

Network Neutrality Is About Money, Not Packets

Henning Schulzrinne
Columbia University

The topic of network neutrality has occupied engineers, economists, and telecommunication lawyers since at least 2003. Network engineers often largely perceive this as yet another debate about quality-of-service, but this constitutes probably the least interesting part of the problem and has served more often to distract from other, more important, issues. Almost all network neutrality problems are rooted in economic incentives, and thus are likely to require approaches related to pricing and enhancing competition.

Network neutrality (NN) is probably the most widely known term in telecommunication policy, both within the engineering community and in the wider well-informed public. However, its most common definition, such as the one in the call for papers for this issue, of “principle by which all traffic on the Internet must be treated equally” is misleading and does not address key economic, engineering, and policy concerns. Indeed, the “all traffic is equal” has often been used to make advocates for the underlying, more general, ideas seem unreasonable. For example, proponents of laws and regulations addressing the underlying concerns are often accused of not wanting to support potentially societally valuable services like telemedicine or connected vehicles. Going back to the earliest writings on the subject¹ as well as the key provisions of the 2010 and 2015 FCC Open Internet rules, I propose a set of three principles that, in my view, address the likely harms that consumer broadband Internet access providers can cause to the broader Internet ecosystem, without constraining new applications or business models or creating a strawman of prohibiting quality-of-service differentiation altogether, in the small number of cases where this is helpful, feasible, and economically viable.

Even though the term NN conjures the unhelpful notion of equal packet treatment, other terms, such as the 2010 and 2015 FCC terminology of the “protecting the Open Internet” have not garnered widespread recognition and are similarly ill-defined. Thus, I will continue to use the term NN, but hope to make the definitions and concerns both more expansive and more precise. However, it might be better to think of NN as the rough equivalent of a fiduciary duty, i.e., acting in their subscribers’ best interest rather than (just) their own.

THE INTERNET ECOSYSTEM

I will use the terms ISP, broadband Internet Access Service [BIAS (This acronym is used in FCC rules, and the word choice is not completely accidental)], and (access) provider interchangeably. Along with all regulatory discussions, the focus is on Internet access for consumers, defined by the FCC in 2010² as: “A mass-market retail service by wire or radio that provides the capability to transmit data to and receive data from all or substantially all Internet endpoints, including any capabilities that are incidental to and enable the operation of the communications service, but excluding dial-up Internet access service.” In particular, this excludes Internet access provided to large businesses and offered incidentally to other nontelecommunication services, such as in coffee shops or hotels and by airlines, as well as so-called *specialized services* that use Internet protocols, but do not reach the Internet at large. For example, such specialized services include carrier-provided voice-over-IP (VoIP) or cable TV carried over IP within the BIAS access network. The “mass market” restriction in the definition also removes the services provided by an enterprise for its employees from consideration, so that, for example, blocking sites considered NSFW is not a violation of NN by any regulatory or other discussion that I am aware of.

For purposes of this paper, I divide the consumer Internet ecosystem into four broad categories, in line with most of the literature (this corresponds roughly to the 2015³ and 2010 Open Internet Report and Order², the latter, however, omits the backbone network), even where the terminology differs: *Access providers* offer “last-mile” access to fixed (residential) or mobile users. Content, application, service, and device providers, sometimes called “*edge providers*” offer the services that are carried on the access network or attach to it. With IoT and mobile devices, the hardware–software boundary is fluid, so it makes sense to treat them as one category. *Backbone providers* offer wide-area, including international, connectivity. Since many of the concerns are similar, I also consider the content distribution networks to be a part of that category, as both traditional long-haul ISPs and content distribution networks (CDNs) connect to access networks in similar locations and often aggregate content. However, CDNs would be unlikely to be considered telecommunications services so that their legal treatment may differ particularly in interconnection disputes. Finally, consumers purchase content, services, and devices, and may also generate content.

In vertically integrated companies, the same corporate entity may offer access, backbone and some subset of content and applications. As this paper will illustrate, vertical integration of the telecommunications and service or content makes providing an open Internet much harder, as the possible choices are reduced to prohibiting differentiation by quality or volume or imposing at least a limited form of price regulation, including zero pricing.

I start by describing the goals that often motivate the concerns about NN in “What are The Goals?,” summarize some of common arguments raised against NN in “Arguments Against Network Neutrality,” very briefly summarize FCC activities in this space in “FCC Open Internet Efforts and Reversal,” and then make the argument that discussions NN are about economic incentives, not traffic engineering problems (see “Network Neutrality Concerns are Usually About Money, not Packets or Queues”), including for quality-of-service (QoS) (see “QoS is Economically Hard, not Technically”). I generalize the NN discussion to three principles that transcend the changing concerns about blocking, “fast lanes,” interconnection and privacy, to name a few, in “Three Network Neutrality Principles”. Longer term, NN proponents, and carriers are faced with the more fundamental challenge of differentiating their service classes as base service levels reach above 25 Mb/s and become symmetric (see “Future Challenge: Customer Differentiation”).

WHAT ARE THE GOALS?

NN is often seen as some abstract good that needs to be protected, without actually specifying the more general societal goals that it should serve. Like most economic policies, the end goal should be maximizing consumer welfare in economic terms as well as foster civic participation, human expression, and personal growth as as noneconomic goals, among other. (Whether Internet access achieves these goals is beyond the scope of this discussion; removing NN as a goal seems unlikely to address any related concerns).

The economic argument for NN is based on the argument that consumer welfare is maximized by allowing for unfettered competition, in this case, primarily for content, applications, services, and devices, as well as for infrastructure services like backbone networks and CDNs. More competition increases value to the end user by fostering innovation and reducing prices, thus increasing the consumption of services and build-out of faster networks. NN proponents often refer to this as the “virtuous circle” argument. Allowing access service providers to either prevent new entrants or increase prices reduces demand for such services, and thus directly reduces consumer welfare and indirectly prevents the emergence of new services that have to rely on scale and universal reachability to be viable.

Some disagree with this emphasis on subscriber and edge-provider interests, making the point that carriers deserve profits at least as much as edge providers. However, possible market power allowing supracompetitive pricing for some edge providers does not appear to entitle access providers to the same.

NN does not directly address competition in telecommunications services, although Title II of the 1996 U.S. Telecommunications Act originally included strong provisions for allowing competing providers to access last-mile communication infrastructure and thus lowering the barriers to entry for additional providers. In a more indirect way, transparency (see “Transparency”) efforts offer a way for consumers to make more informed choices among competing last-mile carriers, thus potentially increasing competitive intensity.

The idea of a neutral platform or carrier goes back to the common-law notion of a common carrier⁴, with railroads, airlines, shipping companies, and classical (telephony) telecommunications carriers serving that role in other industries. In all of these cases, common carriers enable commerce and travel whose total value far exceeds the cost of transport, and where competition on some or all routes may be limited.

ARGUMENTS AGAINST NETWORK NEUTRALITY

The arguments against NN tend to fall into following six categories.

Legal: Some claim that BIAs are not telecommunication services and thus are beyond current legal frameworks, e.g., of the United States telecommunication laws. (In general, regulatory bodies such as the FCC can only promulgate regulations that fall within the scope of authority granted by laws.)

General cost of regulations: Because of general concerns about the cost of complying with regulations, not specific to telecommunications, it is argued that Internet access services should not be subject to sector-specific regulation.

Lack of demonstrated need: Opponents claim that there have been a few violations of NN since the advent of consumer broadband services in the 1990s, and thus NN is a heavy hammer in search of a nail.

Alternative approaches: Opponents claim that antitrust law, other consumer protection laws, or public disapproval can deal with any anticonsumer or anticompetitive behavior by carriers. Advocates see general rules rather than postfacto adjudication of behavior as necessary to give edge providers a reasonable certainty that they will not be blocked or otherwise harmed by the behavior of last-mile providers.

Investment impact: Either the regulatory uncertainty, the fear of rate regulation or the prohibition of certain business models would make telecommunication services less attractive for investment and thus reduce the capital investment, particularly in high-cost areas. As a related argument, some claim that the market position of over-the-top providers of content and services allows them to pay less than their fair share of the cost of deploying and operating networks, thus reducing the return on investment of network operators.

Impact on new applications: NN may prevent the emergence of new applications, particularly those that require better-than-best-effort service, e.g., for telemedicine or other real-time services.

The FCC report and order along with the filings in the 2017 *Restoring Internet Freedom* proceeding are a good source of detailed legal, technical, and economic arguments. This paper cannot address most of these arguments at length, particularly since they are often legal and economic, but we will take a closer look at the impact on new applications in the following sections.

FCC OPEN INTERNET EFFORTS AND REVERSAL

To illustrate how NN has been codified in practice, I briefly describe the four FCC NN efforts, spanning the years 2004–2017, without describing the various court cases that often motivated the actions.

In 2004, FCC Chairman Powell (https://apps.fcc.gov/edocs_public/attachmatch/DOC-243556A1.pdf), appointed by President G. W. Bush, enumerated four Internet freedoms: “Freedom to access content, . . . use applications, . . . to attach personal devices, . . . to obtain service plan information.”

Concerns about access networks stifling competition go back much further, well before the Internet, to the earliest days of the voice telephone network. For example, the 1913 Kingsbury Commitment (https://www.corp.att.com/history/milestone_1913.html) required AT&T’s long distance network to connect to competing local telephone companies. Device attachment has long been contentious, and was only resolved in the United States after the 1968 Carterphone court decision and the so-called Part 68 rules for terminal attachments in 1977.

In 2010, the FCC Open Internet rules² mandated transparency, no blocking, and no unreasonable discrimination. The latter rule read “A person engaged in the provision of fixed broadband Internet access service, . . . , shall not unreasonably discriminate in transmitting lawful network traffic over a consumer’s broadband Internet access service. Reasonable network management (RNM) shall not constitute unreasonable discrimination.” The transparency and no-blocking rules applied to fixed and mobile Internet access, while the nondiscrimination rule only applied to fixed networks.

In 2015, after a loss in federal court, the FCC revised³ the Open Internet rules and based them on Title II of the 1996 Telecommunications Act that had previously been applied mainly to TDM networks, i.e., categorized Internet access as a telecommunication service. It also extended the nondiscrimination rule to wireless Internet access, added interconnection as an area of observation and potential intervention, and added a general conduct rule. The transparency rules were strengthened to include more performance metrics and information about consumer pricing.

In 2017, the new the chairman appointed by the incoming administration reversed course and declared ISP to be information services not subject to FCC regulation⁵, except the 2010 transparency rules. As for the previous iterations, the new rules will be contested in federal court.

NETWORK NEUTRALITY CONCERNS ARE USUALLY ABOUT MONEY, NOT PACKETS OR QUEUES

In network engineering, NN is often seen as a technical problem. However, so far, disputes about details of traffic management have never arisen on purely technical grounds. Without fail, the conflict is about economic incentives. Indeed, since the earliest efforts of codifying NN, carriers were granted exceptions for RNM. While the RNM terminology sounds hopelessly vague to engineers, the general notion of applying to standard, widely used techniques has been sufficient and has not required paneling an FCC engineering review board to review carrier router configurations.

The economic incentives for carriers fall into three categories, with the first two overlapping:

Reducing competition for vertically integrated services: If a carrier offers a service competing with third-party providers, it has an incentive to make its own service more attractive than that of the competition. It can do that by artificially reducing the quality or reliability of the third-party service, imposing financial charges on those services or backbone providers they use, or by outright blocking.

Two-sided market: BIAS providers interact with consumers as well as Internet backbone providers and content distribution networks (CDNs); some of the CDNs are operated by content providers. Since those providers cannot offer and charge for their service except by using the local access network, the BIAS provider has the opportunity to extract some or all of the third-party surplus by charging backbone operators and CDNs for access to their network. If the backbone operator and CDNs have pricing power, they will likely pass on these costs to the provider of services. While harder to implement in practice, the BIAS

provider could also directly charge providers of services or content for access to its subscribers. This incentive applies whether the provider is vertically integrated or not and overlaps with the first incentive. Providers may argue that charging backbone or content providers allows them to offer lower prices to their residential subscribers.

Leveraging customer information: BIAS providers know a lot about their customers.⁶ For example, they can observe all their traffic flows, except those using Virtual Private Network (VPNs), and their Domain Name System (DNS) queries. (Recent interest in privacy-protecting DNS queries using HTTPS may change this, but since DNS queries are currently unencrypted, they can be easily tracked even if the query is answered by a third-party resolver). Since they know the identity and location of the subscriber, they can monetize this data, e.g., by operating advertising exchanges. Again, this incentive does not depend on vertical integration, but vertical integration makes the data more valuable. For example, an ISP can use the Internet browsing data gathered to tailor advertisements in its streaming video product.

In all these cases, access providers may be deterred to act on these incentives by the fear of losing customers, if there is sufficient competition, or of customers choosing to subscribe to lower-cost slower services or simply foregoing the service altogether if they find it insufficiently attractive. If the provider is vertically integrated, it might be willing to suffer some loss in BIAS revenue if that protects other sources of revenue. Typically, smaller ISPs are less likely to be able to leverage their position; not surprisingly, for example, interconnection disputes have largely occurred between the largest consumer ISPs on one side and backbone ISPs and CDNs on the other, even though smaller providers typically bear higher costs for middle-mile connectivity.

A number of high-profile cases provide examples illustrating these motivations:

Service blocking: In 2005, Madison River, a small DSL provider, blocked the SIP port to discourage competition with its own telephony service. Comcast interfered with peer-to-peer video by issuing TCP RST packets in 2007. (It is disputed whether this was motivated by reducing network congestion or video competition.)

Interconnection: In 2014, Netflix and Level3 complained that their interconnection with major ISPs such as Comcast and Verizon FiOS suffered from poor performance.

Volume-based discrimination: Zero rating exempts certain services or traffic types from counting toward the bandwidth cap or volume-based charges. In some cases, the content provider can pay for that privilege to encourage use of its service or content. AT&T, in particular, has created such a third-party-paid service, but the only users appear to be services offered by AT&T.

Privacy: Verizon was fined by the FCC for tracking users with so-called supercookies (<https://www.fcc.gov/document/fcc-settles-verizon-supercookie-probe>).

Generally, the economic incentives and thus NN concerns shift over time. For example, for the early years of home broadband services, Wu¹ documents a number of other practices such as prohibitions against serving content, using VPNs or connecting Wi-Fi routers that now seem quaint. Many cellular data packages restrict tethering, i.e., using the cellular data connection to connect laptops, by volume or speed, which seems closer to violating the notion of application and device neutrality (see “Future Challenge: Customer Differentiation”).

QOS IS ECONOMICALLY HARD, NOT TECHNICALLY

Before discussing the issue of QoS in the context of NN, it is worth pointing out that despite decades of effort, there has been little consumer-level deployment of QoS in wireline networks, even though most routers, both core and edge, implement a variety of QoS mechanisms. QoS is quite common, with simple classes, for LTE networks, with a separate low-latency QCI (QoS Class Identifier) for conversational voice offered by the carrier, i.e., a specialized service in an NN terminology.

However, no carrier seems to be offering end-to-end, multi-provider QoS guarantees even to businesses, and abolishing notions of NN seems unlikely to change this. Thus, QoS is likely best seen as a last-mile, local service. This also implies that for practical purposes, specialized services and QoS are currently largely synonymous. If, however, QoS-sensitive services use backbone services that do

not suffer from congestion, other QoS-enabled services become possible, even without end-to-end QoS mechanisms.

Unlike for other services that offer differentiated grades of service, such as airlines, QoS for wireline networks is best seen as an insurance product, i.e., most of the time, best effort and QoS-enhanced service will be indistinguishable in a well-run network that actually offers the speed advertised. Except for satellite, most US residential networks fall into that category, as the FCC Measuring Broadband America report⁷ shows. Thus, a clever user would only turn on QoS services in the rare instances when congestion causes high packet loss, thus implementing a version of Paris Metro Pricing.⁸ This will likely lead to unpredictable “surge” pricing, and be about as popular as similar mechanisms for ride sharing services. Given those difficulties and since users highly value predictability,⁹ service providers may want to offer a “QoS retainer” or business-class service that charges a higher monthly rate and offers priority to some subset, by volume or rate, of the user’s application traffic, but without the ability to switch services and charges on demand. Indeed, the use of service level agreements for some commercial Internet access services already effectively provide a QoS guarantee.

There are only a few practical ways to charge for intranetwork QoS guarantees, each with their own competitive and neutrality challenges:

Carrier-provided application: Since the access provider offers the application, it can easily assign the application to a separate traffic class. Specialized services fall into this category, but this clearly disadvantages competing services from other providers.

Negotiated with third party: A third-party service wanting QoS guarantees could negotiate terms and conditions with carriers. However, there are roughly 2000 service providers in the United States alone (<https://www.ustelecom.org/broadband-industry/broadband-industry-stats/providers>), making the transaction costs of individualized negotiation nearly insurmountable, unless the service provider restricts QoS guarantees to the largest ISPs. A total of 90% of households are served by 14 ISPs.

Paid by third party automatically: In theory, third parties could pay terminating providers to offer traffic priority or some version of a committed information rate, but this would likely require an end-to-end signaling mechanism or the cooperation of most backbone providers as intermediaries to avoid creating settlements among thousands of carriers.

Subscriber paid: Subscribers could, as described above, pay for QoS for whatever flows they desire, likely mediated through the application needing QoS and a common API. Since subscribers also maintain a business relationship with their carrier and can authenticate themselves, this is relatively straightforward.

Zero cost: As noted, an access provider could make guaranteed QoS available as a part of the subscription, whether for the whole headline download speed or restricted to a fraction of the bandwidth or limited by volume. Since many QoS-sensitive services such as interactive games, VoIP, and video conferencing operate at speeds well below 1–2 Mb/s, this may be sufficient for consumer applications.

The implied assumption of many “NN prevents QoS” arguments appears to be that only the first option is technically viable, which appears to be clearly unfounded. Any option that involves payment by the subscriber or provider of the service raises concerns about economic discrimination, particularly if the ISP offers a competing service. Thus, rates for such services, if above zero, may need to be monitored for anticompetitive behavior.

With these examples, three principles emerge, namely end-user choice, economic neutrality, and transparency.

THREE NETWORK NEUTRALITY PRINCIPLES

I believe that three principles can guide open Internet policies that are likely to weather changes in technology, business models and economic incentives: a foundation of user choice, supported by economic neutrality and virtual structural separation and amplified by transparency.

Principle 1: End-User Choice

In the original paper by Wu,¹ the principle of NN is defined as “absent evidence of harm to the local network or the interests of other users, broadband carriers should not discriminate in how they treat traffic on their broadband network on the basis of inter-network criteria.”

Rather than focusing on network behavior only, it may be more helpful to consider *end-user choice* as the principle for deciding whether a particular traffic management or other policy is reasonable if it cannot be readily justified as protecting the network against attack or abuse.

User choice can trivially address concerns about traffic filtering, e.g., for blocking content objectionable to the subscriber. User choice also addresses the objection that NN prohibits QoS differentiation. Among the options discussed in “QoS is Economically Hard, not Technically,” if end users can choose whether to apply QoS to any application of their choice, e.g., by five-tuple (protocol, network source and destination addresses and ports) or DiffServ code point, and pays a uniform rate, the new service becomes available to all content, application, and service providers. This moves beyond a simplistic all-traffic-the-same to a user-choice model, removing incentives for the carrier to favor their own or particular third-party services. For example, a standard service plan could include a limited amount of high-priority data to avoid per-byte metering.

Implementing such a mechanism in practice would not be trivial, as it would require standardized signaling mechanisms that allow the end-user application to indicate which traffic flows should receive better QoS. Needless to say, QoS signaling for end users has not gained traction, whether by explicit signaling^{10,11} or by DiffServ code points. DiffServ is implemented in most routers, but agreeing on uniform code points and a signaling mechanism from the end user to the network remains a standardization and implementation challenge that software-defined networks may finally be able to address, combined with a simple user-to-network API that conveys flowspecs.¹²

Thus, QoS differentiation is not hard, from a policy perspective, in a neutral, consumer-driven network. However, charging raises potential concerns that I discuss in the following, as they apply more broadly.

User choice also provides a basis for deciding what user data providers should be able to monetize, suggesting opt-in as the choice-enabling approach for protecting privacy.

Principle 2: Economic Neutrality and Virtual Structural Separation

Until a few years ago, the delivery of content and the content itself were largely separate businesses, even for service providers that offered both, such as cable companies. Internet service was not a substitute for cable TV, and they used separate frequency bands on the coax cable. Now, almost all large US home broadband providers own video distribution services, and, in many cases, content producers. Cord cutting and IP-based delivery of linear TV have started to fundamentally change the business model and increased incentives for broadband providers to use their broadband services to favor their video distribution, or, failing that, to extract additional revenue from third-party providers.

In markets where the provider of the infrastructure with significant market power also competes for end-user services, a traditional remedy is structural separation. For example, British Telecom has a separate unit, Openreach, that provides fiber and copper transport to BT and other ISPs that offer consumer Internet access services. Elsewhere, the local regulator sets terms and conditions for the incumbent facilities-based provider, so that competing providers can offer their services on equal footing. This model existed briefly in the United States after the 1996 Telecommunication Act, but has now essentially disappeared.

This principle of competitively neutral charging is most challenging to apply for metered service, i.e., charging customers per gigabyte. Unlike for other utilities such as gas, electricity, and water, where there is a relatively high volume-dependent cost, the variable cost of delivering bits on the last mile is a tiny fraction of the typical residential wireline broadband bill. If providers charge per gigabyte for Internet access, but do not charge for the delivery of their own linear TV (“cable TV”) services, whether delivered via the Internet, a separate IP network, or even traditional digital cable, they enjoy

an inherent competitive advantage over over-the-top (OTT) video providers. This inherent advantage explains why zero-rating, i.e., the exclusion of certain content from metering or capping, has become one of the key controversies within the NN debate, and is probably the single largest disagreement, other than concerns about regulatory authority, in attempts to legislate NN in the United States.

In practice, as soon as payment is involved, vertically integrated providers still need to be prohibited from exercising a version of margin squeeze.¹³ For example, a provider could charge a supra-competitive price for QoS-enabled traffic or per gigabyte of zero-rated traffic, and offer the user a refund, i.e., the application operated by the service provider pays back the user. Since the QoS or volume charge and application fee are paid to the same entity, the QoS surcharge is only bounded by the cost of the application or content itself. Thus, for vertically integrated providers, the NN-respecting choice is between a technical prohibition of service differentiation and some form of price regulation, whether ex-ante or based on complaints. Naturally, this only applies to the part of the application space that is vertically integrated. For a BIAS provider that also offers entertainment video services, for example, the charge for telemedicine applications need not be regulated.

The FCC evaluated a variety of zero-rating plans in 2016¹⁴ and found least concern with zero-dollar versions that exempted all similar content from volume-based charges.

Zero rating may diminish in salience, however, as most United States smartphone plans have moved to versions of nominally “unlimited” plans that throttle or deprioritize rather than charge for data volumes exceeding a cap. It is less clear how data caps and zero rating will evolve for wireline services. Currently, a number of providers are implementing soft caps at around 1 TB (<https://arstechnica.com/information-technology/2017/08/at-least-196-internet-providers-in-the-us-have-data-caps/>). For reference, the median usage at Comcast was 151 GB per month in June 2018 (<https://www.xfinity.com/support/articles/data-usage-average-network-usage>); a user who substitutes HD OTT viewing for all traditional TV viewing, at around 4.5 h per day in the United States,¹⁵ would consume between 150 and 450 GB per month, but ultra HD could push this to well above 1 TB per unique viewer. This does not account for multiple viewers within a household watching different programs at the same time.

Unfortunately, there does not appear to be any empirical economic research on any impact the principles of economic neutrality had or would have.

Principle 3: Transparency

While transparency does not figure in the original Wu paper nor is typically discussed in the NN engineering or economics literature, it plays an important role in making NN both enforceable and tractable. Since most subscribers are not technical experts, disclosing network management practices makes it more likely that technical experts, whether at consumer advocates or edge providers, can review such practices and raise concerns with the provider or, if necessary, with a regulator. Generally, in most countries, companies can get into legal trouble for what, in the United States are called “unfair and deceptive trade practices,” if they lie about their actions even if the action itself would be legal if disclosed.

Disclosure of network performance, e.g., through the FCC Measuring Broadband America program (<https://www.fcc.gov/general/measuring-broadband-america>), is also seen as a means of increasing competition, by providing consumers with information to pick the best-performing ISP or to switch to a competitor. Generally, only wireline performance is being measured by regulators in a systematic way, as the geographic and temporal variability of mobile data speeds is still much higher.

The 2017 “Restoring Internet Freedom” order,⁵ maintained transparency, in its more modest 2010 form, as a requirement; the order claims that it will prevent ISPs from violating any voluntary no-blocking or nondiscrimination commitments.

FUTURE CHALLENGE: CUSTOMER DIFFERENTIATION

As long as telecommunication networks have been around, carriers have tried to separate their customer base into consumer and business users, often using relatively small differences to charge very different rates for similar services. In that approach, they resemble airlines and, in earlier eras,

railroads, where, say, upholstery and availability of food are used to separate customers by willingness to pay. For many years, telephone companies offered business lines that were essentially the same as residential lines, with the carriers trying to ensure that businesses would not be able to purchase residential service.

In broadband networks, speed, reliability, and some business-related features are used to segregate different classes of users. For example, residential services were generally slower, asymmetric in speed, and did not offer a static IP address. With symmetric IPv6 services offering 100 Mb/s or above speeds with static network addresses, many of these natural differentiators have become less relevant, so businesses might be tempted to purchase two consumer-grade services and get a more reliable service than buying one, significantly more expensive, business-class service. For residential service, there is limited willingness to pay for service above 25 Mb/s, although initial experiences hint at gigabit speeds as a status symbol. Thus, unless extremely high-bandwidth applications emerge, the old model of offering a basic service suitable for light users of email and text web pages and a “family-size,” video-capable service is no longer feasible.

Generally, public interest organizations and providers of OTT services have been opposed to metering bandwidth, for the reasons noted previously, but as download speed recedes as a natural differentiator, offering different products becomes harder. Indeed, we may converge on the pricing model offered by Google Fiber and some rural electric cooperatives, with only two tiers, at 100 Mb/s and 1 Gb/s. If the market is not fully competitive and thus the service offered at prices above the marginal cost, providers will then be incentivized to offer discounts to low-income consumers.

Except for consumers with low demand or financial constraints, substituting cellular mobile service for fixed residential service will likely take longer, maintaining the service differentiation

While recently, as 4G/LTE connections approach or exceed the speed of some wireline residential connections, cellular providers have combined “unlimited” data plans with restrictions on tethering, making it difficult or impossible to substitute a cellular data plan for a wireline Internet connection if one wants to use devices other than a smartphone. Since the average usage of cellular data, not counting Wi-Fi, around 3 GB/month (<https://www.fiercewireless.com/wireless/how-much-cellular-and-wi-fi-data-are-smartphone-users-consuming-and-which-apps-verizon-0>) in late 2016, is much less than HFC usage, at around 151 GB per month in late 2017 (i.e., the average is likely significantly higher than that), cellular providers are unlikely to be able to support wireline-sized monthly usage with current networks.

Thus, proponents of an open Internet and the principles discussed previously are challenged to consider customer differentiation. Is restricting tethering for “unlimited” plans a violation of NN by application? Will the notion of separate plans for business and residential use disappear for connections at or below 1 Gb/s? Can cellular (5G?) services provide competitive substitutes for wireline services?

CONCLUSION

NN is best seen not as rules for packet treatment, but as constraints on the economic behavior of access providers. Thus, it is quite possible to offer broadband services with differentiated service classes that satisfy the goals of nondiscrimination on economic grounds. Similarly, zero-rating and interconnection can be offered on reasonable and nondiscriminatory terms, without significantly impacting the ability of ISPs to operate a sustainable business. However, data-based analyses of NN policies or their absence are lacking.

REFERENCES

- [1] T. Wu, “Network neutrality, broadband discrimination,” *J. Telecommun. High Technol. Law*, vol. 22, p. 141, 2003. [Online]. Available: <http://dx.doi.org/10.2139/ssrn.388863>
- [2] Federal Communications Commission, “Preserving the open Internet,” Report and Order, Washington, DC, USA, Dec. 2010. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/FCC-10-201A1_Rcd.pdf

- [3] Federal Communications Commission, “Protecting and promoting the open Internet,” Report and Order, Washington, DC, USA, Mar. 2015. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/FCC-15-24A1.pdf
- [4] Robert Cannon, “Common carriers,” Dec. 2017. [Online]. Available: http://www.cyberte telecom.org/notes/common_carrier.htm
- [5] Federal Communications Commission, “Restoring Internet freedom,” Declaratory ruling, report and order, and order, Washington, DC, USA, Jan. 2018. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/DOC-347927A1.pdf
- [6] Federal Communications Commission, “Protecting the privacy of customers of broadband and other telecommunications services,” Report and order, Washington, DC, USA, Nov. 2016. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/FCC-16-148A1.pdf
- [7] Federal Communications Commission, “2016 Measuring Broadband America fixed broadband report,” Washington, DC, USA, Dec. 2016. [Online]. Available: <http://data.fcc.gov/download/measuring-broadband-america/2016/2016-Fixe d-Measuring-Broadband-America-Report.pdf>
- [8] A. Odlyzko, “Paris metro pricing for the Internet,” in *Proc. 1st ACM Conf. Electron. Commerce*, Denver, CO, USA, 1999, pp. 140–147. [Online]. Available: <http://doi.acm.org/10.1145/336992.337030>
- [9] A. Odlyzko, “Internet pricing and the history of communications,” *Comput. Netw.*, vol. 36, no. 5, pp. 1389–1286, 2001. [Online]. Available: [https://doi.org/10.1016/S1389-1286\(01\)00188-8](https://doi.org/10.1016/S1389-1286(01)00188-8)
- [10] “Resource ReSerVation protocol (RSVP)—Version 1 functional specification,” Internet Engineering Task Force, RFC 2205, Sep. 1997. [Online]. Available: <http://doi.org/10.17487/RFC2205>
- [11] R. Hancock, G. Karagiannis, J. Loughney, and S. V. den Bosch, “Next steps in signaling (NSIS): Framework,” Internet Engineering Task Force, RFC 4080, Jun. 2005. [Online]. Available: <http://doi.org/10.17487/RFC4080>
- [12] J. Wroclawski, “The use of RSVP with IETF integrated services,” Internet Engineering Task Force, RFC 2210, Sep. 1997. [Online]. Available: <http://www.rfc-editor.org/rfc/rfc2210.txt>
- [13] G. Gaudin and D. Mantzari, “Margin squeeze: An above-cost predatory pricing approach,” *J. Competition Law Economics*, vol. 12, no. 1, pp. 151–179, 2016. [Online]. Available: <http://dx.doi.org/10.1093/joclec/nhv042>
- [14] Federal Communications Commission, “Wireless Telecommunications Bureau Report: Policy review of mobile broadband operators’ sponsored data offerings for zero-rated content and services,” Jan. 2017. [Online]. Available: https://apps.fcc.gov/edocs_public/attachmatch/DOC-342987A1.pdf
- [15] The Nielsen Company, “The Nielsen total audience report: Q2 2017,” New York, NY, USA, Nov. 2017. [Online]. Available: <http://www.nielsen.com/us/en/insights/reports/2017/the-nielsen-total-audience-q2-2017.html>

ABOUT THE AUTHORS

Prof. Henning Schulzrinne is a Levi Professor of Computer Science at Columbia University. He received the Ph.D. degree from the University of Massachusetts in Amherst, Massachusetts. He was a member of the technical staff at AT&T Bell Laboratories and Associate Department Head at GMD-Fokus (Berlin), before joining the Computer Science and EE departments at Columbia University. He served as chair of Computer Science from 2004 to 2009 and as Engineering Fellow, Technical Advisor and Chief Technology Officer of the Federal Communications Commission (FCC) from 2010 until 2017. Protocol standards codeveloped by him, including RTP, RTSP and SIP, are now used by almost all Internet telephony and multimedia applications. He is a Fellow of the ACM and IEEE.