] Although it is important to protect and encourage garage innovation, it is critical to sustain ISPs' incentives to invest so that major new applications, which typically require institutional efforts, will have a capable network as their platform. The NN proposals try to protect the garage innovation by banning QoS-based product differentiation or even traffic prioritization altogether. They solve one problem of the dilemma but worsen the other more serious problem. They may even defeat themselves in the sense that they will impede institutional application innovation. The key policy challenge in the NN debate is to strike a balance between incentives and monopolies. This challenge, however, is a familiar issue in many intellectual property laws.

IV. A TECHNICAL AND MIDDLE-GROUND SOLUTION

Based on the analyses above, a working middle-ground solution to the NN debate needs to (1) allow ISPs to serve an emerging QoS market; (2) sustain and encourage garage innovation; (3) give consumers meaningful protection; and (4) treat all ISP customers, including the ICPs, fairly. This paper now proposes such a solution and explains why it meets these four objectives.

A. Protecting Garage Innovation under QoS Provision

Robert Atkinson and professor Weiser published a moderate proposal in 2006 that addressed some of the objectives enumerated above.[31] Extending Weiser's earlier idea, they proposed to charge the FCC with an antitrust-like regulatory power to protect consumers. They also proposed to mandate ISPs to use "some not insignificant portion of the broadband bandwidth" to provide basic Internet services. This last idea and some more sophisticated versions of it, however, have been well known in the networking community for many years,[32] a fact suggesting that the current NN debate has not attracted enough attention from the technical community.

This paper argues that QoS provision can co-exist with garage innovation protection. More specifically, a certain fraction of network bandwidth can be reserved to protect garage innovation, and the rest of bandwidth can be used for QoS provision. This is technically feasible, as will be explained below.

B. How It Works

Understanding and appreciating the idea above requires a detailed discussion of link scheduling algorithms. As discussed earlier, a router controls packet queueing delays mainly via link schedulers. By controlling the sending-order of packets, a link scheduler effectively distributes the link bandwidth among applications. This can be better understood by studying the traffic control mechanism at freeway entrances in many metropolitan areas. At such an entrance, two or more ramps lead to a single on-ramp of a freeway. During rush hour, one of the ramps is an express lane for car pools. A traffic light controls the ramps and one car goes per green signal at the car's ramp. By controlling the interval lengths between the green signals at each ramp, the traffic light can assign different fractions of the highway passage to the ramps, and the car-pool ramp can receive a faster passage. However, any other ramp can still receive a fraction of the passage and will not be starved. In computer networking, such a scheduling scheme is known as Weighted Fair Queueing (WFQ), a breakthrough in QoS research.[3] Very sophisticated link schedulers based on WFQ can deliver very flexible QoS services, although it is in general difficult to implement any complex scheduling algorithms such as WFQ.[33]

In theory ISPs can dedicate all or most of their bandwidth to QoS provision; other applications not paying premiums may only be served on a "best effort" basis, which means their packets will consume the residual bandwidth, if any, in a FIFO order. The residual bandwidth can go down to zero in the worst case. This situation is similar to a highway on-ramp where no signal exists for the car-pool ramp and the signals at other ramps are always red if at least one car exists on the car-pool ramp. So in theory the car-pool ramp can take almost all the passage and starve the other ramps. Many NN proponents have challenged such a situation vigorously. Indeed, as in the blocking case, even for applications not paying premiums, ISPs have an implied-warranty duty under contract law to avoid such a starvation or near-starvation. With advanced link schedulers, however, ISPs can eliminate such starvation by reserving a non-trivial fraction of their bandwidth to provide the "typical" services of today, although all applications not paying premiums need to share this reserved bandwidth, and the bandwidth tragedy-of-the-commons problem may still exist among those applications.

C. A Counter-argument and a Rebuttal

NN proponents may argue that this bandwidth reservation scheme effectively downgrades the services of those non-premium-paying applications from their current levels. This argument, while valid, is economically misplaced, as explained below.

For a simplified illustration, assume that the reserved fraction of bandwidth is set at 50%. With the remaining 50% of bandwidth to serve the QoS market, ISPs can increase their profits and then invest in a faster Internet in response to greater QoS demands. With a QoS market taking off, such a feedback loop is positive and the ISPs could triple the capacity of the Internet within a certain period of time. This calculation is realistic because a QoS-enabling network will incubate many newer QoS-based applications demanding for larger network capacities. While this positive feedback occurs, the reserved bandwidth will also increase three-fold and will then be 50% larger than what it is now. In contrast, if the ISPs are discouraged from investing, the capacity of the Internet may stay relatively flat for a long time. Clearly, the proposed scheme not only can sustain garage innovation, but also can promote it via steadily driving Internet evolution forward. Thus, Internet traffic prioritization can both coexist with and encourage Internet innovation, including network innovation, institutional application innovation, and garage application innovation.

D. The Other Objectives to Be Achieved

Protecting consumers and enforcing fair dealing across all ISP customers are not problems if the market is competitive and has FCC policies as well as antitrust laws watching in the background, where competition will assure fair treatment of all customers. Perhaps new Internet service contracts with non-flat-rate billing will be written, but competition will prevent ISPs from overcharging specific customers. It is possible, and indeed likely, that when the market reaches its equilibrium, ICPs will pay more than what they do now, even without requesting QoS. This will be, however, because the ICPs currently are enjoying a historical pricing-model lock-in and treating the flat-rate bandwidth as a commons, not because they will receive discrimination in the future.

A counter-argument for this last analysis is that since the current competition in the broadband access market is limited, there is no guarantee of fair dealing. This paper argues otherwise. First, as other commentators have argued, current FCC policies (as suggested by the *Madison River* case), newer antitrust-like FCC policies, or even antitrust laws themselves can help enforce fair dealing. Second, the limited competition in the current broadband access market should not be taken as a given; newer broadband access technologies such as wireless, power-line, metro-Ethernet or optical-fiber are technically available now, although with small penetration rates and high initial costs. The policy-making focus should be on solving the competition problem by stimulating those new technologies to establish a more competitive market, rather than artificially neutralizing the problem by stifling the evolution of the Internet via regulation.

V. CONCLUSION

The network neutrality debate is complicated. Understanding it requires a solid understanding of the technical details of the Internet and some economic aspects of Internet evolution. As a living engineering miracle, the Internet has never been neutral and has never been designed to be neutral. Rather, it has been designed to be practical and it continues to evolve in a practical way. Many of the current arguments in the debate are misplaced, prejudiced or hyperbolic. The fundamental policy goal should be striking a balance between securing incentives for network innovation as well as institutional application innovation, and protecting garage application innovation. Driven by their respective financial interests, Internet service providers and Internet content providers essentially dispute, under the name of network neutrality, their legacy Internet service contracts, which are increasingly problematic with today's technical and economic realities on the Internet. This paper has proposed a technically feasible middle-ground solution to the debate. The solution is to use bandwidth reservation to protect garage innovation under QoS provision.

REFERENCES

[1] Anne Broache & Declan McCullagh, CNET News.com, *Playing favorites on the Net?*, Dec. 21, 2005, http://news.com.com/2100-1028_3-6003281.html

[2] Declan McCullagh, CNET News.com, *House plans vote on Net neutrality*, June 7, 2006, http://news.com.com/2100-1028_3-6080983.html

[3] JAMES F. KUROSE & KEITH W. ROSS, COMPUTER NETWORKING (3rd ed. 2005)

[4] Domenico Ferrari, *Client Requirements for Real-Time Communication Services*, 26 IEEE COMMUNICATIONS MAGAZINE 65

[5] IETF RFC 3935, *A Mission Statement for the IETF*, http://www.ietf.org/rfc/rfc3935.txt

[6] IETF RFC 4677, *The Tao of IETF*, http://www.ietf.org/rfc/rfc4677.txt; Active IETF Working Groups, http://www.ietf.org/html.charters/wg-dir.html

[7] RFC Index, http://www.ietf.org/iesg/1rfc_index.txt

[8] Declan McCullagh, CNET News.com, *Telco Agrees to Stop Blocking VoIP Calls*, Mar. 3, 2005, http://news.com.com/2100-7352_3-5598633.html; Federal Communications Commission, *Consent Decree*, DA 05-543, http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-05-543A2.pdf

[9] CBC News, *Telus Cuts Subscriber Access to Pro-Union Website*, July 24, 2005, http://www.cbc.ca/canada/story/2005/07/24/telus-sites050724.html

[10] Stefanie Olsen, CNET News.com, *AOL Charged with Blocking Opponents' E-mail*, Apr. 13, 2006, http://news.com.com/2100-1030_3-6061089.html

[11] Hiawatha Bray, *Telecoms Want Their Products to Travel on a Faster Internet*, BOSTON GLOBE, Dec. 13, 2005

[12] Arshad Mohammed, *Verizon Executive Calls for End to Google's "free lunch"*, WASHINGTON POST, Feb. 7, 2006; Anne Broache, CNET News.com, *Net Neutrality Debate Still Simmers*, July 11, 2006

[13] Anne Broache, CNET News.com, *Tech Manufacturers Rally against Net Neutrality*, Sept. 19, 2006, http://news.com.com/2100-1028_3-6117241.html

[14] SavetheInternet.com, http://www.savetheinternet.com

[15] Hands Off The Internet, http://www.handsoff.org

[16] Tim Wu & Lawrence Lessig, *Ex Parte Letter to the Federal Communications Commission*, CS Docket No. 02-52, Aug. 22, 2003, http://www.freepress.net/docs/wu_lessig_fcc.pdf

[17] Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. Telecomm. & High Tech. L. 141

[18] LAWRENCE LESSIG, THE FUTURE OF IDEAS (2001)

[19] J. H. Saltzer et al., *End-to-end Arguments in System Design*, 2 ACM TRANSACTIONS ON COMPUTER SYSTEMS 277 (1984)

[20] Christopher Yoo, *Would Mandating Broadband Network Neutrality Help Or Hurt Competition? A Comment on the End-to-End Debate*, 3 J. TELECOMM. & HIGH TECH. L. 23 (2004); Christopher Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1 (2005); Christopher Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847 (2006)

[21] James B. Speta, *Handicapping the Race for the Last Mile? A Critique of Open Access Rules for Broadband Platforms*, 17 YALE J. ON REG. 39 (2000)

[22] Adam D. Thierer, *Are "Dumb Pipe" Mandates Smart Public Policy? Vertical Integration, Network Neutrality, and the Network Layer Model*, 3. J. TELECOMM. & HIGH TECH. L. 275 (2005)

[23] Philip J. Weiser, *Toward A Next Generation Regulatory Regime*, 35 LOY. U. CHI. L.J. 41 (2003)

[24] Garrett Hardin, *The Tragedy of the Commons*, 162 SCIENCE 1243 (1968)

[25] Dynamic Platform Standards Project, *Introduction and Summary for Congressional Staff*, http://www.dpsproject.com/CongressSummary.html

[26] IETF RFC 791, *Internet Protocol* (1981), http://www.ietf.org/rfc/rfc0791.txt

[27] IETF RFC 2460, *Internet Protocol, Version 6 (IPv6) Specification* 25 (1998), http://www.ietf.org/rfc/rfc2460.txt

[28] RADIA PERLMAN, INTERCONNECTIONS: BRIDGES, ROUTERS, SWITCHES, AND INTERNETWORKING PROTOCOLS (2nd ed. 1999)

[29] IETF RFC 1105, *A Border Gateway Protocol (BGP)* (1981), http://www.ietf.org/rfc/rfc1105.txt

[30] CARL SHAPIRO & HAL R. VARIAN, INFORMATION RULES: A STRATEGIC GUIDE TO THE NETWORK ECONOMY (1998)

[31] Robert D. Atkinson & Philip J. Weiser, *A Third Way on Network Neutrality*, 13 THE NEW ATLANTIS 47 (2006)

[32] ATM Forum Technical Committee, *Traffic Management Specification Version 4.0* (1996), http://www.mfaforum.org/ftp/pub/approved-specs/af-tm-0056.000.pdf

[33] Itamar Elhanany et al., *Packet Scheduling in Next-Generation Multiterabit Networks*, 34 IEEE COMPUTER 104 (2001)