Traffic Management and Net Neutrality in Wireless Networks

Scott Jordan, *Member, IEEE*

***Abstract*—Many wireless ISPs limit the applications that may be used on wireless devices. In the United States, Congress is debating whether wireless network subscribers should have the right to use applications of their choice. We examine whether wireless ISPs should be able to limit applications. We address how wired and wireless networks differ with respect to traffic management, and conclude that wireless networks often require stronger traffic management than wired networks at and below the network layer. We propose dual goals of providing a level playing field between applications offered by ISPs and those offered by competing application providers and guaranteeing wireless ISPs the ability to reasonably manage wireless network resources. We consider three scenarios for how applications may be restricted on wireless networks, and find that none achieves both goals. We review United States communications law, and conclude that ISPs should be prohibited from giving themselves an unfair competitive edge by blocking applications or by denying QoS to competing application providers. We propose a set of regulations based on network architecture and communication law that limits an ISP's ability to restrict applications by requiring an open interface between network and transport layers. We illustrate how ISPs may deploy QoS within such a regulatory framework, and how this proposed policy can achieve our goals.**

***Index Terms*—Open wireless architecture; computer network management; telecommunication services; public policy; law.**

I. Introduction

THE academic literature on wireless networks focuses almost exclusively on network architecture, protocols, and applications. However, the variety of applications developed and deployed in wireless networks depends not only upon technical developments but also upon the actions taken by wireless Internet Service Providers (ISPs) and by wireless application providers. Currently, many wireless ISPs limit both the wireless devices that may be used by subscribers and the applications that may be used on wireless devices. Both of these types of limits have come under debate. Some believe that wireless network subscribers should have the right to use any compatible wireless device of their choice and the right to use any application of their choice.

Manuscript received November 30, 2010; revised May 12, 2011. The associate editor coordinating the review of this paper and approving it for publication was D. Hutchison.

S. Jordan is with the Department of Computer Science, University of California, Irvine (e-mail: sjordan@uci.edu).

Portions of this work have appeared in the 2011 IEEE Conference on Wireless Communications and Networking Conference. Portions of this work also appear in [1]; [1] is directed at the policy community and focuses on development of statute language; this article is directed at the networking community and focuses on the relationship of policy to network architecture. This work has been supported by the National Science Foundation.

Digital Object Identifier 10.1109/TNSM.2011.110311.100093

In the United States, these questions are currently being debated within Congress. Congress has the power to pass laws that determine what behavior is allowed by communications providers, and the Federal Communications Commission (FCC) has the power to create detailed regulations to implement laws passed by Congress. Similar questions are also being considered in Canada, Europe, Japan, and elsewhere (see e.g. [2]), but the different regulatory approaches of these countries may necessitate different approaches and are thus outside the scope of this paper.

We believe that this debate should be informed by relevant aspects of wireless network architecture. This paper addresses the question of whether wireless subscribers should have the right to run applications of their choice on their wireless devices, or whether and to what extent wireless ISPs should be able to limit applications or their behavior. To address these questions, we base our reasoning on both wireless network architecture and communications law.

To policymakers, this debate is a part of a larger debate over *net neutrality*. Net neutrality represents the idea that Internet users are entitled to service that does not discriminate on the basis of source, destination, or ownership of Internet traffic. Net neutrality has typically been discussed in the context of wired networks, with a focus on the wired public Internet. Those who support wired net neutrality believe that there is a danger that ISPs who offer applications may discriminate in favor of themselves over competing application providers, see e.g. [3][4]. Examples of such vertical integration may include cable ISPs that discriminate in favor of their own Voice-over-IP (VoIP) packets over competing VoIP providers' packets, and telephone ISPs that discriminate in favor of their own video over IP packets over competing video providers' packets. Those who oppose wired net neutrality believe that any such danger does not represent a market failure and that net neutrality regulation will be counterproductive, see e.g. [5]. There are also some who argue for a co-regulatory solution wherein a regulatory body (e.g. the FCC) should set forth broad terms that should govern cooperation, but should abstain from detailed regulation and instead oversee the work of a private sector self-regulatory body [6][2]. Good overviews of the arguments on both sides can be found in [7], [8], and [9].

Recently, the question has arisen over whether and how net neutrality should apply to wireless networks. In wireless networks, similar anticompetitive concerns may apply. Cellular network ISPs may have the incentive to discriminate in favor of their own video packets over competing video providers' packets. In addition, they may have the incentive to discriminate against any applications that compete with their

primary revenue streams, including competing voice and text-messaging applications that run over the Internet Protocol (IP).

In 2009, the FCC proposed a set of regulations to implement a form of net neutrality. The proposed regulations are included in a document called the *Notice of Proposed Rulemaking on Open Internet Practices* [10] (hereafter referred to as the NPRM). The NPRM not only states the proposed regulations, it also asks the public to comment on many aspects of them in order to provide the FCC advice before the regulations are finalized. Of interest here is that the NPRM asks many questions about whether and how to apply net neutrality to wireless networks. It asks whether net neutrality should be applied differently to wireless networks than to wired ones. It notes that wireless networks face special challenges due to attenuation and interference, and that they determine how users and devices share scarce resources through control over the frequency, time, and power of wireless devices' transmissions. It asks whether wireless devices and/or wireless networks merit different treatment. In 2010, the FCC issued the set of final regulations in the Open Internet Order [11]. The regulations, discussed in more detail later in this paper, distinguish between *fixed wireless* networks and *mobile wireless* networks.

With respect to device attachment, the NPRM asks whether subscribers should be able to connect wireless devices of their choice providing that they do not "harm the network". It ponders whether wireless ISPs should allow attachment of any device with a compatible air interface, including tethering, and if so how to prevent harm.

With respect to services, the NPRM considers which applications or services should be covered by a net neutrality requirement. It considers whether to exclude *specialized services*, defined as IP-based offerings such as voice and subscription video provided over the same networks used for broadband Internet access. It proposes that a nondiscrimination principle should only apply to Internet services, and thus exclude voice and short messaging services (SMS). However, it ponders what to do in 4th generation (4G) wireless networks capable of supporting voice, video, and data services on a converged platform architecture.

With respect to traffic management, the NPRM notes that wireless capacity may be more limited than wired capacity, and that demands can vary dynamically and widely among users. It discusses that wireless networks must be designed to deal with wide variations in signal levels across the service area as well as interference from other devices. It ponders whether wireless networks are more sensitive to user behavior. It asks how these differences in technical characteristics affect the reasonableness of various network management practices, e.g. whether it is reasonable for a wireless ISP to block capacity hungry applications. This inquiry builds upon a previous inquiry by the FCC on what traffic management practices are reasonable forms of network management [12], which had been motivated by Comcast's practice in 2007 of using reset packets to terminate selected peer-to-peer connections [13]. The NPRM also asks what impact tethering will have on wireless network congestion, and what network management measures are reasonable in this context.

The key question that this paper attempts to address is whether wireless net neutrality should be different than wired net neutrality because of the different technologies used. There is little academic literature that directly addresses net neutrality in wireless networks. Wu [14] started much of the current debate, where he focused on whether subscribers should be able to attach wireless devices of their choice. Wu argued for the extension of the FCC's Carterfone rules [15] to wireless networks. The Carterfone rules (as later implemented by Part 68 of the FCC's rules) required telephone companies to allow subscribers to use customer premises equipment (e.g. telephones) of their choice on the public switched telephone network unless the equipment harms the network. Wu's proposed extension of these rules to wireless networks include a prohibition on locking of devices to a carrier and allowing attachment of compatible and non-harmful devices. To allow such attachment, he proposes that industry or the FCC should define a basic air interface for wireless devices. Wu also argues for the application of net neutrality to wireless networks (which at the time meant application of the FCC's Internet Policy Statement [16]), and states that carriers should meter and charge for bandwidth usage rather than block particular applications. Wu also argues for wireless carrier disclosure of limits, including locks, protocol or application disabling, and bandwidth limits. Finally, he recommends that carriers and equipment manufacturers should work towards standardization of application development platforms.

In response, Hahn et. al. [17] claim that attachment of devices and Quality-of-Service (QoS) are separate issues. Having previously opposed net neutrality as a method to regulate QoS [18], in this paper they argue against many of Wu's proposals. First they argue that there is sufficient wireless competition to avoid market failure and that innovation in wireless devices and applications is thriving. Next they provide an economic analysis and argue that the results show that the benefits of device subsidies, device exclusivity, and limits on devices and on applications outweigh the costs of each.

Both of these papers focus primarily on the device attachment issue. Neither focuses on the *differences in traffic management* between wired and wireless networks, and hence on potential differences with respect to QoS. Recently, Marsden [19] discusses a number of outstanding issues of wireless net neutrality.

This paper focuses on the question of whether wireless subscribers should have the right to run applications of their choice on their wireless devices, or whether and to what extent wireless ISPs should be able to limit applications or their behavior. The analysis is based on both network architecture and communications law.

A few basic hypotheses are necessary to focus the analysis. First, the paper focuses on whether some form of net neutrality is required to ensure a level playing field between application providers who also serve as ISPs and application providers who do not serve as ISPs; other rationales for net neutrality are not considered here. Second, it is assumed that the primary method of potential discrimination is the use of QoS mechanisms such as packet prioritization or bandwidth reservation.

The primary focus is thus placed on applications and traffic management, rather than device attachment. We are concerned

with which applications or services should be covered by a net neutrality requirement, and whether this requires carving out an exception for *specialized services*. We are concerned with whether the challenges of wireless signals and mobility merit different traffic management techniques, and how these techniques may affect net neutrality.

There seem to be several parts of this issue. First, should *reasonable network management* or *specialized services* be defined differently for wireless networks than for wired networks? If so, do these differences in definitions merit *different treatment* with respect to net neutrality? The resolution of this issue depends in part on whether wireless networks require qualitatively different *types of network management* than wired networks. Second, should wireless network operators have a different ability to *restrict applications* used on their networks than wired ISPs? The resolution of this issue depends in part on whether applications have a greater ability to negatively interfere with desired network operation in wireless networks than in wired networks. Third, should wireless network operators have a different ability to *restrict devices* used on their networks than wired ISPs? The resolution of this issue depends in part on whether wireless devices have a greater ability to negatively interfere with desired network operation than their wired counterparts.

The focus in this paper is on the first two issues (reasonable network management and restricting applications); the third issue (restricting devices) requires additional consideration of legal issues of interconnection that is outside the scope considered here. We focus here on the technical aspects of these issues. With respect to all three issues, there are clearly economic and legal aspects that should be considered. For instance, competition may be substantially different in wireless networks than in wired networks. However, these economic and legal aspects are beyond the scope of this paper.

Section II reviews the pertinent aspects of wireless network architecture and addresses how wired and wireless networks differ with respect to traffic management. We find that although wireless networks often require stronger forms of traffic management than wired networks, these differences occur at the network layer and below. Section III reviews the pertinent communications law required to analyze the issue. We look to communications law for guidance in what types of network management and traffic discrimination should be allowed or prohibited. We conclude that ISPs should be prohibited from giving themselves an unfair competitive edge by blocking applications that compete with their own or by denying access to or charging unfair prices for QoS to competing application providers. Section IV lays out three scenarios currently under consideration for how applications may or may not be restricted on wireless networks. We find that none accomplishes our dual goals of providing a level playing field between applications offered by ISPs and those offered by competing application providers and guaranteeing wireless ISPs the ability to reasonably manage wireless network resources.

In section V, we propose a set of regulations that limits a wireless ISP's ability to restrict applications. The new policy is based on the layered structure of wireless networks and on communications law. We argue that since the differences between wired and wireless networks lie in lower layers, net neutrality in both wired and wireless networks can be effectively accomplished by requiring an open interface between network and transport layers, and that such a layer interface is consistent with current law. We illustrate how ISPs may deploy QoS within such a regulatory framework, and how pricing can be used to define a contract with users that removes the need for network control over user applications. Finally, section VI compares our proposed policy to the three scenarios discussed previously. We find that none of these three scenarios effectively accomplishes both goals, while our proposed policy does.

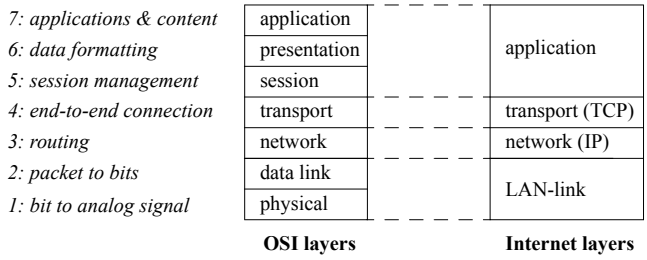

Fig. 1. OSI and Internet layered models.

II. Traffic Management in Wireless Networks

In this section, we review the pertinent aspects of wireless network architecture, address how wired and wireless networks differ with respect to traffic management, and discuss how these differences impact the issue of net neutrality.

A. Network architecture

We will consider differences between wireless and wired networks by the layer, starting at the bottom. We will denote layers using the OSI layer model [20], which is the reference model for layered architectures developed by the International Standards Organization, as pictured in figure 1.

The design of physical layer protocols is very tightly connected to the transmission medium, since each transmission medium has different characteristics of how signals propagate through the medium [21]. Wireless networks often face greater challenges from attenuation, noise and interference than do wired networks. In response, wireless networks often use more complex methods for modulation. Wireless networks also often use modulation schemes that support more complex multiple access, e.g. code division multiple access or a combination of time and frequency division multiple access. Finally, wireless networks use power control to cope with fluctuating attenuation and multipath levels.

The design of data link layer protocols also differs substantially between wired and wireless networks [22]. Interference from other users in wireless networks is usually worse than in wired networks, because there are typically more wireless devices within hearing distance of each other than wired devices share a wire. As a result, the data link layer protocol is also carefully matched to the type of network, including whether it is wired or wireless.

Portions of the design of network layer protocols may depend on whether the network is wired or wireless. Both

types of networks use IP network layer protocols and/or Signaling System 7 (SS7) network layer protocols for addressing and routing. Wireless networks, however, have additional network layer tasks. Portions of Radio Resource Management (RRM) are implemented at the network layer to allocate wireless network resources amongst cells. In addition, wireless networks with mobile users must accommodate handoffs.

In contrast, the design of the transport through application layers is largely similar in wired and wireless networks. Both types of networks use TCP and/or SS7 transport layer protocols for flow control, although sometimes variants of Internet protocols are created for wireless implementation, e.g. variants of TCP to cope with wireless attenuation and multipath. Similarly, both types of networks use higher layer Internet and/or SS7 protocols for managing calls or connections, authenticating users, presenting data, and interfacing to applications.

In summary, wireless networks differ substantially from wired networks at the network layer and below. However, they differ in much more limited manners at transport layer and above.

B. Traffic management

We now turn more specifically to traffic management practices in wired versus wireless networks. The traffic management requirements of wireless networks can differ from those of wired networks for two broad reasons. First, wireless networks face greater challenges due to the wireless medium and often due to mobility of their users. These challenges include attenuation, multipath, interference, and handoffs. Second, many wireless networks (especially cell phone networks) rely heavily upon their ability to offer satisfactory performance for telephone calls.

As a consequence, wireless networks use several traffic management techniques that are uncommon in wired networks. To accommodate voice calls, cell phone networks reserve capacity for each call or connection or otherwise limit the total traffic in the network, so that the negative impact of users upon each other is controlled and limited [23]. Capacity reservation for voice calls is accomplished using a few different techniques. First, power control is used to maintain a constant quality in the connection. Second, the network limits the number of number of voice users in each cell in order to maintain a minimum performance level per user. When a user migrates from one cell to another, the network attempts to allocate resources in the next cell; if resources are not available in the next cell and if the user requires too much transmission power to maintain a connection with the current cell (thus causing too much interference for its neighbors), the call is terminated. None of these techniques are yet common in the wired Internet.

These wireless traffic management techniques were designed for the dominant application on cell phone networks – telephone calls. However, there is a long term trend toward technology convergence. Both texting and Internet access are now key applications supported by cell phone networks, and subscribers pay significant sums for each. Cell phone networks are thus migrating in their architecture to look more like the Internet. In parallel, on the Internet real-time applications such as telephone calls, video conferencing, and gaming are growing in popularity. In response, Internet architecture is expected to migrate in its architecture to more efficiently support these applications. Cable tv networks now support both Internet access and telephone calls. Its architecture is also migrating toward one that can gracefully support all three sets of applications.

In order to support texting and Internet access on cell phone networks, and to create wireless local area network protocols such as Wi-Fi, traffic management techniques were required that are appropriate for delay-insensitive applications on wireless networks [24]. These techniques have some commonality with traffic management techniques used in wired networks for delay-insensitive applications, but they must also cope with the variability of the wireless medium. Thus, they also borrow ideas from traffic management techniques used in wireless networks for real-time applications. Like their wired network brethren, wireless traffic management for delay-insensitive applications allows for a wide variety of throughput per connection, so that a wide variety of applications can be efficiently supported. This is often accomplished using scheduling in the data link layer protocol. Like their wireless real-time application brethren, wireless traffic management for delay-insensitive applications dynamically adjusts transmission power on the basis of attenuation and multipath, and may use RRM to balance load between cells and to support handoffs.

It is thus true that wireless networks require some different types of traffic management than wired networks – due to both the nature of the wireless medium and the greater focus of some wireless networks on real-time applications. In wireless networks, users compete for resources with other users within hearing range, rather than only on the same wire. The greater variability of signals transmitted wirelessly instead of on wires would render real-time applications useless without some type of QoS implemented to smooth out the variations. A portion of these QoS techniques must be applied in very quick response to variations in the wireless signal, and thus must be automated in the wireless device and in network equipment.

In addition, wireless network capacity is usually more expensive than wired network capacity [23]. The cost of wired networks is usually dominated by the cost of purchasing and installing transmission lines, and to a lesser extent by the cost of network devices such as routers. Wireless networks incur similar costs, with the cost of transmission lines replaced by the cost of obtaining spectrum. In addition, wireless networks may incur costs for leasing sites for base stations. However, the capacity of wireless networks is almost always significantly less than the capacity of wired networks that cover the same geography, due to the relative scarcity of spectrum. This decreased capacity translates into a higher shadow cost for bandwidth[1]. As a result, there is often an incentive for traffic management techniques in wireless networks to be more efficient than in wired networks.

[1]Note, however, that in low density areas wireless networks may be lower cost than wired networks, due to the ability of a single wireless base station to replace a large number of long transmission lines.

C. Impact on net neutrality

We now turn more specifically to how these differences impact the issue of net neutrality. We conclude that wireless networks often are justified in using different traffic management techniques than wired networks, but only at and below the network layer. It follows that these differences merit a definition of *reasonable network management* that recognizes the differences in lower layer traffic management requirements and techniques. However, it also follows that lack of differences in upper layers merits a definition of *reasonable network management* that enables competition at the application layer in both types of networks. Wireless networks should address their greater challenges at lower layers either by exercising stronger lower layer traffic management techniques than used by wired networks or by exercising reasonable techniques used in wired networks but to a greater extent. However, wireless networks are not justified (by technical differences) in implementing different traffic management practices above the network layer than those used in wired networks.

Any net neutrality position applied to wireless networks, whether pro or con, should reflect the differences between wired and wireless networks at or below the network layer, and should reflect the similarities between wired and wireless networks above the network layer. Proponents of wireless net neutrality should accept that wireless networks require stronger forms of traffic management at or below the network layer, and should focus on regulation that ensures a level playing field between providers of various applications (including those that require QoS). Opponents of wireless net neutrality should accept that stronger traffic management at or below the network layer does not justify different treatment above the network layer, and should focus on defining reasonable network management in a manner that acknowledges this.

Since such lower layer techniques include the reservation and prioritization methods that generated much of the initial net neutrality debate, it remains a challenge to construct a net neutrality policy that can be consistently applied to both wired and wireless networks. We turn to this challenge in the remaining sections of the paper.

III. United States Communications Law and Net Neutrality

Before proceeding further, it would be helpful to most readers to give some background of communications law in the United States, and the arguments for and against net neutrality.

A. Communications law

U.S. federal communications law was separately developed for telephone networks and for cable video networks. Only recently has communications policy started to address the Internet. The relevant Federal law is contained in the Communications Act of 1934, as updated by various Acts since its original introduction. The Act is separated into titles, with Title I addressing general matters, Title II addressing telephone networks, and Title VI addressing cable video networks.

Title I is largely concerned with establishing the FCC and its operation. The FCC has the authority to create regulations that interpret and implement laws passed by Congress relating to communications. Title I, however, has also been interpreted by the courts to give the FCC limited ancillary authority to create regulations on types of communications not covered by the other titles. As a result, whenever communications are not deemed to fall under other titles, the only authority the FCC has stems from Title I. In addition, Title I includes a provision which instructs the FCC to forbear from applying regulations unless they are in the public interest and required to ensure just and reasonable practices.

Title II states much more detailed provisions applicable to common carriers, namely telephone companies. Before 1996, common carrier telecommunications policy was based on the assumption that local phone service was offered by a set of local monopoly carriers[2]. The Telecommunications Act of 1996 [27] removed barriers between local and long distance telephone service, and allowed competition between local and long distance carriers. However, communications law regarding common carriers has always principally focused only on voice service, and has not effectively addressed video or data service. Section 201 of Title II requires common carriers to offer service upon request and to interconnect with other carriers, and mandates that charges and practices be just and reasonable. Section 202 bars common carriers from unjust or unreasonable discrimination and from giving undue or unreasonable preference. Section 203 requires common carriers to post public lists of their charges. Section 205 gives the FCC authority to proscribe just and reasonable charges. Sections 251 and 252 regulate interconnection, and require an incumbent local exchange carrier to interconnect with other carriers on rates, terms, and conditions that are just, reasonable, and nondiscriminatory and in a manner that is at least equal in quality to that provided to itself. In addition, there are many other provisions in Title II that are not pertinent to the discussion here.

Title VI regulates cable carriers; it was largely laid down in the Cable Communications Policy Act of 1984 [28] and the Cable Television Consumer Protection Act of 1992 [29]. The policy was based on the likelihood of a set of local monopoly carriers; there was little video competition before the introduction of direct broadcast satellite video service. Communications law regarding cable carriers has always principally focused only on video service, and has not effectively addressed voice or data service. While most of the title is not pertinent to the discussion here, section 628 discusses the relationship between cable carriers and content providers. In particular, it addresses cases of vertical integration, in which a cable carrier has an attributable interest in a content provider. The section prohibits unfair methods of competition or unfair acts or practices, and certain types of discrimination in prices, terms, and conditions of sale.

In contrast, there is no similar body of federal communications law that directly addresses Internet access and services. In the absence of explicit statutes, it has been up to the FCC to determine the applicability of the Communications Act to the Internet. During the last forty years, this issue has come up repeatedly. In 1970, the FCC wrestled with the

[2]Common carriage law predates the Communications Act of 1934; see e.g. [25] and [26] for background on common carriage law and competition law and their relationship to net neutrality.

convergence between communications and computation in the consideration of the applicability of the Communications Act to data processing services. They implemented regulations [30] that delineated between data processing services and communication services on the basis of the device, i.e. whether a computer was used for communication or for the processing of information. The communications market continued to be subject to the common carrier provisions in the Communications Act, while the data processing services market was viewed as competitive and thus only subject to Title I. Telephone carriers who wanted to enter the data processing market were mandated to establish separate subsidiaries for this purpose.

By 1980, convergence between communication and computation had proceeded to the point where delineation on the basis of device had become unworkable. In response, the FCC replaced this framework with regulations [31] that delineated on the basis of the service. *Basic services* were defined as pure transmission capabilities and the data processing capabilities required to support them, whereas *enhanced services* were defined as processing that changed the format of the information or provided additional information. In the Telecommunications Act of 1996, the terms changed from *basic services* to *telecommunication services* and from *enhanced services* to *information services*, but the delineation was similar. Telecommunication services continued to be regulated by Title II. In contrast, telephone carriers who wanted to offer information services were only mandated to provide a certain degree of open access to competing providers of these services under Title I.

In a sense, therefore, until recently the portion of the Internet that resembles (or is identical to) telephone networks has often been considered to be within the domain of telephone common carrier law. However, the portion of the Internet that resembles applications has been considered to be outside this domain. However, in the last decade, the FCC declared that Internet access is an information service, and is thus not subject to common carrier regulation [32], [33]. This decision removed from Internet access several prohibitions on discrimination included in common carrier regulation. Similarly, Title II applies to telephone services offered by wireless carriers, but does not apply to wireless data services.

B. Net neutrality

This lifting of title II constraints on unreasonable discrimination, together with the development of QoS, triggered the push for net neutrality. Proponents of net neutrality (generally, application providers and consumer groups) argue that without a prohibition on discrimination, ISPs may charge application providers discriminatory prices for access to dedicated bandwidth or for QoS, or may outright block access to certain applications or websites, and that such activity will inhibit development of new Internet applications (see e.g. [34], [35], and [36]). Some proponents believe that ISPs should not be allowed to charge for priority treatment of traffic on the Internet portion of their service offerings (see e.g. [35]). When applying their position to wireless networks, many proponents of net neutrality argue that wireless networks face the same dangers as wired networks (see e.g. [35] and [36]). They are particularly concerned when wireless ISPs restrict the applications used on wireless devices, especially voice and video over IP. Many proponents believe that wireless providers should not restrict applications and should treat all applications equally (see e.g. [35], [36], and [34]).

Opponents of net neutrality (generally, ISPs) argue that there is no current problem, that competition is sufficient to ensure that commercially negotiated arrangements for bandwidth or QoS will not negatively impact consumers, and that any regulation will discourage investment in network infrastructure (see e.g. [37] and [38]). When applying their position to wireless networks, opponents of net neutrality see few differences. Opponents argue that there is a greater need for traffic management on wireless networks than on wired networks, and that this further undermines the case for wireless net neutrality (see e.g. [37] and [39]). In addition, they argue that there is greater competition amongst wireless providers, and thus even less need for net neutrality regulation (see e.g. [39] and [38]).

Some of the questions related to net neutrality take on additional importance in wireless networks. In particular, we focus on the question of whether users have the right to run any software of their choice on wireless devices. We look to communications law for guidance as to what types of network management and traffic discrimination should be allowed or prohibited.

Title VI of the Communications Act prohibits cable carriers who vertically integrate with content providers from using unfair methods of competition and from using certain types of discrimination in prices, terms, and conditions of sale. These same concerns apply to the Internet, both in its wired and wireless versions. In particular, we see the principal danger coming from potentially unfair methods of competition between application providers and ISPs who also offer applications. ISPs who offer applications based on QoS may give themselves an unfair competitive edge by denying access to QoS to application providers.

Title II of the Communications Act prohibits common carriers from using unjust or unreasonable discrimination and from giving undue or unreasonable preference. It also mandates that charges and practices be just and reasonable. These principles should also apply to the Internet, including to wireless Internet Service Providers. Many forms of traffic management, including QoS and application blocking, can be interpreted as forms of discrimination or preference. These traffic management practices and any charges associated with them should conform to these principles. The key test here has historically been one of "reasonableness", which is interpreted and implemented by the courts and the FCC. This same test should apply here.

Title I requires the FCC to forbear from applying regulations unless they are in the public interest and required to ensure just and reasonable practices. This forbearance principle should also apply to the Internet. Thus, regulation should not be applied if and when the wireless Internet access market becomes sufficiently competitive such that the dangers of vertical integration disappear.

We conclude that ISPs should be prohibited from giving

themselves an unfair competitive edge by blocking applications that compete with their own or by denying access to or charging unfair prices for QoS to competing application providers. Outside of these concerns, we find that communications law allows communications providers latitude to determine their own network management practices.

IV. Scenarios

In this section, we present and discuss three scenarios currently under consideration for how applications may or may not be restricted on wireless networks. The first two scenarios broadly represent proposals by those who oppose wireless net neutrality and those who support wireless net neutrality. The third scenario represents the FCC Open Internet Order. There are many other possible wireless net neutrality policies other than these three.

A. Scenario 1: No restrictions on traffic management

In the first scenario, commonly supported by those who oppose wireless net neutrality, ISPs are allowed to restrict applications at their discretion. A user's terms for cellular data service often include restrictions on the applications that may be used in conjunction with the service. For instance, both the AT&T Wireless data service and the Verizon Wireless mobile broadband terms and conditions proscribe permitted and prohibited uses[3]. Both AT&T Wireless and Verizon Wireless explicitly *permit* Internet browsing, email, and intranet access. Both AT&T Wireless and Verizon Wireless explicitly *prohibit* certain applications, including "server devices or host computer applications", "auto-responders, cancel-bots, or similar automated or manual routines which generate excessive amounts of net traffic or which disrupt net user groups or email use by others", and "software that maintains continuous active Internet connections when a computer's connection would otherwise be idle or any keep alive functions" unless they adhere to the ISP's data retry requirements. In addition, both companies prohibit running any application that tethers a wireless device to other devices such as computers, unless the subscriber has purchased that option. AT&T Wireless also specifically prohibits Web camera posts and broadcasts, automatic data feeds, automated machine-to-machine connections, peer-to-peer (P2P) file sharing, downloading movies using P2P file sharing services, redirecting television signals for viewing on personal computers, operation of servers, telemetry devices, supervisory control and data acquisition devices, and use as a substitute or backup for private lines, landlines or full-time or dedicated data connections. Verizon Wireless permits uploading, downloading and streaming of audio, video and games, and VoIP, whereas AT&T Wireless's terms and conditions prohibit all applications that are not specifically permitted, and thus would probably be interpreted as prohibiting these applications.

In addition, under the first scenario, ISPs are allowed to determine which applications may obtain access to QoS. Commonly, the ISPs own voice and video applications will rely on QoS, but voice and video applications provided by competitors or offered over the public Internet will be treated as best effort traffic.

The justification given by wireless ISPs for such application restrictions is usually capacity issues[4]. AT&T Wireless justifies a portion of its list of prohibited applications stating that they "cause extreme network capacity issues and interference with the network". Verizon Wireless similarly prohibits use of its data plans in a manner that "interferes with network's ability to fairly allocate capacity among users, or that otherwise degrades service quality for other users". Both AT&T Wireless and Verizon Wireless reserve the right to protect their wireless networks from "harm, compromised capacity or degradation in performance".

Such limits on applications are usually in addition to overall limits placed on data usage per month. AT&T Wireless currently[5] offers plans for smartphones without tethering with limits of 200MB or 2GB per month, and a plan for smartphones with tethering with a limit of 4GB per month. Verizon Wireless currently[6] offers an unlimited data plan for smartphones without tethering, and a plan for smartphones with tethering with a limit of 2GB per month. All plans with limits include per MB or per GB charges for overages. The pricing of these plans is also device dependent. In addition, AT&T Wireless explains that some variations in price are "based on the transmit and receive capacity of each device".

This first scenario has severe drawbacks. The ISPs do not explain how they define capacity issues or interference, what constitutes fair allocation of capacity, or how they determine which applications are permitted or prohibited. The prohibited applications seem to have in common that they tend to have higher duty cycles than permitted applications, e.g. prohibited applications may transmit or receive at a lower rate but for longer periods of time than permitted applications. It is not apparent, however, that prohibited applications necessarily use more capacity, cause more congestion, or use more network resources than permitted applications. It is also not apparent what the relationship is between the two traffic management methods of application restriction and usage limits.

The language used by wireless ISPs in their terms and conditions does not provide much clarity. No definitions are given of what constitutes "excessive amounts of net traffic" or "extreme network capacity issues". Furthermore, references to degrading "service quality for other users" or disrupting "net user groups or email use by others" would seem to be unenforceable, since all applications compete with each other for resources on the air interface during congested times. Finally, AT&T Wireless's approach of prohibiting all applications which are not expressly permitted is too strict to be reasonable.

B. Scenario 2: Open air interface

A second scenario, often supported by those who support wireless net neutrality, consists of a mandate that air

[3]All examples of AT&T Wireless and Verizon Wireless terms and conditions are as of May 9, 2011.

[4]Wireless ISPs also prohibit use for illegal purposes, including spam, sending viruses, denial of service, and illegal security breaches. However, discussion of illegal purposes is outside the scope of this paper.

[5]As of May 9, 2011.

[6]As of May 9, 2011.

interfaces be both standardized and open. The air interface would probably consist of the physical and medium access control[7] layers, as currently represented in 3G by the CDMA and GSM standards. Openness is accomplished by requiring wireless ISPs to allow attachment of any device that utilizes a compatible standardized air interface. Compatibility would be regulated by a requirement to obtain either a wireless modem or a SIM card from the ISP.

This scenario mainly focuses on the "any device" principle, which is outside the scope of this paper. However, it is often envisioned by proponents of wireless net neutrality that allowing attachment of any device would also allow the subscriber to run any application. This translation from *any device* to *any application* is based on the presumption that the ISP would have no ability to restrict applications on a device that they do not sell.

This second scenario also has several problems. First, ISPs subsidize most devices that they sell, and although under this scenario they must allow subscribers to use other devices, they may choose not to subsidize devices that they do not sell. If ISPs can continue to restrict applications on devices they sell, openness of applications is limited. Second, the approach does not directly address the use of QoS. An ISP would retain the ability to use QoS as it wishes. As a result, an ISP would be allowed to determine which applications may obtain access to QoS, e.g. the ISPs own voice and video applications may rely on QoS, but voice and video applications provided by competitors or offered over the public Internet may be treated as best effort traffic.

C. Scenario 3: No unreasonable discrimination for Internet services

A third scenario, perhaps intended to be a compromise between proponents and opponents of wireless net neutrality, is given by the FCC Open Internet Order. The regulations distinguish between fixed and mobile wireless broadband Internet access services.

ISPs that offer fixed wireless or wireline broadband Internet access services are subject to a set of rules for Internet services. A *no blocking rule* prohibits them from preventing any user from running any lawful application, unless the blocking is for the purposes of "reasonable network management". A *no unreasonable discrimination rule* prohibits them from unreasonably discriminating amongst lawful applications, unless the discrimination is for the purposes of reasonable network management. A network management practice is reasonable "if it is appropriate and tailored to achieving a legitimate network management purpose, taking into account the particular network architecture and technology of the broadband Internet access service". Such purposes include network security, parental and spam filters controlled by the user, and congestion mitigation. Practices that allow end-user control, are transparent, are based on use-agnostic discrimination, and/or are recognized best practices are more likely to be accepted as reasonable. In contrast, a network management practice that prioritizes traffic based on payment from a third party is unlikely to be considered reasonable. The dependence on network architecture allows for differences between fixed wireless and wireline networks, but these will be handled in the future on a case-by-case basis.

In contrast, ISPs that offer mobile wireless broadband Internet access services are subject only to weaker rules for Internet services. The *no blocking rule* only prohibits them from preventing any user from access lawful websites and running applications that compete with the ISP's voice or video telephone services, unless the blocking is for the purposes of reasonable network management. This no blocking rule thus allows a mobile wireless ISP to block a range of other applications. The FCC did not apply any *no unreasonable discrimination rule* to mobile wireless broadband Internet access services.

None of these rules are applied to "specialized services", which are likely to include facilities-based VoIP and IP-video, even when they are offered over the same networks and same capacity as "Internet services".

As a consequence, the *no blocking rule* may allow fixed wireless ISPs to prohibit applications only when they can demonstrate that such prohibitions are reasonable practices to reduce congestion, but may allow mobile wireless ISPs to prohibit any application that doesn't compete with their own voice or video telephone services. The prohibitions by many current mobile wireless ISPs of particular applications will have to be judged under these regulations on a case-by-case basis. Fixed wireless ISP traffic management practices may be deemed to be unreasonable if they do not treat applications that generate similar traffic loads similarly, while mobile wireless ISP traffic management practices are likely to be granted considerably more latitude. However, we expect that AT&T Wireless's approach of prohibiting all applications which are not expressly permitted would not be deemed reasonable.

In addition, the *no unreasonable discrimination rule* may require fixed wireless ISPs (but not mobile wireless ISPs) to give access to QoS to all applications that require it, if QoS is offered in the wireless network for non-specialized services.

This third scenario thus has severe drawbacks of its own. It is likely that many wireless ISPs will respond to these regulations by offering QoS only for specialized services. As a consequence, the specialized services exception will allow a wireless ISP to offer QoS only to its own voice and video applications, and to block access to QoS to competing application provider's voice and video applications. This lack of access to QoS will limit the viability of applications that require QoS and that are not offered by the ISP itself.

D. Summary

Our two goals, as discussed above, are to provide a level playing field between applications offered by ISPs and those offered by competing application providers, and to allow wireless ISPs to reasonably manage wireless network resources. None of these three scenarios effectively accomplishes both goals. In particular, none of the three scenarios accomplishes the goal of providing a level playing field. The first scenario (allowing ISPs to restrict applications at their discretion) does not address the first goal at all, except through competition

[7]The medium access control layer is the bottom portion of the data link layer.

between ISPs. The second scenario (mandating standardized open air interfaces) attempts to address the first goal, but falls short because it focuses on devices rather than applications. The third scenario (prohibiting unreasonable discrimination for non-specialized services) also attempts to address the first goal, but falls short because it allows the key applications to be placed under specialized services where they can avoid competition. The third scenario also may unreasonably limit a wireless ISP's ability to manage wireless network resources for non-specialized services.

Since all three scenarios fall short of accomplishing these goals, we will propose an alternative approach.

V. A Proposal for an Open Layer Interface

In this section, we propose a new policy that limits a wireless ISP's ability to restrict applications while guaranteeing it the ability to implement reasonable traffic management. In the first subsection, we propose a layer-based delineation between services that can only be provided by ISPs and services that can be provided by either ISPs or by application providers. In the second subsection, we use this delineation to propose a wireless net neutrality policy by requiring an open layer interface between these two sets of services. This proposal is presented in the context of communications law in the United States and may or may not apply to the regulatory approaches used in other countries.

A. Defining the relevant layer interface

Much of the difficulty with formulating communications policy that effectively addresses these goals stems from the FCC's decision to classify Internet access solely as an information service. To any Internet researcher, clearly the Internet consists of both telecommunication services (the lower layers of the network) and information services (the upper layers). We thus look to the academic literature on communications policy for a better way. A number of papers have suggested using network layers as a tool to formulate communications policy. Lessig [40] considers a model consisting of physical, logical, applications, and content layers. He argues that the physical and logical layers of the Internet have historically been neutral. He believes it is acceptable for the physical layer to be closed, but proposes that the logical layer should be open and act as a commons. Werbach [41] uses a similar set of four layers: the physical layer corresponds to OSI layer 1, the logical layer to OSI layers 2-6, and the application and content layers share OSI layer 7. He argues that communications policy should be formulated around these layers, with open interfaces between them. Solum and Chung [42] propose a six layer model, and argue that communications policy should attempt to respect the integrity of layers and to place regulation at or near the layer where the problem occurs. Whitt [43] suggests a four layer model similar to Werbach's, and presents principles concerning how layers should inform policy formulation.

We see great power in a layered approach to communications policy. However, unlike these papers, we believe that the key distinction should be between *network layer and below* versus *transport layer and above*. The purpose of such a distinction is to support our goals of (1) providing a level playing field between applications offered by ISPs and those offered by competing application providers and (2) allowing wireless ISPs to reasonably manage wireless network resources. With respect to the first goal, Internet layering and the end-to-end principle suggest that a network function should be implemented at or below the network layer *only if* it cannot be implemented effectively above the network layer. (See [26] for an extensive argument that this broad interpretation of the end-to-end principle maximizes innovation and social welfare.) With respect to the second goal, in section II, we found that although wireless networks often require stronger forms of traffic management than wired networks, these differences occur at the network layer and below. In contrast to our focus on network layer and below versus transport layer and above, most of the other layered models discussed above merge the Internet network and transport layers into a single logical layer, and therefore fail to distinguish between network layer functions provided in the access network and transport layer functions provided only at endpoints.

The delineation between network layer and below versus transport layer and above also can be used to distinguish between methods of discrimination that pose competitive concerns because they must be offered by ISPs (e.g. QoS) and those that do not because they can be offered by an ISP or by an application provider.

Although this layered approach has great power in a policy framework, we believe the key distinction in policy should really be between functionality that must be provided *within the access network by the ISP* and functionality that can be provided *elsewhere in the Internet*. A purely layer-based delineation is not sufficient to provide this distinction. We propose the following definitions be used in regulations to legally define the scope:

(1) INTERNET INFRASTRUCTURE SERVICES.-The term 'Internet infrastructure services' means all services- (A) over a network that uses a public right-of-way; and (B) that reside at or below the network layer or are required to manage the network.

(2) INTERNET APPLICATION SERVICES.-The term 'Internet application services' means all services-(A) over a network that uses a public right-of-way; (B) that are not infrastructure services; and (C) that do not fall under title VI of the Communications Act.

The definition of Internet infrastructure services starts with services at and below the network layer, and adds higher layer services required for an ISP to manage its network (e.g. DHCP). The definition of Internet application services starts with services above the network layer, subtracts services required for an ISP to manage its network, and also subtracts cable video services (which are adequately addressed by Title VI). Both terms only apply to networks that use a public right-of-way; those that do not are private networks and are not subject to regulation.

These terms agree with the intent of the older terms telecommunication services and information services. Internet infrastructure services do not change the content of information, similar to telecommunication services. In contrast,

Internet application services create, store, or change the presentation of information, similar to information services.

Internet infrastructure services can only be provided by carriers, and must be provided by each carrier within their autonomous system. They require large investments into loops or wireless spectrum and switches or routers. These large initial fixed costs of the business are high relative to the costs per incremental customer served, and thus Internet infrastructure services have a natural economy of scale that serves as a high barrier-to-entry. This high barrier-to-entry leads few carriers to offer service in any particular geographical region.

In contrast, Internet application services can be provided by carriers or by many other application providers on the Internet and can be placed at many locations within the Internet. Such services include email, web-hosting, caching, voicemail, and the portions of VoIP and IPtv that can be offered by independent application providers. Internet application services usually have fixed costs that are small relative to incremental costs, and thus there is usually a low barrier-to-entry, which leads to a competitive market with a large number of application providers.

Some protocols that have usually been implemented in the access network (e.g. DNS, pop, and smtp) are thus classified as Internet application services. We believe this is appropriate, since these functions could be offered by an application provider other than the access carrier.

B. Requiring an open layer interface

Since the differences between wired and wireless networks lie in lower layers, net neutrality in both wired and wireless networks can be simply and effectively accomplished by requiring an *open interface* between Internet infrastructure services and Internet application services. We propose that ISPs be prohibited from using Internet infrastructure to produce an uneven playing field in Internet applications. We also propose that ISPs should be guaranteed the right to use desirable forms of network management.

To accomplish this, we propose the following regulations:

A) Any QoS mechanisms that an ISP implements in Internet infrastructure services must be made available to application providers and peering ISPs without unreasonable discrimination.

B) ISPs are prohibited from refusing to provide enabling Internet infrastructure services to competing application providers in order to differentiate the ISP's own application offerings.

C) ISPs are prohibited from providing Internet infrastructure services to competing application providers at inflated prices in order to favor the ISP's own application offerings.

D) ISPs are prohibited from making exclusive deals to provide enabling Internet infrastructure services to certain application providers.

E) ISPs have the right to apply reasonable traffic management mechanisms.

F) ISPs have the right to make reasonable arrangements with consumers, application providers, and peering ISPs for Internet infrastructure services in a manner that does not conflict with the above regulations.

G) Forbearance of these regulations should be applied where sufficient competition exists.

Together, these seven regulations accomplish our two goals. Regulations A through D accomplish the goal of ensuring a level playing field. Regulations E through G delineate reasonable ISP use of traffic management. Details on how to implement each of these regulations are given in [9].

We now illustrate how ISPs may deploy QoS within such a regulatory framework. The goal of an open layer interface should be to allow application access on a level playing field to lower layer QoS. The interface also represents a contract between the wireless carrier and the user. The wireless carrier offers certain services (perhaps at a specified price) and the user requests (or purchases) those services desired at the volume desired. The central idea is that the wireless carrier will use traffic management techniques to ensure that services are rendered and hence contracts are satisfied.

Since the services are offered at an interface between the network and transport layers, they cannot be based on the application. In contrast, many current plans are not application-agnostic and are hence not consistent with an open layer interface. Some plans for smartphones include unlimited amounts of data, but restrict use to certain devices (e.g. prohibit tethering to a laptop) and to certain applications (e.g. permit web browsing and email, but prohibit file sharing, streaming, and VoIP). The goals of traffic management can be more efficiently accomplished by using an application-agnostic interface that allows users to choose their own applications and to match these applications to QoS options based on price.

That said, we do not expect that the vast majority of users will want to directly match applications to QoS themselves. Instead, we expect that most QoS options offered by ISPs will be designed to support particular applications efficiently, e.g. high throughput service for file sharing, traditional best effort service for email and web browsing, bounded delay service for streaming, and guaranteed low delay service for VoIP and video conferencing. This will allow a user to determine which QoS package to purchase based on a user's network use. We expect that common applications will have a default QoS selection appropriate for most users. However, this default could be overwritten by savvy users, e.g. if a user wants traditional best effort service for file sharing, she can elect this option.

For guidance over likely pricing strategies, we look to current pricing plans for mobile wireless smartphones. Part of the determination of price relates to usage of resources in lower layers. In wireless networks the key transmission resources are bandwidth and transmission power. The bandwidth and transmission power allocated to a user or a flow determines performance at lower layers, usually described as in terms of throughput, delay, and packet loss. Various lower layer traffic management practices control the bandwidth and transmission power allocated to each user or flow in order to create the desired variety of QoS service offerings. The price for each QoS service offering thus is likely to depend in part on the resources used to achieve that QoS. However, another part of the determination of price relates to demand for the service; the high price per unit volume for text messaging is due to this factor. Pricing plans are simplified representations of these

factors. Two forms are currently common. In the quota form, a user can purchase access to a specified amount of a service at a specified price, e.g. 5GB/month of best effort service for $60/month. In the volume form, a user can purchase access on a per volume basis, e.g. $50/GB. The two form are often combined, e.g. 5GB/month for $60/month with excess volume charged at $50/GB[8].

We would expect that pricing for QoS could be implemented using an open layer interface in a similar form. A wireless ISP may choose to offer add-on packages to a basic Internet access service. For instance, it may offer a guaranteed low delay service (designed for VoIP) at a price of $.01/min/(8kbps), where 8kbps is chosen to accommodate one VoIP stream. Alternatively, it may choose to offer a high throughput service, transmitted at times low congestion (designed for file sharing) at a price of $10/month for 5GB/month with excess volume charged at $5/GB. In neither case, however, could the ISP charge based on the application itself. Alternatively, it may offer these services combined with existing tiers, e.g. the guaranteed low delay service may be available to consumers who subscribe to a higher broadband Internet tier. Finally, our policy allows for reasonable application provider payment for QoS on behalf of a user. For example, a VoIP application provider other than the user's ISP may wish to purchase guaranteed low delay service from the user's ISP on behalf of the user, and then bundle this low delay service with the application provider's VoIP service sold to the user.

Finally, it should be noted that placing the contract at the interface between the network and transport layer in no way unreasonably restricts an ISP's ability to charge additional amounts for services that are not part of OSI layers 1-3. An ISP may continue to subsidize the purchase price of a wireless device and to recoup that subsidy over time. An ISP may also offer high layer services at additional cost, e.g. voice mail, voice dialing, navigation, ring tones, and roadside assistance.

VI. A Comparison of our Proposed Policy to the Three Scenarios

In this section, we compare our proposed policy to the three scenarios given above in section IV.

A. Revisiting scenario 1: No restrictions on traffic management

In the first scenario, ISPs are allowed to restrict applications at their discretion. This scenario is commonly used to restrict the applications that may be used by subscribers of cellular data service, e.g. to prohibit all applications other than Internet browsing, email, and intranet access. We reject this approach as unreasonably restrictive. ISPs may more directly limit the traffic a subscriber transmits and/or receives by placing limits or charges on the volume and/or rate of transmission. Such limits would be allowed under our proposed policy. Furthermore, they would remove the need for the imprecise language currently used in subscriber terms and conditions to limit user traffic, e.g. "excessive amounts of net traffic" and "extreme network capacity issues". By placing limits directly on traffic rather than on applications, public policy would remove impediments to application development.

In addition, under the first scenario, ISPs are allowed to determine which applications may obtain access to QoS. Commonly, the ISP's own voice and video applications will rely on QoS, but voice and video applications provided by competitors or offered over the public Internet will be treated as best effort traffic. In contrast, our proposed policy would mandate that any ISP using QoS for its own applications must offer it without unreasonable discrimination to application providers and peering ISPs. This approach will encourage the offering of end-to-end QoS, and thus also remove impediments to development of applications that rely on QoS.

B. Revisiting scenario 2: Open air interface

In the second scenario, there would be a mandate that ISPs support a standardized open air interface. The air interface would probably consist of the physical and medium access control layers, as currently represented in 3G by the CDMA and GSM standards. Openness is accomplished by requiring wireless ISPs to allow attachment of any device that utilizes a compatible standardized air interface. Compatibility would be regulated by a requirement to obtain either a wireless modem or a SIM card from the ISP. The assumption here is that allowing attachment of any device would also allow the subscriber to run any application. However, we envision that ISPs would continue to restrict applications on the devices that they subsidize. In contrast, our proposed policy does not allow such application restriction on any device regardless of whether it is provided by the ISP. In addition, under the second scenario, a wireless ISP would retain the ability to use QoS as it wishes, whereas our proposed approach would prohibit unreasonable discrimination using QoS.

We also believe that the requirement of an open layer interface between Internet infrastructure services and Internet application services is less intrusive than the requirement of a standard air interface for wireless devices. An open layer interface between the networking and transport layer only requires standardization of the *Application Programming Interface* (API) for services offered by the networking and lower layers, while a standardized air interface requires standardization of all physical and data link layer protocols and of their use.

C. Revisiting scenario 3: No unreasonable discrimination for Internet services

In the third scenario, fixed wireless ISPs are prohibited from unreasonably discriminating amongst applications except for specialized services, where specialized services include VoIP and IP-video services provided over the same networks used for broadband Internet access. This rule for fixed wireless ISPs is more restrictive than our proposed policy for non-specialized services. It may be interpreted as requiring fixed wireless ISPs to provide QoS without charge to all applications that require it, if QoS is offered in the wireless network for non-specialized services. In contrast, our proposed policy would allow an ISP to sell QoS at reasonable prices. In

[8]These are the charges for one of the AT&T Wireless Data Connect plans for full Internet access from a laptop, as of May 9, 2011.

contrast, there is no such mandate for mobile wireless ISPs, and thus our proposed policy is certainly more restrictive than the FCC regulations for mobile wireless ISPs.

In addition, under the third scenario, it is likely that many wireless ISPs will respond to these regulations by offering QoS only for specialized services. As a consequence, the specialized services exception would allow a wireless ISP to offer QoS only to its own voice and video applications, and to block access to QoS to competing application provider's voice and video applications. In contrast, our proposed policy does not create an artificial distinction between specialized and non-specialized services, and requires offering of QoS whenever an ISP uses it for its own applications.

Furthermore, with an open layer interface, there is no need to define *harmful interference*, since the network maintains the ability to control how users compete for resources with each other. An open layer interface does not impede an ISP's ability to reasonably accomplish this. There is also no need to define what constitutes the *Internet* portion of a provider's offerings, as the layering model applies both to Internet protocols and to telephone network protocols. All of a wireless provider's offerings are thus included without reference to the application or technology. There is also no need to define what constitutes *specialized services*, as the open layer interface requires access to lower layer QoS mechanisms that enables real-time applications, rather than carving out real-time applications as an exception. The open layer interface thus encourages competition in specialized services, rather than inhibiting such competition. Finally, requiring an open layer interface in both wired and wireless networks avoids the need to differentiate net neutrality policy on the basis of the technology used. However, it does allow for differences between wired and wireless networks on the basis of competitive markets, since forbearance may be granted in some markets but not in others.

D. Implications and compatibility

An open layer interface also would bear on the question of whether users have the right to run the software of their choice on wireless devices. Because an open layer interface allows any application provider to offer Internet application services without unreasonably impacting the network, there is no reasonable justification on the basis of traffic congestion for limiting applications on any device. The impact of an application would be controlled by the ISP at the interface through limits and/or charges placed on traffic. Applications that transmit high volumes of traffic would either purchase this capacity at standard rates or would consume a high proportion of a pre-purchased traffic quota. Applications that require QoS would similarly either purchase the required QoS at standard rates or would consume some portion of pre-purchased QoS. The interface is only concerned with traffic volume, QoS, and payment; it is not concerned with what particular application this traffic is destined to. From the point of view of the lower layers, therefore, the system is application-agnostic.

This interface-based approach is consistent with current wireless device and operating system architecture. There are a number of different operating systems that are used on wireless devices. However, they are all built using layered architectures. The operating system limits the ability of applications to access lower layer protocols. Indeed, an API provides an interface between applications and the operating system. The API defines what services the operating system offers to applications and how to access them. As in wired networks, common wireless network device operating system APIs offer access to lower layer functionalities at the networking layer or above; they do not offer direct access to protocols at the physical and data link layers. This architecture implies that direct access to all lower layer protocols is not necessary for net neutrality; all that is required is an *open layer interface*. Similarly, the architecture implies that lower layer protocols have no need to know what application a packet belongs to; they only need know what type of traffic management to apply.

In conclusion, we believe that our proposed policy provides a level playing field between applications offered by ISPs and those offered by competing application providers, and allows wireless ISPs to reasonably manage wireless network resources, whereas none of the three other approaches effectively accomplishes both goals.

VII. Conclusion

The principal question addressed in this paper is whether and to what extent wireless ISPs should be able to limit applications or their behavior. The analysis is based on both wireless network architecture, which dictates the form of traffic management methods, and communications law, which dictates what various communications providers are allowed to do. From communications law, we found that ISPs should be prohibited from giving themselves an unfair competitive edge by blocking applications that compete with their own or by denying access to or charging unfair prices for QoS to competing application providers. We examined three scenarios currently under consideration for how applications may or may not be restricted on wireless networks, and found that none accomplishes these goals. From network architecture, we found that wireless networks often require stronger traffic management than wired networks at and below the network layer. We used this finding to construct a set of regulations that limits an ISP's ability to restrict applications by requiring an open interface between network and transport layers. We illustrated how ISPs may deploy QoS within such a regulatory framework, and how this proposed policy can achieve the dual goals. We did not address here several important aspects – in particular, we did not consider potential differences in competitive markets (which may lead to forbearance of our proposed net neutrality rules), and we did not consider the related issue of device attachment.

References

[1] S. Jordan, "The application of net neutrality to wireless networks based on network architecture," *Policy Internet*, vol. 2, 2010.

[2] C. T. Marsden, *Net Neutrality: Towards a Co-Regulatory Solution*. Bloomsbury, 2010.

[3] M. A. Lemley and L. Lessig, "The end of end-to-end: preserving the architecture of the Internet in the broadband era," *UCLA Law Rev.*, vol. 48, pp. 925–972, 2001. Available: http://ssrn.com/abstract=259491.

[4] T. Wu, "Network neutrality, broadband discrimination," *J. Telecommun. High Technol. Law*, vol. 2, pp. 141–178, 2003. Available: http://ssrn.com/abstract=388863.

[5] C. S. Yoo, "Beyond network neutrality," *Harvard J. Law Technol.*, vol. 19, pp. 1–77, 2005. Available: http://ssrn.com/abstract=742404.
[6] P. Weiser, "The future of Internet regulation," *UC Davis Law Rev.*, vol. 43, pp. 529–590, 2009.
[7] D. D. Clark, "Network neutrality: words of power and 800-pound gorillas," *International J. Commun.*, vol. 1, pp. 701–708, 2007.
[8] J. M. Peha, "The benefits and risks of mandating network neutrality, and the quest for a balanced policy," *International J. Commun.*, vol. 1, pp. 644–668, 2007.
[9] S. Jordan, "A layered network approach to net neutrality," *International J. Commun.*, vol. 1, pp. 427–460, 2007.
[10] FCC, "FCC 09-93, Open Internet NPRM," Oct. 2009. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-09-93A1.pdf.
[11] ——, "FCC 10-201, Open Internet Order," Dec. 2010. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-10-201A1.pdf.
[12] ——, "FCC 07-31, Broadband Market Practices Notice of Inquiry," 2007. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-07-31A1.pdf.
[13] Comcast Corporation, "Comments of Comcast Corporation before the Federal Communications Commission in the matter of broadband industry practices (WC docket no. 07-52)," Feb 2008. Available: http://fjallfoss.fcc.gov/ecfs/document/view?id=6519840991.
[14] T. Wu, "Wireless Carterfone," *International J. Commun.*, vol. 1, pp. 389–426, 2007.
[15] FCC, "FCC 68-661, Carterfone Order," 1968.
[16] ——, "FCC 05-151, Internet Policy Statement," 2005. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-151A1.pdf.
[17] R. W. Hahn, R. E. Litan, and H. J. Singer, "The economics of 'wireless net neutrality'," *J. Competition Law Economics*, vol. 3, no. 3, pp. 399–451, 2007.
[18] R. W. Hahn and R. E. Litan, "The myth of network neutrality and what we should do about it," *International J. Commun.*, vol. 1, pp. 595-606, 2007.
[19] C. Marsden, "European law and regulation of mobile net neutrality," *European J. Law Technol.*, vol. 1, 2010.
[20] H. Zimmermann, "OSI reference model-the ISO model of architecture for open systems interconnection," *IEEE Trans. Commun.*, vol. 28, no. 4, pp. 425–432, Apr. 1980.
[21] A. S. Tanenbaum and D. J. Wetherall, *Computer Networks*. Prentice Hall, 2011.
[22] A. Leon-Garcia and I. Widjaja, *Communication Networks: Fundamental Concepts and Key Architectures*. McGraw-Hill, 2004.
[23] J. Zander, S.-L. Kim, and M. Almgren, *Radio Resource Management for Wireless Networks*. Artech House Publishers, 2001.
[24] J. F. Kurose and K. W. Ross, *Computer Networking: A Top-Down Approach*. Pearson Education, 2010.
[25] B. A. Cherry, "Misusing network neutrality to eliminate common carriage threatens free speech and the postal system," *Northern Kentucky Law Review*, vol. 33, pp. 483–511, 2006.
[26] B. van Schewick, *Internet Architecture and Innovation*, 2010.
[27] U.S. Congress, "S.652.PP, Telecomunications Act of 1996, 104th Congress, Public Law 104-104," 1996. Available: http://thomas.loc.gov/cgi-bin/bdquery/z?d104:s.00652:.
[28] ——, "S.66.PP, Cable Communications Policy Act of 1984, 98th Congress, Public Law 98-549," 1984. Available: http://thomas.loc.gov/cgi-bin/bdquery/z?d098:s.00066:.
[29] ——, "S.12.PP, Cable Television Consumer Protection Act of 1992, 102nd Congress, Public Law 102-385," 1992. Available: http://thomas.loc.gov/cgi-bin/bdquery/z?d102:s.00012:.
[30] FCC, "First Computer Inquiry, Final Decision, 28 FCC2d 267," 1971.
[31] ——, "Second Computer Inquiry, Final Decision, 77 FCC2d 384," 1980.
[32] ——, "FCC 02-77, Cable Modem Service Order," 2002. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-02-77A1.pdf.
[33] ——, "FCC 05-150, Internet over Wireline Facility Order," 2005. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/FCC-05-150A1.pdf.
[34] Center for Democracy & Technology, "Comments of the Center for Democracy & Technology before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[35] Free Press, "Comments of Free Press before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[36] Google Inc., "Comments of Google Inc. before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[37] AT&T Inc., "Comments of AT&T Inc. before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[38] Verizon and Verizon Wireless, "Comments of Verizon and Verizon Wireless before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[39] CTIA - The Wireless Association, "Comments of CTIA - The Wireless Association before the Federal Communications Commission in the matter of preserving the open Internet (GN docket no. 09-191) and in the matter of broadband industry practices (WC docket no. 07-52)," Jan. 2010.
[40] L. Lessig, "The Internet under siege," *Foreign Policy*, vol. 127, pp. 56–65, 2001.
[41] K. D. Werbach, "A layered model for Internet policy," *J. Telecommun. High Technol. Law*, vol. 1, pp. 37–67, 2002. Available: http://ssrn.com/abstract=648581.
[42] L. B. Solum and M. Chung, "The layers principle: Internet architecture and the law," *U San Diego Public Law Research Paper No. 55*, 2003. Available: http://ssrn.com/abstract=416263.
[43] R. S. Whitt, "A horizontal leap forward: formulating a new communications public policy framework based on the network layers model," *Federal Commun. Law J.*, vol. 56, pp. 587–672, 2004.

Scott Jordan (S'86-M'90) received the B.S./A.B., the M.S., and Ph.D. degrees from the University of California, Berkeley, in 1985, 1987, and 1990, respectively. From 1990 until 1999, he served as a faculty member at Northwestern University. Since 1999, he has served as a faculty member at the University of California, Irvine. During 2006, he served as an IEEE Congressional Fellow, working in the United States Senate on Internet and telecommunications policy issues. His research interests currently include net neutrality, pricing and differentiated services in the Internet, and resource allocation in wireless multimedia networks.