3-1-2011

Rough Consensus and Running Code: Integrating Engineering Principles into Internet Policy Debates

Christopher S. Yoo
University of Pennsylvania Law School, csyoo@law.upenn.edu

SYMPOSIUM INTRODUCTION

Rough Consensus and Running Code: Integrating Engineering Principles into Internet Policy Debates

Christopher S. Yoo*

On May 6–7, 2010, the University of Pennsylvania's Center for Technology, Innovation and Competition hosted the conference, "Rough Consensus and Running Code: Integrating Engineering Principles into the Internet Policy Debates."[1] This conference brought together members of

* Professor of Law, Communication, and Computer & Information Science and Founding Director of the Center for Technology, Innovation and Competition at the University of Pennsylvania. Special thanks to David Clark, Jonathan Smith, and Anna Gavin for their help in putting this conference together and to the staff of the *Federal Communications Law Journal* for their willingness to publish this special conference issue.

1. The full program and video of the panels are available at *Rough Consensus and Running Code: Integrating Engineering Principles into the Internet Policy Debates*, CENTER FOR TECH., INNOVATION & COMPETITION (2010), http://www.law.upenn.edu/cf/institutes/ctic/conferences/internetpolicy.html.

the engineering community, regulators, legal academics, and industry participants in an attempt to provide policymakers with a better understanding of the Internet's technical aspects and how they influence emerging issues of broadband policy.

At various points during the recent debates over broadband policy, observers both inside and the outside the government have acknowledged that the debate has yet to reflect a full appreciation of the engineering principles underlying the Internet and the technological opportunities and challenges posed by the existing architecture. The level of discourse is reminiscent of the days when economic arguments first began to be advanced in during regulatory proceedings, when participants in policy debates lacked a sufficient vocabulary and an understanding of the underlying intuitions to engage in a meaningful discourse about the relevant insights.

The conference's title, "Rough Consensus and Running Code,"[2] also emphasizes that network engineering has long been a pragmatic rather than a theoretical discipline that does not lend itself to abstract conclusions. Network engineers recognize that there is no such thing as the perfect protocol. Instead, optimal network design varies with the particular services, technologies, and flows associated with any particular scenario. In other words, network engineering is more about shades of gray than absolutes, with any solution being contingent on the particular circumstances and subject to change over time as the underlying context shifts. Policymaking is better served by an understanding of the relevant tradeoffs than by categorical endorsements of particular architectural structures as being the foundation for the Internet's success.

Another side effect of the lack of technical sophistication in the current debate is a tendency to defer to opinions advanced by leading members of the engineering community. People without technical backgrounds often regard strong statements of scientific conclusions as possessing a high degree of conclusiveness. Yet anyone who reads broadly in the technical literature quickly realizes that members of the engineering community often disagree sharply over the best way to move forward and that many seemingly authoritative declarations are actually positions in technical debates that are hotly contested and still ongoing. Just as in

2. For the seminal statement, see David Clark, *A Cloudy Crystal Ball – Visions of the Future*, 24 PROC. INTERNET ENGINEERING TASK FORCE 539, 543 (1992), https://www.ietf.org/proceedings/24.pdf ("We reject: kings, presidents and voting. We believe in: rough consensus and running code.").

economics and law, where there are often as many different positions as there are people offering opinions, so too in network engineering. At the same time, many areas over which policymakers are now struggling are regarded by the engineering community as completely uncontroversial and long settled.

Understanding how technical considerations should influence Internet policy thus requires a better understanding of the principles on which the Internet is based and an appreciation of the current areas of agreement and dispute within the engineering community. Toward this end, the conference program brought together engineers representing the full range of views on various issues currently confronting policymakers, as well as industry participants who have actual experience in deploying and running networks.

I. TUTORIAL

The conference began with a tutorial designed to provide an introduction to the basic engineering concepts underlying the Internet and to provide a flavor of the tradeoffs underlying the architectural choices. Major topics included the differences between host-to-host protocols, such as the Transmission Control Protocol (TCP) and the User Datagram Protocol (UDP); the edge-based approach currently used to manage network congestion, known as Additive Increase Multiplicative Decrease (AIMD); the deployment of active queue management techniques such as Random Early Discard (RED); the role of Classless Inter-Domain Routing (CIDR) to solve emerging routing problems; the challenges posed by network address translators (NATs); the role of the Border Gateway Protocol (BGP) in routing traffic; and the history of scheduling through techniques such as Integrated Services (IntServ), Differentiated Services (DiffServ), MultiProtocol Label Switching (MPLS), Explicit Congestion Notification (ECN), and emerging techniques such as Low Extra Delay Background Transport (LEDBAT). It offered some observations about current demands that the Internet is not designed to perform well, such as cost allocation, efficiency, security, mobility, and multicasting. It also offered some examples of how architectural decisions that are locally rational can create unexpected and potentially problematic interactions as traffic scales.

II. THE CONTINUING DEBATE OVER NETWORK MANAGEMENT AND QUALITY OF SERVICE

Over the past two decades, some engineers have proposed a series of enhancements to the Internet's architecture to provide more reliable quality of service than the current "best efforts" architecture permits.[3] Other engineers believe that instead of deploying new forms of network management, the better solution is simply to add more capacity.[4] This panel reexamined this debate in light of recent changes to the technological and competitive environment.

David Clark, who served as DARPA's chief protocol architect during the 1980s and currently serves as senior research scientist at the Computer Science and Artificial Intelligence Laboratory at MIT, expressed annoyance that the term "management" had been co-opted in the current debate, given that networks have always been managed. He also criticized the term "network neutrality" given that the Internet is not now and never has been neutral.[5] Instead, the issue is how to manage scarcity, which leads to congestion. Interestingly, the latency that degrades the performance of many time-sensitive applications is often caused by routers deployed by end users in their home networks (a phenomenon called "self congestion") in ways that is alleviated, but not eliminated, by increasing the bandwidth of the access link. It can also arise in other locations on a steady state or intermittent basis. Clark also indicated that concerns about strategic uses of discrimination to create artificial scarcity are overblown, in part because network providers do not need quality of service (QoS) techniques to create scarcity and in part because providing QoS would help innovation. The QoS techniques designed into the protocols that run the Internet ensure that decisions about prioritization are made by end users rather than network operators.

Deke Kassabian, senior technology director for networking and telecommunications at the University of Pennsylvania, described how

---

3. For textbook discussions of these proposals, see, e.g., 1 DOUGLAS E. COMER, INTERNETWORKING WITH TCP/IP: PRINCIPLES, PROTOCOLS, AND ARCHITECTURE 510–14 (5th ed. 2006); JAMES F. KUROSE & KEITH W. ROSS, COMPUTER NETWORKING: A TOP-DOWN APPROACH 602–04, 660–72 (5th ed. 2010).

4. *See, e.g.*, COMER, *supra* note 3, at 511; KUROSE & ROSS, *supra* note 3, at 603, 629–31.

5. *See* David Clark, Written Statement to the En Banc Public Hearing on Broadband Network Management Practices Before the FCC (Feb. 25, 2008), http://www.fcc.gov/broadband_network_management/022508/clark.pdf ("The Internet is not neutral, and has not been neutral for a long time.").

network architectures of large research universities are designed. Penn ensures that its user community has flexible and affordable access to network capacity by maintaining a private line connection to the nearest carrier hotel, where it can obtain easy access to a wide variety of service providers. In terms of performance management, Penn's basic approach is to add bandwidth rather than actively manage QoS. Penn does engage in some bandwidth management, however, by limiting students' Internet access on a per-address basis as well as capping the total amount available to students. Penn occasionally protects other users by limiting the bandwidth consumed by major research projects, sometimes diverting network intensive research projects onto Internet2's Interoperable On-demand Network (ION), which can establish dedicated circuits on a temporary basis.[6] In terms of security, rather than relying on a border firewall, Penn minimizes the impact on other users by deploying security as close as possible to the asset being protected through hardened server configurations, dedicated firewalls in front of a server, or broader use of authentication. Kassabian summarized the essence of this approach captured with the mantra, "open networks, closed servers, protected sessions."

Paul Dauby, vice president and chief operating officer of the Perry-Spencer Rural Telephone Cooperative (PSC), described the efforts of a remarkable rural cooperative serving six counties in southwest Indiana. Despite serving a territory with only 10.3 access lines per square mile and 2.98 subscribers per route mile, PSC supports a dazzling variety of services.[7] It offers digital subscriber line (DSL) service to all of its customers; fixed wireless broadband through unlicensed spectrum;[8] fiber-to-the-home to 560 customers in areas where it operates as a competitive local exchange carrier (CLEC);[9] limited multichannel video to its broadband customers via a virtual local area network (VLAN); and a ten-gigabit regional Ethernet transport that serves area hospitals. In order to make wireless broadband work on unlicensed spectrum, it limits the

6. In private conversations, Kassabian indicated that Penn also prioritizes traffic associated with public safety communications and environmental controls.

7. By way of comparison, Dauby indicated that if a city with the geographic footprint of Washington, D.C., had equivalent subscriber density as the service area in which PSC operates as an incumbent local exchange carrier (ILEC), it would only have seven hundred total subscribers.

8. PSC uses its wireless network for backhaul as well as for providing direct end user connections.

9. Dauby reports that PSC recently received a $29 million grant from the Rural Utilities Service to provide fiber-to-the-premises to its ILEC customers as well.

bandwidth available to peer-to-peer applications, restricting them to no more than ten sessions. PSC currently does not rate limit its wireline offerings despite the fact that it pays transit costs that are several times the cost in larger cities. The advent of over-the-top video is placing increasing financial pressure on their ability to continue its policy of nondiscrimination.

Paul Misener, vice president for global public policy at Amazon.com, remarked about what he saw as a surprising level of agreement on network neutrality. Specifically, both sides of the debate agree that openness is good, that a fair amount of concentration exists at the edges, and that switching costs restrict end users' ability to change providers. In addition, the industry had been in a state of détente during which few untoward activities had occurred, which he attributed to the network providers' fear of regulation. He argued that topological solutions—such as moving servers nearer to end users, buying private line service to closer interconnection points, and contracting with content distribution networks (CDNs) like Akamai---did not violate network neutrality so long as they involve new investments that are incremental to the facilities used to provide existing services. During the question and answer session, he argued that networks should be permitted to favor time sensitive applications such as voice over Internet protocol (VoIP) over less time sensitive applications such as file transfers.

III. CHANGING TECHNOLOGY AND THE LIMITS OF THE LAYERED AND END-TO-END MODELS

Network engineers have long explored alternatives to the layered, edge-based approach that dominates the network's current architecture.[10] This shift is motivated in part by one of the most distinctive characteristics of networks, specifically the interactions between individual flows and the underlying protocols as networks scale. It also reflects the emergence of management and security solutions that require the aggregation of information about the behavior of multiple endpoints and flows. This panel, chaired by the late W. David Sincoskie, professor of electrical and computer engineering and director of the Center for Information and Communication Sciences at the University of Delaware, who tragically

10. *See, e.g.*, R. Bush & D. Meyer, *Some Internet Architectural Guidelines and Philosophy*, IETF RFC 3439, at 7 (rel. Dec. 2002), http://tools.ietf.org/pdf/rfc3439; *The Rise of the Middle and the Future of End-to-End: Reflections on the Evolution of the Internet Architecture*, IETF RFC 3724 (J. Kempf & R. Austein eds., rel. Mar. 2004), http://tools.ietf.org/pdf/rfc3724.

passed away on October 20, 2010, explored the implications of those changes. Sincoskie shared anecdotes of his experiences in the telecommunications industry. He also offered the observation that the Internet is no longer end-to-end and that layering is an abstract concept that when strictly enforced does not perform well in reality.

Matt Mathis, who recently served as senior networking engineering specialist at the Pittsburgh Supercomputing Center, explained how new implementations designed to make TCP run faster are causing congestion in parts of the network. For example, the auto-tuning feature of Windows Vista, Windows 7, Linux, and Mac O/S causes end users running those operating systems to obtain a greater proportion of the available bandwidth than end users running older versions of Windows, such as Windows XP. In addition, TCP allocates bandwidth in inverse proportion to the roundtrip time of the underlying TCP connection. This allows end users located relatively close to their data to consume up to ninety percent of the capacity of the relevant link. Also, the new implementations are designed to expand their transmission windows until they fill all of the available links. Thus, unlike previous implementations of TCP, new implementations inevitably create congestion at some location in the network. This makes performance unstable and unpredictable and makes it extremely difficult for network providers to outbuild the load, particularly when applications are designed to prefetch data. The result is that the network has to play a more active role in allocating network capacity through techniques such as weighted fair queuing.

Jason Livingood, executive director for Internet systems engineering, National Engineering and Technical Operations, Comcast Cable Communications, noted the vehement disagreement among engineers over the relative merits of edge-based versus network-based solutions, pointing out that the decision the two approaches should not be regarded as a binary choice. Instead, engineering's emphasis on tradeoffs and optimality means that any particular solution makes sense for particular circumstances and is necessarily subject to change over time. He gave several examples of functions that previously were provided by the hosts operating at the edge of the network were migrating into the core—including cloud computing, antispam filtering, congestion management, security, and some type of relay to provide global access to content during the transition from IPv4 to IPv6. Other developments were shifting functions in the opposite direction, such as the Session Initiation Protocol (SIP), which was shifting primary responsibility for the functions traditionally associated with telephone

switches operating in the core of the network into the hosts operating at the edge.

Kevin Werbach, assistant professor of legal studies and business ethics at the Wharton School of the University of Pennsylvania, observed that the layered approach that the engineering community uses to frame network design contrasts sharply with the siloed, technology-specific approach reflected in the federal statutes governing communications law. In addition, he pointed out that the layered model does not prescribe certain architectures and that the real world frequently does not conform to the theoretical model. He identified several risks in the current debate, including superficially applying engineering concepts to policymaking, thinking in terms of absolutes, and oversimplifying. He also pointed out a number of ways in which the network has changed since the Internet's primary protocols were designed in the 1970s, including the growing importance of wireless networks, cloud computing, online gaming, video, the Internet of things, and the Internet as a platform for commerce, advertising, and media distribution. He called for a better understanding of the incentives of network players and the relationships between them, better translation of engineering principles into the legal discourse, and more complete data to serve as the basis for decisionmaking.

I served as the fourth panelist and began by pointing out engineers disagree sharply over the relative merits of layering and the end-to-end argument. Moreover, while the policy debate tends to equate layering with ensuring that the lower layers and the core of the network remain relatively "dumb," the engineering community tends to regard the layered stack as following an "hourglass" model that recognizes both that the upper and lower layers of the network are often quite complex and that only the middle layer primarily responsible for addressing that must be kept simple. In addition, contrary to what others suggest, layers do not operate completely independently. Many common protocols cross layers, and interactions across layers have led to the development of active queue management and other core-based solutions to ensure that network resources are allocated fairly. Moreover, because routers operating in the core of the network are able to see what multiple end users are doing, they are often in a better position to implement certain security and congestion management techniques. Lastly, protocol layering can create a design hierarchy that promotes innovations that are consistent with the hierarchy while simultaneously discouraging innovation that is inconsistent with the

hierarchy.[11]

IV. ARCHITECTURE AND NETWORK SECURITY

The engineering community has long recognized that the anonymity and connectionlessness of the Internet's original architecture limits the network's ability to meet end users' growing need for security. The conference's third panel, chaired by Matthew Blaze, associate professor of computer and information science at the University of Pennsylvania, explored ways in which the current architecture can support network security as well as technical changes under consideration that could enhance its ability to do so.

Andrea Matwyshyn, assistant professor of legal studies and business ethics at the Wharton School, emphasized the importance of taking human considerations into account when designing network security. Instead of reflexively regarding failures as the result of user error, exemplified by the oft-used acronym PEBKAC ("problem exists between keyboard and chair"),[12] security systems should take into account the fact that even the best intentioned end user is imperfect and should reflect the way people interact with technology. Network engineers should also assume that every security system can and will be broken, and they should proactively incorporate response plans for when this inevitably occurs. They should also remember that end users are capable of understanding how to respond to problems—if solutions are clearly explained to them. Network security would also be improved by more frequent interactions between engineers and lawyers, and by bearing in mind that security is governed by a wide range of competing legal regimes—including (but not limited to) contract, intellectual property, telecommunications regulation, and consumer protection laws.

Edward Felten, professor of computer science and public affairs and director of the Center for Information Technology Policy at Princeton University, analyzed the security implications of the decision to place functions in the network's endpoints or in the network's core. As an initial matter, Felten emphasized that end users are not the only endpoints and that many functions that end users regard as being in the network (such as cloud computing, email servers, and other third-party intermediaries) are, from

11. I expand on these ideas in Christopher S. Yoo, *Protocol Layering: A Study in Incorporating Engineering Insights into Internet Policy*, 60 DUKE L.J. (forthcoming May 2011).

12. Another commonly used acronym is PICNIC ("problem in chair, not in computer").

the standpoint of network architecture, simply other endpoints. Moreover, the most threatening and most visible security problems (including malware such as botnets and spyware, server attacks, and phishing and other attempts to deceive end users) generally arise on the end hosts. Network-oriented security threats exist, such as attacks on the routing infrastructure, and networks can analyze traffic to help detect security problems. That said, end hosts can view traffic after it has been unpacked from any archives, decompressed, decrypted, and reassembled. This often places them in a better position to implement security, especially because they can conduct dynamic analysis of code while it is operating instead of simply static code residing on a hard drive.

Jonathan Smith, Olga and Alberico Pompa Professor of Engineering and Applied Science at the University of Pennsylvania, noted that while the end-to-end argument decentralizes innovation, it also decentralizes responsibility for security enforcement. Moreover, the network's current bias toward allow-by-default facilitates connection, such a default may no longer be the correct architecture in a network that has become a distributed system increasingly populated by security threats. In addition, although layers create opacity that makes programming easier by reducing what a programmer needs to know about how other layers are configured, hiding information about what is going on in lower layers may possibly be problematic from the standpoint of trust. Smith identified a number of solutions that are not working, including passwords, public key infrastructures, software updates, measures to protect the routing infrastructure (such as IPSEC, DNSSEC, and BGPSEC), firewalls, intrusion detection systems, and rate throttling defenses. Instead of implementing these ineffective solutions, network architects should improve the infrastructure for authentication and attribution, build automated trust systems, provide for a degree of cross-layer transparency through structures such as a knowledge plane, shift to deny-by-default, and make both the edge and the core more extensible.

V. KEYNOTE ADDRESS BY PAUL MOCKAPETRIS

The dinner keynote address delivered by Paul Mockapetris, inventor of the domain name system and currently chairman and chief scientist at Nominum, Inc., noted that the success of the network is often attributed to what is often called Metcalfe's law, which holds that a network's utility is proportional to the square of the number endpoints in the network. This

implies that a network's value grows quadratically as it expands.[13] So long as the value grows faster than the cost, networks keep growing wonderfully. The problem is that in the modern world, being part of a larger network does not necessarily confer benefits to the extent that it provides connections to hackers and other security threats. One solution is to use the DNS to begin tracking reputation data about particular actors. Although some industry observers raise concerns about placing critical information that needs to be secure into the DNS, this objection overlooks that fact that critical information is already in the DNS. Although some people argue that smart DNS services deviate from the simplicity of the hourglass model often used to describe the Internet, in reality, we already have multiple hourglasses to deal with different types of transmission technologies.

Mockapetris closed by offering a few observations about network neutrality, arguing that it should be illegal for parties to give users applications that act against their interests without making clear what those applications are doing, wondering if such safeguards are best served by an architecture that does not reveal who is serving as the counterparty and market maker in any particular transaction, as is the case in the current network architecture.

VI. NEW APPLICATIONS, NEW CHALLENGES

Emerging applications, such as Internet protocol television (IPTV) and gaming, are placing demands on networks that are quite different from the flows generated by the applications that dominated the early Internet, such as email and web browsing. This panel, moderated by Saswati Sarkar, associate professor of electrical engineering at the University of Pennsylvania, explored the pressures that these new applications are creating on the network architecture as well as the technological options to adapt to these changes.

I provided an overview of the technical and policy challenges confronted by IPTV. Some IPTV providers employ dedicated or prioritized connections between the central office and the end users' premises. "Over-the-top" services, such as Netflix and YouTube, rely on the public Internet to transport their packets on a best efforts basis. Over-the-top services

13. *See* George Gilder, *Metcalf's Law and Legacy*, FORBES ASAP ARTICLES BY GEORGE GILDER, BASED ON CHAPTERS IN HIS FORTHCOMING BOOK – TELECOSM (Sept. 13, 1993), http://www.seas.upenn.edu/~gajl/ggindex.html (click on "Metcalfe's [sic] Law and Legacy" hyperlink).

employ a wide variety of techniques to provide the QoS needed to support video, including content delivery networks and adaptive streaming (which adjusts video resolution quality in light of the available bandwidth). In addition, IPTV providers must decide which platforms to support, both in terms of devices (such as PCs, Blu-ray players, gaming consoles, and smart phones) and encoding formats, which often incorporate varying maximum transfer rates. In order to obtain access to content, IPTV providers must also protect content against illegal copying, either through digital rights management (DRM) or filtering and must anticipate likely reactions to these measures, such as encryption, darknets, and greater exploitation of the analog hole. In addition, the growing importance of video has renewed interest in using multicasting to distribute mass media content. IPTV is also limited by legacy regulatory requirements, such as mandates for public, educational, and governmental (PEG) channels.

Paul Mitchell, general manager for regulatory and standards at Microsoft, discussed some the challenges confronted by game consoles such as the Xbox, Microsoft's effort to use high performance computing, home-theater quality graphics and audio, and network connectivity to provide an interactive, immersive game experience. The feature designed to allow users to communicate with each other while gaming drew the attention of regulators interested in determining whether this feature represented a telecommunications service.[14] Microsoft has now combined the Xbox with other products as services (such as Windows Phone 7, Microsoft Communicator, and the Kin smartphone) to allow voice communications and the sharing of video and audio across a wide variety of platforms. In many countries, however, regulatory restrictions prevent end users from taking advantage of the full range of these features. Another challenge is finding ways to make DRM interoperable. Regarding network neutrality, although that all networks are managed, they should be managed in predictable ways. Mitchell also provided a demonstration of adaptive streaming and described the challenges of supporting features such as closed captioning on a wide range of devices and encoding formats.

14. Xbox has also become a platform for distributing Netflix. In earlier conversations, Mitchell also discussed how regulators also inquired whether Xbox's Party Mode, which allows friends in separate locations to watch the same video at the same time, represented a cable service. Telephone Interview with Paul Mitchell, General Manager for Regulatory and Standards, Microsoft Corp. (Apr. 9, 2010). Microsoft has subsequently taken steps to turn the Xbox into a platform for subscription television service. Nick Eaton, *Microsoft Considering TV Service on Xbox*, MICROSOFT BLOG (Nov. 20, 2010, 11:05 AM), http://blog.seattlepi.com/microsoft/archives/229997.asp.

Joe Weinman, vice president for strategy and business development at AT&T Business Solutions, observed that the future demand for video distribution appears to be effectively insatiable, driven by new technologies such as ultra HD, multiscreen video for immersive virtual environments, 3D video, and the incorporation of video into social networking. At the same time, chip manufacturers are producing new products that make mobile video increasingly feasible. Other technologies that will increase the demand for bandwidth include Javascript and XML (Ajax), which triggers request for data when a mouse is moved or a keystroke is struck, such as popup information when a mouse hovers over a link. Other technologies that will increase the demand for bandwidth include sensor networks, cloud computing, and the emergence of households as de facto data centers in their own right. Solutions such as rate adaptation are useful stopgap measures, but may not work well when multiple users adapt in the same way at the same time. More problematically, rate adaptation addresses congestion by degrading the end users' experience rather than by ensuring that end users have access to the network resources needed to run highly interactive, latency-sensitive, and bandwidth-intensive applications.

Marjory Blumenthal, associate provost for academics at Georgetown University and former executive director of the National Academy of Sciences' Computer Science and Telecommunications Board, commented on all of the presentations. She noted the uncertainty implicit in the wide variety of predictions about the future of video, which range from the wildly optimistic to the severely pessimistic, and raised the possibility that adaptive technologies may represent a reasonably effective compromise that sufficiently preserves the end user experience. Regulatory requirements such as PEG can vary widely across different areas.[15] Others such as the Communications Assistance to Law Enforcement Act (CALEA)[16] can lead to unintended consequences.[17] In addition, the increasing cost effectiveness of filtering technologies, the ability to protect against illegal downloads through man-in-the-middle strategies, and the importance of proprietary DRM standards are changing the role of Internet service providers (ISPs). Lastly, the remote storage of data implicit in cloud computing puts someone other than the end user in charge of determining whether particular data is saved or lost, which can limit end users' control

15. *See* 47 U.S.C. § 531 (2006).

16. *Id.* §§ 1001–10.

17. *See, e.g.*, Daniel F. Spulber & Christopher S. Yoo, *On the Regulation of Networks as Complex Systems: A Graph Theory Approach*, 99 NW. U. L. REV. 1687, 1719 (2005).

over their own identities.

VII. THE FUTURE IS WIRELESS

As the FCC's proceeding on "Preserving the Open Internet" recognizes, wireless network face challenges that are quite different from wireline networks.[18] This panel, moderated by David Farber, distinguished career professor of computer science and public policy at Carnegie Mellon University, moved beyond the traditional focus on spectrum allocation to consider the unique management challenges that wireless networks confront, paying particular attention to how the physics of wave propagation, differences in network reliability, and the dynamic changes in the routing architecture associated with mobility often require wireless networks to employ network management techniques.

Dirk Grunwald, professor of computer science at the University of Colorado, discussed the difficulties inherent in the physics of wave propagation. Every frequency has different characteristics in terms of attenuation, absorption, and diffraction. Moreover, multipath reflections can cause the same signal to arrive at the same location along two different paths. If they arrive out of phase, they can cancel each other out in the same way that Bose headphones and sound dampening systems in cars operate. This causes signal quality to vary across time and space, demonstrated by how moving a car slightly can dramatically affect the quality of a radio signal. Engineers compensate for these variations by using different modulation schemes, which necessarily provide less bandwidth to distant locations. Network operators must decide in an environment that is constantly changing whether to equalize the performance of nearby and distant links rather than maximize total throughput. Differences in loss rates also affect the performance of TCP, because the average throughput rate is inversely proportional to the square root of the packet loss rate.[19] The solution may be to employ multiple solutions simultaneously allowing cognitive radios to maximize spectrum reuse.

Charles Jackson, a consultant who has previously held staff positions with the FCC, the U.S. House of Representatives, and the U.S. Commerce Department, addressed some of the network-based issues associated with

---

18. Preserving the Open Internet, *Report and Order*, 52 Comm. Reg. (P & F) 1, at paras. 86, 94–95, 103 (2010), *available at* http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf.

19. Matthew Mathis et al., *The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm*, COMPUTER COMM. REV., July 1997, at 67–68.

wireless networking. As an initial matter, wireless networks typically give voice communications priority over data traffic, which is typically less sensitive to latency. Preventing wireless networks from prioritizing in this manner either holding back reserve capacity that cannot be used for data transmissions or permitting voice service to degrade. The fact that radio links are less reliable than wireline connections has also led wireless networks to deploy smart-link technologies such as Automatic Repeater reQuest (ARQ) to shift responsibility for error recovery from the endpoints to the network. In addition, handset upgrades can often substitute for network investments, since receivers that are more sensitive require less capacity from base stations. Moreover, host-based congestion control depends on an honor system that is breaking down, which is causing networks to take a more active role in allocating bandwidth. Jackson also provided examples where traffic surges from Windows updates or earthquakes led ISPs to throttle certain types of traffic.

Robert Khedouri, chief executive officer of MusicGremlin, Inc., and vice president for services/strategy & planning for mobile network operators at SanDisk, described his experience launching the first MP3 player capable of downloading music directly from WiFi hotspots instead of sideloading it from a PC. MusicGremlin chose to adopt a "closed loop" system in which a single entity guaranteed secure delivery all the way from the content owner to the end user's device, similar to the manner in which Apple's iTunes establishes a closed loop between content owners and PCs. Relying on a closed, integrated system, complete with a vertically integrated music service, allowed MusicGremlin to provide the protection against piracy on which content providers insist. It also allowed the system to offer the value proposition to end users of ensuring seamless transfer with low latency. The company also deployed other bandwidth saving technologies, such as pushing content overnight to users who signed up for playlists, using burstable downloads to conserve on battery life, and caching the entire catalog of songs on every device to reduce search latency. MusicGremlin was acquired by SanDisk in 2008.

Christian Sandvig, associate professor of communication at the University of Illinois at Urbana-Champaign, noted that previous metaphors used to describe wireless technologies provide little insight into emerging aspects of spectrum, such as cognitive radios, smart antennas, and innovative forms of spectrum reuse. In addition, these metaphors fail to capture the variability and sensitivity to local conditions that make the performance of wireless networks so unpredictable, as illustrated by the

following example. While living in London, Sandvig deployed a directional antenna to provide WiFi service to the famous Speakers' Corner in Hyde Park,[20] only to find his signal intermittently negated despite the absence of any direct obstructions. The cause was double-decker buses stopped at a nearby traffic light, which periodically created a multipath reflection that cancelled out the direct signal. In addition, wireless networks face a tradeoff between making wireless devices easier to operate by hiding complexity and increasing wireless networks' configurability. On the one hand, the proliferation of wireless devices has turned consumers into overburdened band managers for their own houses. On the other hand, the advent of sensor networks and other technologies have made it easier than ever for them to adapt to local conditions.

* * *

The presentations and discussions at the conference represented a remarkable exploration of the issues that yielded fresh insights into issues of broadband policy. Indeed, former FCC Chief Economist Gerald Faulhaber congratulated the program for accomplishing something new in telecommunications policy, which he regarded as no mean feat.

The pages that follow contain articles by selected speakers exploring many of the themes raised during the conference. The conference proceedings and this special conference issue represent the first step in what we hope will be a new CTIC-led research initiative designed to better integrate the principles of network engineering into Internet policy debates.

20. For a description of this experiment, see PHILIP N. HOWARD, NEW MEDIA CAMPAIGNS AND THE MANAGED CITIZEN xi–xii (2006).