

The State of Network Neutrality Regulation

Volker Stocker
Weizenbaum Institute for the
Networked Society/TU Berlin
vstocker@inet.tu-berlin.de

Georgios Smaragdakis
TU Berlin
Max Planck Institute for Informatics

William Lehr
MIT
wlehr@mit.edu

ABSTRACT

The Network Neutrality (NN) debate refers to the battle over the design of a regulatory framework for preserving the Internet as a public network and open innovation platform. Fueled by concerns that broadband access service providers might abuse network management to discriminate against third party providers (e.g., content or application providers), policymakers have struggled with designing rules that would protect the Internet from unreasonable network management practices. In this article, we provide an overview of the history of the debate in the U.S. and the EU and highlight the challenges that will confront network engineers designing and operating networks as the debate continues to evolve.

CCS CONCEPTS

• **General and reference** → **Surveys and overviews**; • **Networks** → **Public Internet**; • **Social and professional topics** → **Governmental regulations**;

KEYWORDS

Internet Regulation; Internet Policy; Network Neutrality.

1 INTRODUCTION

Since its invention in the 1960s as a government-sponsored research network in the U.S., the Internet and its associated ecosystem have evolved into a key component of our global communications and computing infrastructure (e.g., [1, 52]). Today, it comprises a complex amalgam of interconnected and complementary computing and communication network resources operating at multiple, partially overlapping layers. Many of these are independently owned and managed, but collectively they support and sustain the user experience we associate with the Internet.

The Internet is widely recognized by policymakers around the world as key infrastructure to support global digital transformation processes as ever more economic and social activity moves on-line and the range and numbers of Internet-connected devices grows. Whereas the legacy telecommunications networks that comprised the Public Switched Telephone Network (PSTN) have long been recognized as comprising essential basic infrastructure and been subject to significant government regulation (and in many countries, were provided as government-owned utilities), the Internet has been mostly unregulated. The Internet emerged as an overlay on legacy telecommunications networks; Internet Service Providers (ISPs) were free to design their service offerings, negotiate interconnection agreements and manage their networks mostly free from significant regulatory constraints.

Over time, the Internet grew in importance and last-mile access providers of telephone and cable networks began to gradually transform their networks into general-purpose Internet Protocol

(IP)-based broadband platforms capable of delivering a diverse array of services, including broadband Internet access. In light of these developments, it was believed by many that a new regulatory framework would be needed to protect the Internet as an open, public platform for access to online content and services, and a platform for innovation. At the same time that access providers were transforming into broadband platform providers, a cohort of new market players such as content and cloud providers [66], and content delivery networks (CDNs) [70] have emerged as key players in the Internet ecosystem. Much of the Internet experience now depends on computing and networking resources contributed by firms like Google, Amazon, Facebook, Netflix, or Akamai. Historically, these newer players have been largely exempt from communications regulatory oversight, while legacy telephone and cable last-mile access providers remain subject to PSTN-era regulations that are poorly matched to current technologies and market conditions.

The effort to craft a new regulatory framework to protect and sustain the Internet as a public network and innovation platform has come to be known as the Network Neutrality (NN) debate. The NN debate is fundamentally a debate about what regulatory rules are needed to protect Internet users and content and application providers from unreasonably discriminatory network management practices by broadband access service providers. In the following, we provide an overview of the history of the debate in the U.S. and the EU and highlight the challenges that will confront network engineers designing and operating networks as the debate continues to evolve.¹ Figuring out who NN rules should apply to and what that may mean for how traffic should be managed is a work-in-process. Specifically, we examine how NN rules can be “bypassed” and explain the challenges participants in the Internet ecosystem may have in seeking to comply with NN rules that are likely to remain ambiguous and evolving. Our goal here is to help educate computer scientists and network engineers who may be unfamiliar with the regulatory history but who may increasingly need to consider NN regulations in the design, operation, and management of networked communication and computing resources that are part of the increasingly complex and socially and economically important Internet ecosystem.

2 THE ROLE OF LEGACY REGULATION FOR THE TELECOMMUNICATIONS SECTOR

There are certain basic infrastructures that are used by virtually every segment of society and the economy. Ensuring universal access to such basic, often critical infrastructure is recognized as a core function for government. That includes such things as access to clean water, reliable electric power, transportation grids, and basic

¹In the meantime, numerous countries worldwide have adopted different versions of network neutrality regulations [18].

telecommunications services [35, 86]. In many countries, the national telecommunication networks (like water and electric power grids) were historically provided via government-owned public utilities. In some countries, like the U.S., telecommunications networks were provided by investor-owned, but still heavily regulated public utilities (e.g., the Bell System).

Government provisioning and heavy regulation of basic infrastructure services has often been justified by the assessment that those infrastructures constituted natural monopolies—or equivalently, that it would cost more and be economically inefficient to have the infrastructure services provided by multiple, competing providers. However, government provisioning brings its own problems, threatening inefficiency, corruption, and distorting private sector incentives to invest and innovate in the infrastructure and in the markets dependent on the government infrastructures. Since the 1960s, privatization and regulatory reform movements in many countries have sought to transfer economic management of many important functions from government control to markets, deregulating sectors ranging from air transport to electric power generation. The belief was that the Darwinian forces of market-based competition would do a better job at directing investment, incentivizing innovation, and driving toward greater efficiency than government command and control regulation. This systemic trend has included the privatization and deregulation of telecommunications infrastructure and services.

In the U.S., competition began to replace regulation first in the markets for telephone equipment (1960s), then long-distance services (1980s), and finally in last-mile local networking services (1990s). In the EU, a similar pattern of steps was followed. By 1998, the telecommunications sector was comprehensively liberalized. The transition towards market liberalization did not eliminate the need for regulation, but rather changed the ways in which such regulatory oversight might be framed and managed. To the extent markets are effectively competitive, they remain subject to general competition laws and policies, the most important of which relate to antitrust law. These address issues such as merger review and specific instances of alleged anti-competitive behavior and may be enforced by general competition authorities (e.g. the Federal Trade Commission or FTC in the U.S. and similar agencies in other countries) and via Court proceedings alleging antitrust violations of general competition law. In addition, sector-specific regulators may be empowered to address issues that require specialized expertise or are of concern for a specific sector.

In the U.S., the Federal Communications Commission (FCC) is the sector-specific regulator that was established in 1934 to oversee the regulation of telecommunications networks and services. The responsibilities of a telecommunications regulator includes ensuring (i) consumer protection; (ii) universal access to basic telecommunication services; and (iii) effective competition wherever feasible in telecommunications markets. To fulfil its mandate, sector-specific regulators like the FCC rely on a portfolio of regulatory tools and instruments, including the ability to mandate disclosure and transparency rules, set price and service terms and obligations, regulate business practices, institute subsidy programs, and convene and manage the adjudication of disputes. The portfolio of tools and actions includes both *ex ante* and *ex post* interventions [46, 112].

To protect other sectors of the economy, manage jurisdictional overlaps, and prevent runaway regulatory authorities, sector-specific regulators are themselves subject to regulatory oversight limiting their authority. In the case of the FCC, that includes the basic legislation that created the FCC (the Communications Act of 1934 as amended), Congress which must approve appropriations and may legislate new authority, the Courts which can interpret existing legislation, and the Executive branch which has authority to appoint the Commissioners that direct the FCC.

Similar frameworks and mechanisms exist in Europe, although the situation is more complex since each of the member states that comprise the European Union (EU) have more significant autonomy and each has its own National Regulatory Authority (NRA). In Europe, there is no single sector-specific regulator like the FCC with community-wide authority. Instead, the Body of European Regulators for Electronic Communications (BEREC) serves an advisory role with representation from each of the member-state NRAs [8, 9].

Today, the Internet has evolved into a broadband-centric ecosystem and platform for innovation that is at the center of the global transformation to a Digital Economy that includes a plethora of significant network-related developments from 5G mobile broadband to cloud computing, from the Internet of Things to augmented reality. Broadband has developed into a general-purpose technology that delivers an evolving range of Internet-based and non-Internet-based (legacy and IP-based) applications [69]. Significantly, the Internet has transformed from being just one of the many applications that were overlaid on and shared the regulated telecommunications infrastructure with the PSTN to be a core data communications platform, which now partially supports global telephony services. Today, telephony is just one of an evolving range of applications that may be overlaid on the Internet. Unsurprisingly, the rise of a broadband-centric Internet ecosystem has given rise to new challenges for the design of an appropriate regulatory framework.

In confronting the challenge of how to transition telecommunications regulations crafted for a legacy PSTN for the broadband Internet, U.S. and EU policymakers have followed somewhat divergent paths. In the U.S., the transformation of legacy cable television and telephone networks into general-purpose broadband networking platforms that could offer similar arrays of video, voice and data services, including Internet access, held the promise of duopoly facilities-based competition in most markets. This challenged traditional notions that last-mile networks were non-contestable natural monopolies (so-called essential facilities or monopolistic bottlenecks [59]) that warranted strict *ex ante* regulations. Instead, it opened the way toward further reliance on competition to ensure desirable market outcomes. Consequently, the FCC embarked on a path to gradually but aggressively deregulate broadband. By 2005, both telephony-based xDSL and cable-modem based broadband services were classified as "information services" and thus freed from the stringent common-carrier regime that had been imposed upon traditional telephony and telecommunications networks under Title II of the U.S. Communications Act of 1934 [3, 54].

EU regulators, confronting the challenges of coordinating policies across multiple NRAs and with larger portions of the population lacking coverage by duopoly facilities-based providers (e.g., legacy cable infrastructure was wholly lacking in Greece), were

slower to deregulate broadband facilities and relied more heavily on service-based competition [19, 78].

The divergence in regulatory approaches resulted in different competitive landscapes in the U.S. and Europe. While end-users in the EU could typically choose between a relatively large number of competing broadband access service providers (providing service either using the facilities of the legacy telephone network provider that remained regulated or cable infrastructures where available), in the U.S., end-users' choices of broadband providers were more limited but included services offered over different underlying facilities-based network providers which were more lightly regulated [19, 48, 79, 102].²

3 THE NETWORK NEUTRALITY DEBATE

The emergence of NN is closely tied to the deregulation process described above. Tim Wu [115, 116], who is often referred to as having coined the term and narrative underlying the concept of NN, traces the origin of the NN debate [117] to rising concerns in the U.S. in the late 1990s that providers of access services to the public Internet (i.e., the legacy cable and telephone companies that owned and operated the last-mile broadband platforms) might use their market power to interfere with the freedom of Internet users and providers of Internet applications, content, and other services to use the Internet, thereby threatening the openness of the Internet. From a conceptual perspective, NN can thus be considered as an attempt to craft a framework for the regulation of last-mile broadband in a post-PSTN world. It shifts the focus of regulatory intervention from mandating non-discriminatory access to the infrastructure needed to provide competitive alternatives for broadband access to imposing rules for how broadband access service providers manage Internet traffic. NN regulations codify non-discrimination rules and aim at establishing a regulatory dividing line to distinguish between desirable, and thus permissible forms of differentiation (or discrimination), and those forms that are considered undesirable and should be forbidden [102].

The challenge for NN regulation is to provide a framework for determining what constitutes reasonable network management, or more specifically, traffic management practices in a complex and evolving technical, business, and policy ecosystem of interconnected computing and network resources, applications, and services. This challenge intersects with the challenge of ensuring appropriate sector-specific regulation for national (and ultimately international) communications infrastructure services.

²In the context of broadband Internet access, facilities-based (or network-based) competition occurs when end users can choose between two or more separate facilities-based networks offering services in the same geographic market. These could be using the same technology (e.g., multiple cable networks) but more typically use different technologies (e.g., telephone and cable network). Service-based competition occurs when there are multiple providers of retail services, and some of these may rely partially or fully on the wholesale network services provided by the underlying facilities-based network providers. Retail-level service-based competition may occur even in the absence of regulatory mandates requiring network operators to provide wholesale network services, although typically regulatory mandates do exist when the network operators are regarded as controlling bottleneck facilities. Service-based competition is complementary to facilities-based competition when facilities-based competition is feasible, but in cases where there is no facilities-based competition, it offers the only viable form of competition and leads to a situation in which multiple broadband Internet access service providers can offer their services even if there is only a single physical access network available [77, 101].

3.1 Network Neutrality in the U.S. and the EU: The Regulatory Status Quo

NN regulations differ across jurisdictions, however, reliance on transparency obligations to help ensure consumer protection are a common denominator. More controversial issues relate to (i) the scope of regulations (i.e., what type of traffic and which market players should be subject to the regulations), and (ii) the definition of what constitutes reasonable and thus permissible network management [102].

In Figure 1 we mark the main milestones in the history of NN regulation in the US and EU.³ The first major step towards crafting NN rules was made in the U.S. in 2005, when the FCC adopted its Internet Policy Statement, setting forth principles intended to ensure and protect consumers' freedom of choice and right to access legal content, use lawful applications, connect safe devices, and select among a competitive selection of choices for service, application, and content providers [38, 90]. Also in 2005, the FCC sanctioned a regional ISP, Madison River Communications which was blocking VoIP services in an effort to preserve the profits they obtained from legacy voice services [36, 37].

In 2007, concerns over broadband access providers abusing their market power were heightened when Comcast was discovered to have been disconnecting end-user BitTorrent sessions without prior notice in an effort to free up network resources being used by the peer-to-peer application [39, 48, 108]. This was widely recognized as an unacceptable network management practice and resulted in the FCC fining Comcast and led to the FCC's first attempt to codify its Internet Policy statement into regulatory rules under its first Open Internet Order (OIO) in 2010 [40]. Because the FCC had previously deregulated broadband services, the 2010 OIO was challenged in U.S. Courts which eventually struck down major provisions of the rules in 2014 [48]. In 2015, the FCC issued a new OIO [41] that largely mirrored the earlier OIO but relied on reclassifying Broadband Internet Access Services (BIAS) as a telecommunications service that was subject to the same Title II legacy regulations that broadband had previously been exempted from, thereby addressing the earlier Court's reasons for striking down the 2010 OIO. With the election of President Trump in 2016, the FCC under the newly Republican-appointed chairman, rescinded the 2015 OIO in their 2018 Order [42], prompting a number of States to enact their own versions of OIO. This resulted in another wave of legal challenges that have culminated in a seminal decision by the D.C. Court of Appeal, which largely upheld the 2018 Order [48, 110].

Meanwhile in the EU, a similar pattern of steps were followed, although in the EU the deregulatory pressure was less acute, in part because facilities-based duopoly competition among BIAS providers was less extensive and because EU member states have generally been more accepting of government regulation than the U.S. The EU adopted a revised Regulatory Framework for Electronic Communications in 2009. A declaration on NN was published, but the proposed means to ensure NN were based on transparency obligations and the possibility to impose "safety net"-style minimum Quality of Service (QoS) standards by NRAs [30, 31].

³For a more detailed overview of these milestones, see [101].

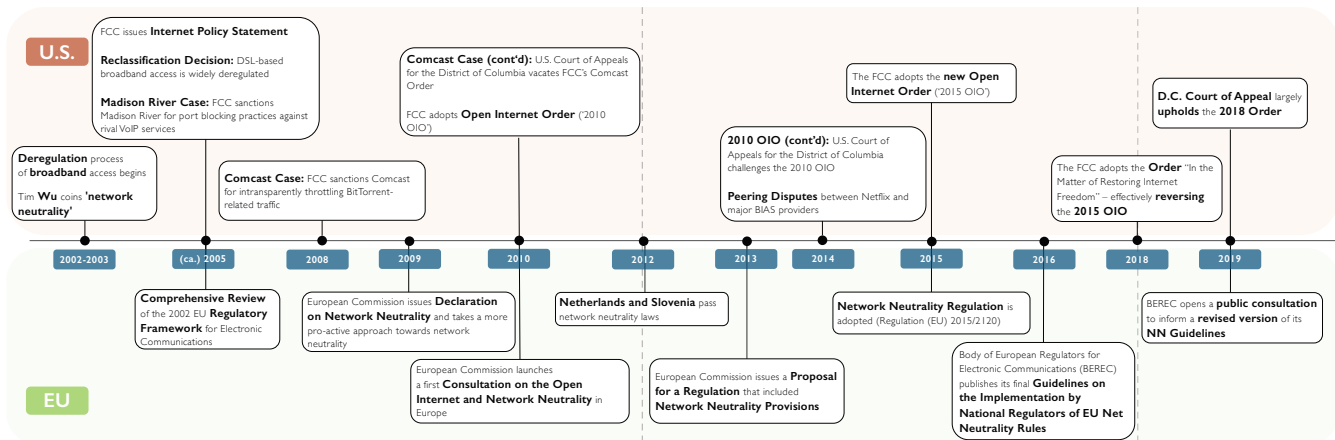


Figure 1: The History of Network Neutrality Regulation: (top) in the U.S. and (bottom) in the EU.

Major steps were to be made in anticipation of, and arguably responding to, undesirable trends towards regulatory fragmentation as indicated by the NN laws that were passed in the Netherlands and Slovenia in 2012 [78, 80], the European Commission (EC) issued a proposal for regulation that contained elements of NN regulations in 2013 [32]. After several further years of debate, the European Commission, the Parliament, and the Council reached a consensus which paved the way for the adoption of Regulation (EU) 2015/2120 [33]. In many respects, this regulation revealed strong similarities with the approach taken by the FCC in their 2015 OIO. In 2016, BEREC adopted guidelines for the implementation of NN regulation within EU member states [5]. On November 28, 2019, BEREC closed a public consultation to inform a revised version of its NN Guidelines [10–12]. A reassessment is under way and a revised second version of the BEREC guidelines is expected by 2020. Meanwhile, the European Electronic Communications Code (EECC) has stipulated that broadband Internet access should be a universal service via Directive (EU) 2018/1972 [34, 102].

3.2 Network Neutrality Rules in Practice

As noted above, NN rules have typically focused on the traffic management practices employed by BIAS providers. For example, paragraphs 104 to 108 of the 2015 OIO [41] in the U.S. required BIAS providers to adhere to three bright-line rules and a business conduct standard that included prohibitions against:

- **Blocking** of lawful content, applications, services, or non-harmful devices
- **Throttling** on the basis of content, applications, services, or devices
- **Paid prioritization**
- **Unreasonable interference or disadvantage** to consumers' and edge-providers' access to the open Internet

In addition, the 2015 OIO imposed transparency requirements on BIAS providers to disclose to consumers and edge providers their network management practices ([41], para. 109). It is noteworthy that these rules did not prohibit BIAS providers from offering multiple tiers of service, with different peak data rates, data caps,

and pricing terms. In prohibiting the blocking of lawful content or applications, the rules explicitly exempted network management practices designed to address digital piracy or access to prohibited content such as child pornography. In prohibiting non-harmful devices, the NN rules also intended to exempt traffic management practices directed at stopping malware or denial of service attacks. Also, in blocking paid prioritization, the rules did not apply to “reasonable network management” practices that might result in traffic prioritization to facilitate efficient sharing of network resources during (transient) states of congestion which might include the prioritization of packets for delay intolerant application traffic (e.g. VoIP) over delay tolerant application traffic (e.g., email).

3.3 Ambiguity of NN Regulation

The two biggest challenges that NN regulations confront relate to (i) determining the scope of regulations, i.e., what type of traffic and market players should fall under the regulation, and (ii) the definition of what constitutes reasonable and thus permissible network management [102], see Figure 2.

Addressing these challenges is complicated by the fact that NN rules are inherently ambiguous. *First*, NN regulation is targeted at protecting the (public) Internet. As such, it does not apply to all IP-based traffic. There are many IP networks operated by enterprises, service providers, and others that are not part of the public Internet, and even when they interconnect with the public Internet, not all of the traffic that is carried is public Internet traffic. As Lehr et al. [69] have shown, it is far from clear what is meant by “the public Internet.”

Second, despite the obvious fact that the experience of Internet users might separately and interactively be shaped by a host of factors and entities that are involved in the end-to-end delivery of Internet traffic, the rules do not apply equally to all entities. Instead, NN rules have focused on regulating broadband access platform providers and their provision of BIAS. Only BIAS traffic is subject to the rules, not other IP traffic or services that may be carried via

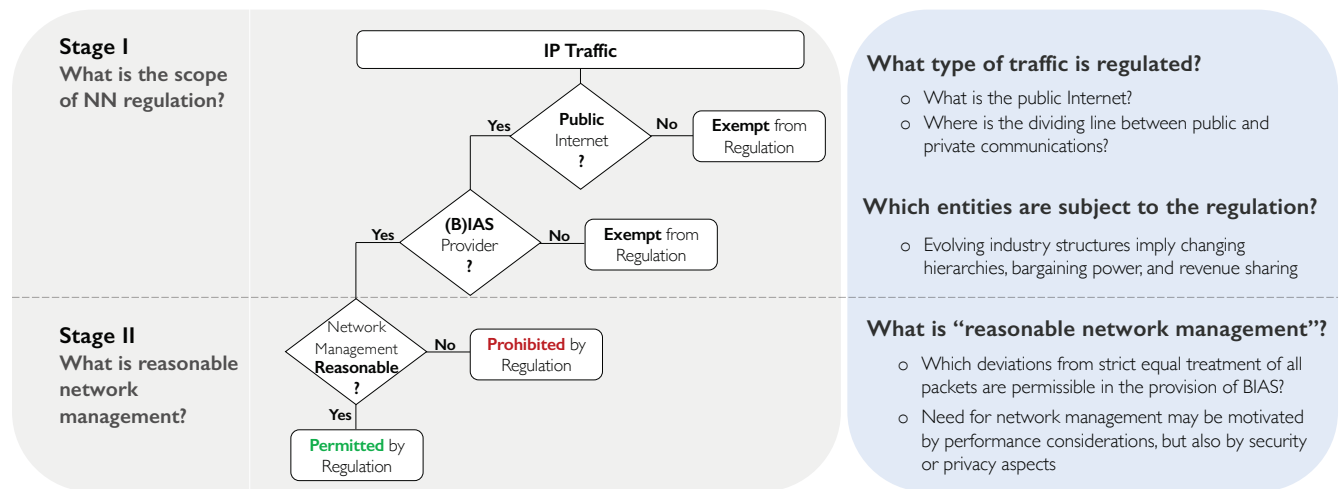


Figure 2: Network Neutrality Regulation: Complexity and Ambiguity.

the broadband networks. This other traffic is usually referred to as specialized services [102].⁴

Third, NN rules differ with respect to the precise requirements of what constitutes acceptable traffic management. Although the general intent of the NN rules is similar across jurisdictions—to enforce a degree of egalitarian treatment of Internet traffic—the actual rules differ with respect to who the rules apply to and precise guidance on their interpretation is often lacking.⁵ Such ambiguity contributes to regulatory uncertainty, but is partially by design. The laws that enfranchise regulatory authorities and the rules promulgated by regulatory agencies are designed to be sufficiently flexible to allow context specific factors to be considered and to provide scope for changing technology, markets, and industry/business structures. Often, the regulators are less informed about the details of specific networking situations or technologies than are the engineers and business decision-makers who own and operate the networks. Regulators want to allow scope for market participants to experiment and innovate, and to earn appropriate rewards (profits) for their innovative activities. The Internet experience is enhanced for all when networking resources can be shared efficiently, malware and harmful traffic can be blocked, and a multitude of heterogeneous applications and services can harmoniously co-exist. Thus, allowing for innovation in network management practices is also important.

A lot of the innovation within the Internet ecosystem—from the edges to the core, from the fundamental physical layer transport technologies (fiber optics, wireless) to higher-level applications support and content delivery networking and the design of client applications—have evolved and contributed to enhancing the Internet experience. Much of this involves employing complex traffic management techniques that might be construed as network management. However, many of the same techniques that may be

employed to enhance the Internet experience might also be used to impair it either selectively (to harm particular end-users, edge providers, or applications) or more generally (result in an overall degradation of the open Internet, potentially to direct traffic to non-Internet services and networks). Unfortunately, given the complexity of what constitutes network management in today’s Internet ecosystem, translating general frameworks such as the 2015 OIO rules or Regulation (EU) 2015/2120 [33] into unambiguous and consensus-based rules that might be followed and enforced by network engineers operating the networks that support and contribute to the Internet is not feasible. What constitutes “reasonable network management” is simply not amenable to a static and unambiguous codification into rules that do not allow scope for contextual interpretation.⁶

4 NETWORK NEUTRALITY LOOPHOLES

NN regulations and their inherent ambiguity give rise to a number of “loopholes” by which the intent or effect of the rules may be bypassed or undermined, see Figure 3. In the following sub-sections we highlight a variety of methods that might be used to render NN regulations less effective, or worse, counterproductive with regards to their intended goal of protecting the public Internet and ensuring reasonable network management. Since many of these “loopholes” accelerate content delivery and improve end-user experience [117], were a revised NN framework to eliminate those loopholes, the end-user experience might suffer.

4.1 Specialized Services

One of the first loopholes arises because NN rules apply to BIAS traffic, but not to other IP traffic classified as “Specialized Services”

⁴The term “specialized services” is used in the U.S. and the EU to refer to services that are not identified as (B)IAS. Some of these may be regulated under different frameworks, e.g., [48]. In the U.S., these services have also been called non-BIAS data services (see also Section 4.1).

⁵For example, it is not clear whether reasonable network management rules should be interpreted as permitting practices which are not explicitly prohibited or prohibiting practices which are not explicitly permitted.

⁶For example, consider the ambiguity in NN rules when applied symmetrically to fixed and mobile broadband. In light of the fact that congestion or availability problems are more likely to arise in the context of wireless mobile broadband services and because mobile telecommunications regulations have not been fully harmonized with fixed telecommunications regulations, mobile BIAS is often afforded greater scope for engaging in network management practices that might be prohibited in the case of a fixed provider.

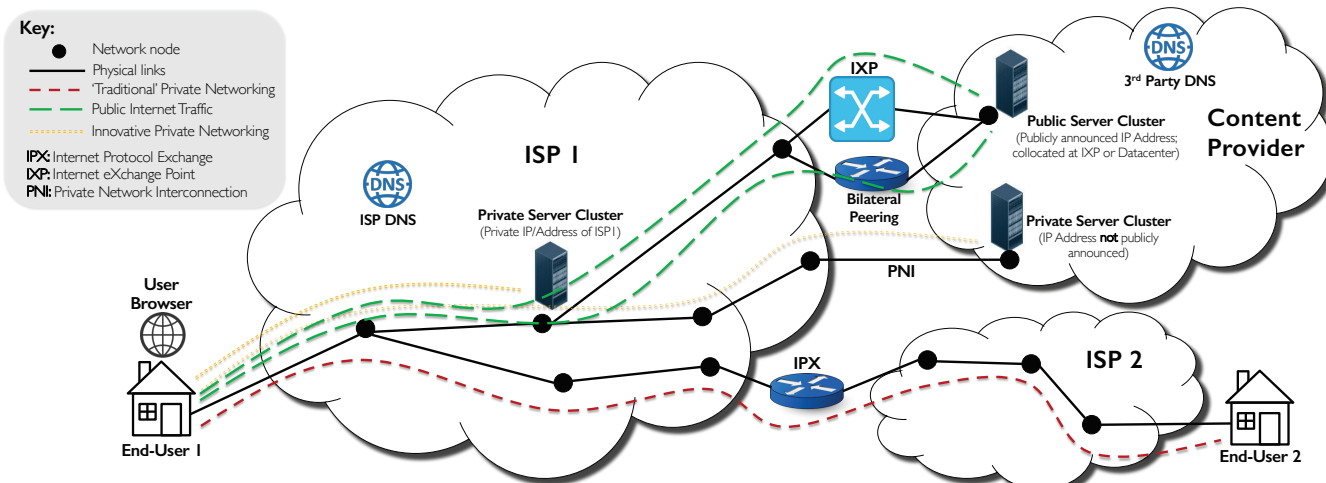


Figure 3: Loopholes for circumventing Network Neutrality regulation.

that are delivered via the same last-mile access networks. ISPs wishing to circumvent NN rules may strategically label IP traffic so that it is not identified as public Internet BIAS traffic, thereby effectively exempting such traffic from application of the NN rules. Specialized services or non-BIAS data services are typically provided via ‘closed’ or ‘private’ IP networks (‘private’ server clusters that serve only the users of an ISP). These networks can be customized to support the optimized delivery of specific QoS-sensitive applications, e.g., IPTV, facilities-based VoIP, or VoLTE. The reach of these services can vary. While IPTV services are inherently local and provided by an ISP exclusively for its customers, the provision of VoIP services requires universal connectivity. As described in [25], IP-based voice services might be carried via private networks using application-specific interconnection points like the Internet Protocol Exchange (IPX). The customization enables the ISP to provide service qualities that could not be achieved via a regulated BIAS and might involve using logically separated network capacities and/or different types of traffic management practices [5].

The bifurcation into a public Internet and non-public IP networks establishes a regulatory split that requires a clear dividing line to determine what the public Internet is, but also where it ends and private (non-Internet) networking begins. Making such a determination, however, presents a difficult challenge [60, 69]. Uncertainty regarding the regulatory treatment of differentially labeled IP traffic (as Internet or non-Internet) contributes to regulatory uncertainty and may create opportunities for regulatory arbitrage. ISPs might argue that, from a technical perspective, the “public Internet” should be defined as comprising all IP prefixes that are advertised to all networks (i.e., ASes) using the BGP-based routing system. If this definition is accepted, then IP networks that advertise prefixes only to a subset of prefixes and carry traffic via routes that are not advertised via the BGP protocol do not provide Internet connectivity. Therefore, they could be considered as non-Internet specialized service networks.

Many edge providers have moved their front-end servers—and thus content, applications, and computation capabilities—closer to end-users (or things). This may be done to economize on transport,

avoid congestion points, and improve latency performance. These servers often communicate to back-end servers in datacenters via private network links [23, 44, 70, 92] rather than via the public Internet. Large content providers have invested in their own private backbone to interconnect their datacenters and, thus, avoid the public Internet [56, 98]. Also, specific 5G-based use cases rely on network slicing and local processing of data in edge clouds [68, 122]. In many instances, communicating endpoints are positioned within the same network (i.e., AS).⁷ As the traffic between relevant endpoints is thus confined within the borders of a single AS, it constitutes an intra-network service that might be considered as private networking. The routing of traffic to a front-end server that is owned by the access ISP or a third party and hosted on the access ISP’s network may also be carried via a private IP network. Arguably, such services may not be dependent on BIAS and may not be regarded as part of the public Internet, thereby exempting them from NN rules.

NRAs might challenge such practices if they determine that the only purpose of such strategies is to circumvent NN rules. However, reaching such a finding would require a context dependent inquiry with an uncertain outcome since clear delineation criteria for classifying traffic as BIAS or specialized services do not yet exist. In view of the dynamic nature of service provision and rapid ecosystem evolution, static criteria for delineating between BIAS and specialized services will be unlikely to keep abreast of the changing nature of the public Internet. Additionally, care must be taken as efficiency and performance might both suffer if the specialized services used to support content delivery were not permitted.

4.2 Content Redirection

By some measures, CDNs are currently responsible for about 80-90% of the traffic that is delivered to residential users [65, 91]. The link

⁷Deploying servers inside an ISP network to deliver content or applications render end-to-end communications an intra-network service, bypassing the “public Internet” and providing possibilities for regulatory arbitrage [69, 103]. For example, Akamai [85], Google Global Cache [51], or Netflix Open Connect [83] pursue such highly distributed CDN strategies. These deployments often use ISP-owned address space [15, 17, 106].

between CDNs and NN has been examined in several publications, e.g., [29, 81, 103, 113, 121]. Highly distributed CDN and (edge) cloud computing approaches help to bring content, applications, and general cloud capabilities (e.g., data processing) closer to end-users. When it comes to server deployment strategies, CDN paradigms vary considerably. While some approaches rely on datacenters that consist of large server farms that are strategically deployed at central communication hubs like Internet eXchange Points (IXPs), highly distributed approaches rely on server deployments within ISP networks. While an overview of the different approaches is provided in [103], intra-ISP deployments eliminate performance bottlenecks at interconnection points, and problems related to cross-domain routing [23, 70]. The latency between end users and a cache inside an ISP is low, sometimes as low as a few msec [92, 109] in metropolitan areas. End user connections are typically terminated at the front-end CDN server and a second connection is established by the front-end server to other CDN servers, i.e., front-servers act as a reverse proxy [44, 85]. In this way, CDNs can innovate and roll out new protocols to accelerate content delivery, by relying on server-to-server communication without requiring any updates from the clients.

Content redirection using CDNs enables a variety of selective redirection strategies that provide additional means for delivering potentially discriminatory services, without engaging in discriminatory traffic management. End-user requests for a specific piece of content or a website can be redirected in a strategic fashion so that the Quality of Experience (QoE) can be varied considerably, even though there is no differential treatment in the transmission of different packets or flows through the network(s). Via the Domain Name System (DNS), the content source can be chosen in a biased fashion. For example, an end-user request can either be matched with a lightly utilized or a congested server. Similarly, it can be redirected to a server that is close or to one that is distant from the requesting end-user's location. In the past, end-users heavily relied on the DNS service provided by their ISPs, and their requests were mapped accordingly [62]. Today, third-party DNS services, e.g., by edge providers like Google DNS or OpenDNS [93], have been gaining in importance. Extended DNS (EDNS) can improve mapping between end-users and CDN servers and is supported by some of the larger CDNs [22]. CDNs may also redirect user requests to servers to achieve their own performance and cost objectives [85]. While these recent developments imply that ISPs might to some extent lose control over the DNS, ISP-CDN collaboration schemes have been proposed and are currently being implemented [91] to improve the coordination between ISPs and CDNs, resulting in a win-win situation for both parties, as well as enhancing the end-user experience.

4.3 Routing and Interconnection

The choice of routing can also be used in a variety of ways to bypass NN regulations. These include routing traffic via private interconnections, manipulation of the route chosen for inter-domain routing (e.g., to differentially select interconnection routes that are more or less congested), or other forms of selective IP-based routing.

Private Interconnections. As explained in previous sections, content hosted on the public Internet and traffic that may eventually be delivered via BIAS may traverse private networks for part or most of its path from source to destination. Private interconnection points involve the physical interfaces (Private Network Interconnections) that connect different networks, allowing them to exchange traffic. The capabilities and arrangements governing these interconnects can vary widely. They may take place at co-location facilities where interconnecting ISPs, large content providers, or clouds [119] maintain the interconnects, at data centers, or at IXPs. While ranging widely in capacity, they can reach up to tens of Gbps for each individual connection [16]. The private interconnections typically rely on reserved/private IP addresses for the router interfaces; these addresses are not advertised in the public routing system for security reasons and to preserve IPv4 address space. While in the U.S., interconnections are often negotiated as customized bilateral agreements, in the EU, a significant fraction of the ISPs use “public peering”—i.e., peering across a Layer 2 shared switch fabric [21, 49]—to interconnect with and receive traffic from content providers [94, 114]. Internet interconnections have remained widely unregulated, which has resulted in an evolving range of agreement types encompassing peering, transit, paid peering, partial transit, and a variety of other arrangements. The management of private networks and private interconnections are not part of the public Internet. Thus, the management of traffic at those interconnection points has not been subject to NN regulations.

Peering and Congestion. Unresolved coordination problems between interconnecting parties might result in situations where degradations in the end-user QoE occur as ISPs strategically (under)invest in interconnection capacity. In the past, interconnection tussles between major ISPs in the U.S. and Netflix had gained public attention [58, 73, 100]. Broadly speaking, ISPs had tried to induce Netflix to accept paid peering arrangements by refusing or strategically postponing the expansion of interconnection capacities called for by the significant increase in asymmetric traffic associated with BIAS subscribers watching Netflix content (i.e., downstream from Netflix servers to broadband subscribers, with relatively little traffic flowing in the opposite direction). The result was prolonged periods of inter-domain congestion and the degradation of end-user QoE. It also may have motivated the FCC to consider interconnections in their 2015 Order (at least on an ex post case-by-case basis) [78]. Congestion arising from such business disputes over who should pay for the increased investment necessitated by growing traffic loads can hardly be considered as a specific network management practice. As such it seems questionable whether such strategies would fall under the scope of NN rules.

Selective IP Routing Policies. Similar to the selective redirection strategies described in Section 4.2, a host of strategies and tools for intra- and inter-domain traffic engineering enable the implementation of selective routing policies. While these might be used to optimize network utilization and performance dynamically, to accommodate changing mixes of application traffic with heterogeneous QoS requirements, the inherently selective nature of such techniques may be used or abused to achieve all kinds of service differentiation. For example, some IP prefixes of servers may be announced or routed differently than others. IP routing provides the means to control the length of the path packets have

to travel within a network or across networks. Especially when this involves transit traffic, this might have a significant impact on end-to-end latency levels. To achieve better performance, prefixes can be aggregated or de-aggregated either via BGP, e.g., [63], or via programmable networks, e.g., SDN and P4 [53, 94, 118]. Corresponding differentiation practices will likely not be considered a violation of current NN rules that focus on how packets are managed by the routers since there is no differentiation between packets within the routers along the path; instead, packets are delivered via different paths—and there might be good reasons for that.

4.4 Middleboxes and Edge Control

A last set of loopholes we describe below is associated with (i) using traffic management via in-network “middleboxes”; (ii) innovative strategies to gain control over and optimizing the behavior of communicating endpoints; and (iii) endpoint-based incentive-compatible prioritization strategies.

Middleboxes and Virtual Network Functions. Middleboxes are devices that are deployed within networks and positioned between communicating endpoints. They can be used by ISPs to enable a wide range of functionality. For example, they may be used for deep packet inspection (DPI), as traffic load balancers, or to implement traffic policing strategies. The scope and flexibility of the implementation of such tools have expanded dramatically with the shift towards programmable networks. Advances in networking technologies enhance the agility, customizability, and adaptability of network management and yield dynamic control over network resources and functions and make it feasible to customize and deploy such tools and (virtual) network functions on-demand anywhere [45, 105]. Middleboxes and virtual network functions provide ISPs with powerful tools for traffic engineering [97]. Much of this can be used to the benefit of end-users and may, in many cases, be partially under end-user control. Intelligent use of these tools can optimize the use of network resources, enhance the user experience, and contribute to network security and reliability. Despite their virtues, these tools may also be used to implement ad hoc discrimination that may favor some users over others or otherwise interfere with end-user preferences for how their traffic is managed [24]. The rise of programmable network functionality increases complexity and adds to the difficulty regulators confront in monitoring traffic management practices [97]. The strategic location of such middleboxes and their associated network functionality can impact whether they are upstream or downstream of BIAS services and whether they are subject to NN regulations. It is challenging to detect the existence of such devices, although measurement work has shown that crowdsourcing techniques can detect where middleboxes may be deployed [55, 111].

Taking Control of Both Ends. An interesting, but perhaps lesser-known, means to introduce traffic differentiation can be achieved by gaining control over both communicating endpoints. If end-users request content from a Google-owned server via a Chrome(-based) browser, the behavior of communicating endpoints is effectively controlled by the same entity: Google. This may enable Google to deliver the requested content from a Google front-end server that is positioned close to the requesting end-user, thus

reducing the reliance on the “public Internet”, and may take advantage of innovative protocols such as QUIC [64, 67] that can optimize end-to-end communications ‘on top’ of the Internet. Applications in smartphones may benefit similarly by interacting with special servers, e.g., Google’s Youtube and the various applications of Facebook. Similarly, based on java script, Netflix can optimize the delivery of their content. A recent development is DNS over HTTPS [14] where the DNS requests are sent directly to the CDN bypassing the ISP DNS. In this case the CDN knows the origin of the request and optimizes the reply that is sent to the end-user. With client-server coordination, it is also possible to achieve higher throughput by establishing multiple connections with the same server or with different servers [72], as well as by utilizing multiple paths [88] to download the content. The techniques described above provide means to achieve traffic optimizations and differentiation without impacting NN rules.

Endpoint-based Incentive-compatible Prioritization. While NN regulations aim to prevent consumer harm by prohibiting ISPs from unreasonably applying network management practices, a differentiation that reflects end-user preferences and relies on end-user control can hardly be considered harmful. In this context, the study presented in [120] provides meaningful insights. It demonstrates that end-users have an incentive to request preferential treatment. By analyzing demand requests of users in more than 160 homes, the study shows that the individual user profile is heavy-tailed and that the ranking of applications differs across different users. End-users may want to prioritize different applications over others. Based on these insights, the same study introduces and evaluates a simple yet effective mechanism that enables end-users to reveal their preferences based on a price and QoS differentiation scheme, which can lead to a win-win situation for the network providers, users, and content providers alike. For a network economic analysis showing the benefits of incentive-compatible price and QoS differentiation strategies, see [60].

5 NETWORK NEUTRALITY MEASUREMENTS

For NN regulations to effectively constrain behavior, it has to be credible to all relevant industry participants that violations can and will be detected and enforced appropriately. If expected false positive or false negative errors are too frequent, then NN regulations will either be ineffective or distort efficient behavior. The latter might manifest itself in inefficient resource usage but also in distorted investment decisions and innovation incentives [4, 102]. Even when QoE impairments can be reliably detected, it is difficult to determine what the cause was or who should be held responsible. The end-to-end delivery chain is often complex, potentially involving a host of traffic management strategies employed simultaneously by different (competing) entities [104]. Crowd-sourcing measurement methods, e.g., speed tests by content providers such as the Netflix ISP speed index [82] and the Google Video Quality Report [50] may be able to identify that service degradation has occurred and yet be unable to pinpoint the root cause of the performance degradation. Also measurements that aggregate the performance at the level of regions do not preclude individual performance experiences varying significantly. Other crowd-sourcing methods may be able to infer some types of ISP traffic management

practices, e.g., Netalyzr [61] and Haystack [13], yet fail to be able to demonstrate that a NN violation has occurred. Recent measurement work has also shown that many of the performance inefficiencies that are experienced are due to the inherent design or interactions of the protocols and systems involved, e.g., DNS [99] or even the home router [107] rather than because of traffic management practices. Thus, attributing all impairments in end-user QoE to NN violations is certainly a gross oversimplification and will often be wrong.

Assuming in a specific case the NN regulations prohibit blocking, throttling, paid prioritization of content as well as unreasonable interference or disadvantage access to consumers and edge providers (see Section 3.2), then in that case, to be able to identify all the above NN violations, it is imperative to have access to all the routers and middleboxes along the path. Such access is not scalable; many routers/middleboxes across a path are operated by different entities and modern routers have multiple queues. Additionally, granting measurement access to routers/middleboxes for ISPs (even via an application programming interface) would pose a significant security challenge that if not appropriately addressed could jeopardize the operation of the network.

In contrast to this, blocking of lawful content ought to be relatively easy to detect. For example, OONI [87] facilitates inferring different types of censorship blocking of content. The detection of throttling and paid prioritization, however, presents a much more difficult challenge. Managing stochastic traffic in a complex network of networks involves decisions by many parties that jointly contribute to determining the realized QoE. Disentangling the effect of particular actions is difficult and usually depends on observing traffic data over longer periods of time. Moreover, persistent impairments are easier to detect⁸ than those that are of short duration and episodic. Measurement techniques based on statistical inference have been proposed to allow identification of traffic differentiation in the Internet [47, 76, 123] and in mobile data networks [71, 96]. However, it is very challenging to validate whether traffic differentiation strategies violate NN. They might as well constitute reasonable network management, adapting to current network conditions or reflecting application requirements or end-user device capabilities. In this context, comparing end-user experiences across different providers for the same application/service might prove useful in detecting NN violations.

For specific protocols it is possible to reliably detect throttling via measurements. For example, Glasnost [28] was able to detect cases of BitTorrent throttling by ISPs, some involved the reset of TCP connections. However, it is challenging to design generic techniques and tools to reliably infer unreasonable interference with, or disadvantage of, specific end-users, edge providers, or applications. One of the key problem here is to define the right metrics to measure. Coming up with an appropriate measurement regime is complex for a number of reasons. Beyond purely technical challenges, the economics and regulatory approach might vary across different jurisdictions. Moreover, given the rapid pace with which the ecosystem evolves, ensuring NN presents a moving target.

While some of the strategies to manage traffic may involve the use of the loopholes we discussed above, identifying that a particular entity has employed a particular network management strategy does not unambiguously indicate which party may have been responsible for the QoE impairment. While loopholes may be used to enhance customer QoE or lower costs, they may also be used to express market power (e.g., by raising rivals costs or inducing more favorable bargaining outcomes by participants with such power). As a consequence, the ability to measure traffic performance from multiple perspectives (end-to-end) and along portions of the path is essential to reliably diagnose the source of impairments. Effecting such measurements requires coordination and sharing across multiple entities. This presents a challenge. Firstly, providers are potentially competing. Secondly, they can be expected to want to avoid being identified as the guilty party (whether rightly or wrongly). Recent studies propose measurement techniques to improve transparency [84, 89], but more work is needed in this direction. Despite the efforts made by the FCC in the context of their Measuring Broadband America initiative [43] and by BERC [2, 6, 7] to develop appropriate third-party measurement tools, these are yet to produce the desired levels of transparency. Current regulations do not require ISPs to provide access to the information and measurement vantage points that are needed to evaluate whether traffic is being managed in ways that are consistent with the goals of NN rules or not. To date, we are aware of only a single case where an ISP was required to provide partial access to router data to independent experts. The mandate was imposed to ensure compliance with the requirements of AT&T's acquisition of Time Warner [26].

6 THE FUTURE OF NETWORK NEUTRALITY REGULATION

As we explained in previous sections, the QoE does not depend solely on the traffic management practices of broadband access providers nor even how traffic is managed on the public Internet. A variety of mechanisms to enhance delivery performance (e.g., the emergence of CDNs) were developed and are in use to protect against malware, enhance reliability and security, and lower network costs. They bypass or augment traditional Internet services and hence are not addressed by existing NN regulations. In most cases, these mechanisms were developed in response to deficiencies in the basic Internet "best-effort" service model and the needs of more demanding applications and usage scenarios. The need for these enhancements and even more extensive capabilities to support ever-more-demanding QoE applications and a wider-range of QoE needs is likely to grow as we transition to next-generation 5G/6G wireless broadband networks. The growing importance and diversity of (distributed) cloud and edge-based capabilities as well as the future role of the IoT, the complexity and heterogeneity of applications and usage environments, and subsequently of networking requirements with regard to aspects like performance, security, reliability, mobility, or energy-efficiency will continue to expand. Not only will these developments further shift and blur the boundaries between network edges and cores, but also between what is perceived as private and public networking.

Regulatory strategies that focus too narrowly on a single network service such as Internet access service, while ignoring other

⁸For example, see the work on inferring persistent congestion [20, 27, 73].

networks or network services that may be used in conjunction with or as substitutes for Internet services will risk being irrelevant with respect to their ability to impact end-users QoE or how traffic is actually managed. Additionally, those same mechanisms that may be used to augment and enhance the QoE of Internet access may be used to bypass narrowly focused NN rules that focus on specific services, management practices, or market participants.

The goal of NN regulations is to protect an open Internet and ensure a good QoE for users, and meeting that challenge in the dynamic and evolving Internet will require a NN framework that is capable of adapting to changing market and technology trends. This will include identifying either implicitly or explicitly the minimum standard of performance that should be delivered from basic Internet access services to meet the fundamental goal of NN regulations. In addition, regulators will need an evolving framework for determining how traffic management practices should be constrained and which market participants are subject to those traffic management constraints.

Recently, NRAs and other stakeholders have made efforts to reassess the regulatory status quo. A recent document by the Internet Society [57] recognized that NN rules that fail to account for the role of edge providers like CDNs, specialized services, (dynamic) routing adjustments and other in-network mechanisms for impacting QoS are unlikely to impose any effective regulatory restraint. In the EU, the recent draft of BEREC's revised NN guidelines [12] indicates a tendency toward granting NRAs more control over all services offered via the broadband platforms of access ISPs. This tendency would imply returning to the sort of control NRAs had over PSTN providers and relying more on ex post enforcement on a case-by-case basis rather than ex ante specific rules. In the U.S., the situation is different and less certain. Currently, there seems to be a tendency to be more de-regulatory but also to rely increasingly on ex post (case-by-case) enforcement. This reflects a recognition that flexibility of network management and diversity of practices that can be used to enhance quality and efficiency makes it difficult to identify specific practices that should be explicitly prohibited or alternatively should be explicitly mandated. For markets, this may be appropriate since it leaves discretion for variation in practices with guidelines that can be enforced when flagrant violations occur. While a partisan divide and changing majorities within the FCC create uncertainty about the regulatory stance towards NN regulation in general, it is currently unclear how Federal and state authority to regulate in this case will be decided. The recent Appeals Court decision affirmed the FCC's ability to reverse its 2015 NN rules (see Section 3.1).

Schulzrinne emphasizes that NN revolves around economic incentives. He argues for updated NN principles, namely (i) end-user choice; (ii) economic neutrality and virtual structural separation; and (iii) transparency [95]. Similarly, Misra *et al.* [74, 75, 81], argue that the NN debate is primarily about economic incentives. The authors propose a non-regulatory alternative, based on incentives, to achieve the goals of NN.

An essential requirement for the effective enforcement of NN rules is the availability of appropriate tools that enable the reliable detection and subsequent sanction of anti-competitive behavior on the Internet (e.g., through the application of ex-post antitrust policies). Appropriate transparency and disclosure mandates that

require ISPs as well as other relevant market players that have an impact on the QoE to disclose their traffic management practices and/or advertise the performance of their public services will prove important for successful NN management. Developing suitable measurement metrics and tools to detect and verify that policy commitments are being met, however, is complex and represents an open research challenge. Making progress toward developing better performance metrics, reporting and assessment capabilities will require cross-disciplinary, multi-stakeholder collaboration involving network operators, equipment vendors, edge providers of content and applications, as well as third-party analysts, end-users and policymakers.

7 CONCLUSION

This paper provides an introduction to why the NN debate is complex and contentious. It also explains why engineers need to be informed and contribute to the debate to ensure they can stay on the right side of the law and help make sure the law (and future policies) stay on the right side of sound network engineering practice. The paper has further elucidated that answering relevant questions inherently and fundamentally requires an interdisciplinary perspective. What needs to be considered is that insufficiently informed and thus, ill-defined regulatory interventions based on an outdated mental model of the Internet will distort investment incentives and harm innovation.

Even though some readers may find it unsatisfactory, the recognition of the complexity of the issues at hand, and how they have (not) been effectively resolved, renders it likely that NN will remain an on-going subject for fierce debate. Importantly, we should not expect network engineers or policymakers to agree on what the best strategies are to preserve an open and innovative Internet, or even what such a goal really ought to mean. Traffic management and the networking capabilities and infrastructure that enables and applications that necessitates it are continuing to evolve and fuel the debate about the sustainability of the current and future Internet.

ACKNOWLEDGMENTS

We would like to thank the fellow panelists, Scott Marcus, Raimo Kantola, and Ignacio Castro, as well as the attendees of the 30th ITS European Conference Panel on Network Neutrality for their feedback on topics presented in this article. We further acknowledge helpful feedback by the editor Olivier Bonaventure as well as by Günter Knieps and Pavlos Nikolopoulos. This work was partially supported by the Federal Ministry of Education and Research of Germany (BMBF) under grant no. 16DII111 ("Deutsches Internet-Institut"), the European Research Council (ERC) grant ResolutioNet (ERC-StG-679158), and the National Science Foundation (NSF) awards CNS-1413973 and CNS-1413905.

REFERENCES

- [1] J. Abbate. *Inventing the Internet*. MIT Press, 1999.
- [2] A. Aldabbagh. Briefing on BEREC's NN QoS Work Policy And Regulatory Aspects of Broadband Services Mapping. ITU-EC, 2016.
- [3] J. M. Bauer. Unbundling Policy in the United States Players, Outcomes and Effects. *Communications and Strategies*, 57(1), 2005.
- [4] J. M. Bauer and G. Knieps. Complementary Innovation and Network Neutrality. *Telecommunications Policy*, 42(2), 2018.

- [5] BEREC. Guidelines on the Implementation by National Regulators of European Net Neutrality Rules. Guidelines, BoR (16) 127, 2016.
- [6] BEREC. Net neutrality measurement tool specification. BoR (17) 179, 2017.
- [7] BEREC. Net Neutrality Regulatory Assessment Methodology. BoR (17) 178, 2017.
- [8] BEREC. Body of European Regulators for Electronic Communications: Tasks. <https://berec.europa.eu/eng/about%5fberec/tasks>, 2019.
- [9] BEREC. Body of European Regulators for Electronic Communications: What is BEREC? <https://berec.europa.eu/eng/about%5fberec/what%5fis%5fberec>, 2019.
- [10] BEREC. Draft BEREC Guidelines on the Implementation of the Open Internet Regulation. Guidelines, BoR (19) 179, 2019.
- [11] BEREC. Public consultation on the document on BEREC Guidelines on the Implementation of the Open Internet Regulation. <https://berec.europa.eu/eng/news,onsultations/Closedpublic,onsultations/2019/6309-public-consultation-on-the-document-on-berec-guidelines-on-the-implementation-of-the-open-internet-regulation>, 2019.
- [12] BEREC. Public Consultation on the draft BEREC Guidelines on the Implementation of the Open Internet Regulation. Public Consultation, BoR (19) 180, October 2019.
- [13] ICSI-UC Berkeley and IMDEA Networks. The Haystack Project. <https://haystack.mobi/>, 2019.
- [14] K. Borgolte, T. Chattopadhyay, N. Feamster, M. Kshirsagar, J. Holland, A. Hounsell, and P. Schmitt. How DNS over HTTPS is Reshaping Privacy, Performance, and Policy in the Internet Ecosystem. In *TPRC* 47, 2019.
- [15] T. Böttger, F. Cuadrado, G. Tyson, I. Castro, and S. Uhlig. Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN. *ACM CCR*, 48(1), 2018.
- [16] Broadband Internet Technical Advisory Group Report (BITAG). Interconnection and Traffic Exchange on the Internet, 2014.
- [17] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *ACM IMC*, 2013.
- [18] A. J. Carrillo. Are There Universal Standards for Network Neutrality? *University of Pittsburgh Law Review*, 80(4), 2019.
- [19] M. Cave and P. Crocioni. Does Europe Need Network Neutrality Rules? *International J. of Communication*, 1, 2007.
- [20] B. Chandrasekaran, G. Smaragdakis, A. Berger, M. Luckie, and K-C. Ng. A Server-to-Server View of the Internet. In *ACM CoNEXT*, 2015.
- [21] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is More to IXPs than Meets the Eye. *ACM CCR*, 43(5), 2013.
- [22] F. Chen, R. K. Sitaraman, and M. Torres. End-User Mapping: Next Generation Research Routing for Content Delivery. In *ACM SIGCOMM*, 2015.
- [23] Y. C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are We One Hop Away from a Better Internet? In *ACM IMC*, 2015.
- [24] D. Choffnes, P. Gill, and A. Mislove. An Empirical Evaluation of Deployed DPI Middleboxes and Their Implications for Policymakers. In *TPRC* 45, 2017.
- [25] KC Claffly and D. D. Clark. Adding enhanced services to the internet: Lessons from history. *J. of Information Policy*, 6(1), 2016.
- [26] KC. Claffly, A. Dhamdhere, D. D. Clark, and S. Bauer. Report of AT&T Independent Measurement Expert Background and supporting arguments for measurement and reporting requirements. UC, San Diego's Center for Applied Internet Data Analysis, 2016.
- [27] A. Dhamdhere, D. D. Clark, A. Gamero-Garrido, M. Luckie, R. K. P. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. C. Snoeren, and kc claffly. Inferring Persistent Interdomain Congestion. In *ACM SIGCOMM*, 2018.
- [28] M. Dischinger, M. Marcon, S. Guha, K. P. Gummadri, R. Mahajan, and S. Saroiu. Glasnost: Enabling End Users to Detect Traffic Differentiation. In *NSDI*, 2010.
- [29] R. F. Easley, H. Guo, and J. Krämer. Research Commentary—From Net Neutrality to Data Neutrality: A Techno-Economic Framework and Research Agenda. *Information Systems Research*, 29(2), 2018.
- [30] European Commission. Directive 2009/136/EC of the European Parliament and of the Council of 25 November 2009 Amending Directive 2002/22/EC on Universal Service and Users' Rights Relating to Electronic Communications Networks and Services (Recast), 2009.
- [31] European Commission. Directive 2009/140/EC of the European Parliament and of the Council of 25 November 2009 Amending Directives 2002/21/EC on a Common Regulatory Framework for Electronic Communications Networks and Services (Recast), 2009.
- [32] European Commission. Proposal for a Regulation of the European Parliament and of the Council Laying Down Measures Concerning the European Single Market for Electronic Communications and to Achieve a Connected Continent (Recast), 2013.
- [33] European Union. Regulation (EU) 2015/2120 of the European Parliament and of the Council of 25 November 2015 Laying Down Measures Concerning Open Internet Access (Recast), 2015.
- [34] European Union. Directive (EU) 2018/1972 of the European Parliament and of the Council of 11 December 2018 establishing the European Electronic Communications Code (Recast), 2018.
- [35] European Union Agency for Network and Information Security. Critical Infrastructures and Services, Internet Infrastructure: Internet Interconnections. <http://enisa.europa.eu/activities/Resilience-and-CIIP/critical-infrastructure-and-services/inter-x/>, 2010.
- [36] FCC. In the Matter of Madison River Communications, LLC and Affiliated Companies. Order, File No. EB-05-IH-0110 (DA 05-543), 2005.
- [37] FCC. In the Matter of Madison River Communications, LLC and Affiliated Companies. Consent Decree, File No. EB-05-IH-0110 (DA 05-543), 2005.
- [38] FCC. In the Matters of Appropriate Framework for Broadband Access to the Internet over Wireline Facilities et al., Policy Statement, CC Docket No. 02-33 et al. (FCC 05-151), 2005.
- [39] FCC. In the Matters of Formal Complaint of Free Press and Public Knowledge against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications. Memorandum Opinion and Order, File No. EB-08-IH-1518, WC Docket No. 07-52 (FCC 08-183), 2008.
- [40] FCC. In the Matter of Preserving the Open Internet; Broadband Industry Practices. Report and Order, GN Docket No. 09-191, WC Docket No. 07-52 (FCC 10-201), 2010.
- [41] FCC. In the Matter of Protecting and Promoting the Open Internet. Report and Order on Remand, Declaratory Ruling, and Order, GN Docket No. 14-28 (FCC 15-24), 2015.
- [42] FCC. In the Matter of Restoring Internet Freedom. Declaratory Ruling, Report and Order, and Order, WC Docket No. 17-108 (FCC 17-166), 2018.
- [43] FCC. Measuring Broadband America. <https://www.fcc.gov/general/measuring-broadband-america>, 2019.
- [44] T. Flach, N. Dukkipati, A. Terzis, B. Raghavan, N. Cardwell, Y. Cheng, A. Jain, S. Hao, E. Katz-Bassett, and R. Govindan. Reducing Web Latency: the Virtue of Gentle Aggression. In *ACM SIGCOMM*, 2013.
- [45] B. Frank, I. Poesse, Y. Lin, G. Smaragdakis, A. Feldmann, B. Maggs, J. Rake, S. Uhlig, and R. Weber. Pushing CDN-ISP Collaboration to the Limit. *ACM CCR*, 43(3), 2013.
- [46] R. Frieden. Ex Ante versus Ex Post Approaches to Network Neutrality: A Comparative Assessment. *Berkeley Technology Law J.*, 30(2), 2015.
- [47] T. Garrett, L. E. Setenareski, L. M. Peres, L. C. E. Bona, and E. P. Duarte. Monitoring Network Neutrality: A Survey on Traffic Differentiation Detection. *IEEE Commun. Surveys & Tutorials*, 20(3), 2018.
- [48] A. A. Gilroy. The Net Neutrality Debate: Access to Broadband Networks. Congressional Research Service Report, updated 15 April, 2019.
- [49] V. Giotas, G. Smaragdakis, B. Huffaker, M. Luckie, and k. claffly. Mapping Peering Interconnections at the Facility Level. In *ACM CoNEXT*, 2015.
- [50] Google. Google Video Quality Report. <https://www.google.com/get/videoqualityreport/>, 2019.
- [51] Google. Introduction to GGC. <https://peering.google.com/#/options/google-global-cache>, 2019.
- [52] S. M. Greenstein. *How the Internet Became Commercial: Innovation, Privatization, and the Birth of a New Network*. Princeton University Press, 2015.
- [53] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. SDX: A Software Defined Internet Exchange. In *ACM SIGCOMM*, 2014.
- [54] T. W. Hazlett and A. Caliskan. Natural Experiments in U.S. Broadband Regulation. *Review of Network Economics*, 7(4), 2008.
- [55] M. Honda, Y. Nishida, C. Raiciu, A. Greenhalgh, M. Handley, and H. Tokuda. Is it Still Possible to Extend TCP? In *ACM IMC*, 2011.
- [56] C-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. AlimiKondapa, N. B. Chandan, B. Sourabh, J. Jay, K. Shiyu, L. Mendelev, S. Padgett, F. Rabe, S. Ray, M. Tewari, M. Tierney, M. Zahn, J. Zolla, J. Ong, and A. Vahdat. B4 and after: managing hierarchy, partitioning, and asymmetry for availability and scale in google's software-defined WAN. In *ACM SIGCOMM*, 2018.
- [57] Internet Society. Net Neutrality Experts' Roundtable Series: Process Report, May 2019.
- [58] kc claffly, D. D. Clark, S. Bauer, and A. Dhamdhere. Policy challenges in mapping Internet interdomain congestion. In *TPRC*, 2016.
- [59] G. Knieps. *Network Economics*. Springer, 2015.
- [60] G. Knieps and V. Stocker. Price and QoS differentiation in all-IP networks. *International J. of Management and Network Economics*, 3(4), 2016.
- [61] C. Kreibich, B. Nechaev, V. Paxson, and N. Weaver. Netalyzer: Illuminating The Edge Network. In *ACM IMC*, 2010.
- [62] B. Krishnamurthy and J. Wang. On Network-Aware Clustering of Web Clients. In *ACM SIGCOMM*, 2001.
- [63] R. Krishnan, H. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao. Moving Beyond End-to-end Path Information to Optimize CDN Performance. In *ACM IMC*, 2009.
- [64] M. Kühlewind. Some updates on QUIC deployment numbers. IETF 106 MAPRG, 2019.
- [65] C. Labovitz. Internet Traffic 2009-2019. APRICOT, 2019.
- [66] C. Labovitz, S. Lelak-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *ACM SIGCOMM*, 2010.

- [67] A. Langley, A. Riddoch, A. Wilk, A. Vicente, C. Krasic, D. Zhang, F. Yang, F. Kournov, I. Swett, J. Iyengar, J. Bailey, J. Dorfman, J. Roskind, J. Kulik, P. Westin, R. Tnetti, R. Shade, R. Hamilton, V. Vasilev, W.-T. Chang, and Z. Shi. The QUIC Transport Protocol: Design and Internet-Scale Deployment. In *ACM SIGCOMM*, 2017.
- [68] W. Lehr. 5G and the Future of Broadband. The Future of the Internet – Innovation, Integration, and Sustainability. Nomos, 2019.
- [69] W. Lehr, D. D. Clark, S. Bauer, A. Berger, and P. Richter. Whither the public Internet? *J. of Information Policy*, 9, 2019.
- [70] T. Leighton. Improving Performance on the Internet. *Comm. of the ACM*, 52(2), 2009.
- [71] F. Li, A. A. Niaki, D. Choffnes, P. Gill, and A. Mislove. A Large-Scale Analysis of Deployed Traffic Differentiation Practices. In *ACM SIGCOMM*, 2019.
- [72] X. Liu, F. Dobrian, H. Milner, J. Jiang, V. Sekar, I. Stoica, and H. Zhang. A Case for a Coordinated Internet Video Control Plane. In *ACM SIGCOMM*, 2012.
- [73] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, and kc claffy. Challenges in Inferring Internet Interdomain Congestion. In *ACM IMC*, 2014.
- [74] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein. On Cooperative Settlement Between Content, Transit, and Eyeball Internet Service Providers. *IEEE/ACM Trans. Networking*, 19(3), 2011.
- [75] R. T. B. Ma and V. Misra. The Public Option: A Nonregulatory Alternative to Network Neutrality. *IEEE/ACM Trans. Networking*, 21(6), 2012.
- [76] A. Maltinsky, R. Giladi, and Y. Shavitt. On Network Neutrality Measurements. *ACM Trans. on Intel. Systems & Technology*, 8(4), 2017.
- [77] J. S. Marcus. Network Neutrality Revisited: Challenges and Responses in the EU and in the US, Study for the IMCO Committee. IP/A/IMCO/2014-02, December, 2014.
- [78] J. S. Marcus. New Network Neutrality Rules in Europe: Comparisons to those in the U.S. *Colorado Technology Law J.*, 14(2), 2016.
- [79] J. S. Marcus, D. Elixmann, and J. R. Carter. The Future of IP Interconnection: Technical, Economic, and Public Policy Aspects, Study for the European Commission. Final Report, 29 January, Bad Honnef: WIK-Consult, 2008.
- [80] C. T. Marsden. Comparative Case Studies in Implementing Net Neutrality: A Critical Analysis of Zero Rating. *SCRIPTed*, 13(1), 2016.
- [81] V. Misra. Net neutrality is all good and fine; the real problem is elsewhere. <http://www.cs.columbia.edu/2014/net-neutrality/>, 2014.
- [82] Netflix. Netflix ISP Speed Index. <https://ispspeedindex.netflix.com/>, 2019.
- [83] Netflix. Netflix Open Connect. <https://openconnect.netflix.com/>, 2019.
- [84] P. Nikolopoulos, C. Pappas, K. Argyraki, and A. Perrig. Retroactive Packet Sampling for Traffic Receipts. In *ACM SIGMETRICS*, 2019.
- [85] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS Oper. Syst. Rev.*, 44(3), 2010.
- [86] US Department of Homeland Security. Critical Infrastructure Sectors: Communications Sector. <https://www.dhs.gov/cisa/communications-sector>, 2019.
- [87] The Tor Project OONI. OONI: Open Observatory of Network Interference. <https://ooni.org/about/>, 2019.
- [88] C. Paasch and O. Bonaventure. Multipath TCP. *Comm. of the ACM*, 57(4), 2014.
- [89] C. Pappas, K. Argyraki, S. Bechtold, and A. Perrig. Transparency Instead of Neutrality. In *HotNets*, 2015.
- [90] M. K. Powell. Preserving internet freedom: Guiding principles for the industry. *J. on Telecom. & High Technology Law*, 3(1), 2004.
- [91] E. Pujol, I. Poese, J. Zerwas, G. Smaragdakis, and A. Feldmann. Steering Hyper-Giants' Traffic at Scale. In *ACM CoNEXT*, 2019.
- [92] E. Pujol, P. Richter, B. Chandrasekaran, G. Smaragdakis, A. Feldmann, B. Maggs, and K. C. Ng. Back-Office Web Traffic on The Internet. In *ACM IMC*, 2014.
- [93] M. A. Sanchez, J. S. Otto, Z. S. Bischof, D. R. Choffnes, F. E. Bustamante, B. Krishnamurthy, and W. Willinger. Dasu: Pushing Experiments to the Internet's Edge. In *NSDI*, 2013.
- [94] B. Schlinker, H. Kim, T. Cui, E. Katz-Bassett, H. V. Madhyastha, I. Cunha, J. Quinn, S. Hasan, P. Lapukhov, and H. Zeng. Engineering Egress with Edge Fabric: Steering Oceans of Content to the World. In *ACM SIGCOMM*, 2017.
- [95] H. Schulzrinne. Network Neutrality Is About Money, Not Packets. *IEEE Internet Computing*, 2018.
- [96] R. Sen, S. Ahmad, A. Phokeer, Z. A. Farooq, I. A. Qazi, D. Choffnes, and K. P. Gummadi. Inside the Walled Garden: Deconstructing Facebook's Free Basics Program. *ACM CCR*, 47(5), 2018.
- [97] A. Shukla and V. Stocker. Navigating the Landscape of Programmable Networks: Looking beyond the Regulatory Status Quo. In *TPRC* 47, 2019.
- [98] R. Singh, M. Ghobadi, K.-T. Foerster, M. Filer, and P. Gill. RADWAN: Rate Adaptive Wide Area Network. In *ACM SIGCOMM*, 2018.
- [99] A. Singla, B. Chandrasekaran, B. Godfrey, and B. M. Maggs. The Internet at the Speed of Light. In *HotNets*, 2014.
- [100] V. Stocker. Interconnection and Capacity Allocation for All-IP Networks: Walled Gardens or Full Integration? In *TPRC* 43, 2015.
- [101] V. Stocker. *Innovative Capacity Allocations for All-IP Networks: A Network Economic Analysis of Evolution and Competition in the Internet Ecosystem*. Nomos, 2020.
- [102] V. Stocker and G. Knieps. Network Neutrality Through the Lens of Network Economics. *Review of Network Economics*, 17(3), 2019.
- [103] V. Stocker, G. Smaragdakis, W. Lehr, and S. Bauer. The Growing Complexity of Content Delivery Networks: Challenges and Implications for the Internet Ecosystem. *Telecommunications Policy*, 41(10), 2017.
- [104] V. Stocker and J. Whalley. Speed isn't everything: A multi-criteria analysis of the broadband consumer experience in the UK. *Telecommunications Policy*, 42(1), 2018.
- [105] R. Stoenescu, M. Popovici, V. Olteanu, J. Martins, R. Bifulco, F. Huici, M. Ahmed, G. Smaragdakis, M. Handley, and C. Raiciu. In-Net: Enabling In-Network Processing for the Masses. In *EuroSys*, 2015.
- [106] F. Streibelt, J. Boettger, N. Chatzis, G. Smaragdakis, and A. Feldmann. Exploring EDNS-Client-Subnet Adopters in your Free Time. In *ACM IMC*, 2013.
- [107] S. Sundaresan, N. Feamster, and R. Teixeira. Locating Throughput Bottlenecks in Home Networks. In *ACM SIGCOMM*, 2014.
- [108] P. Svensson. Comcast Blocks Some Internet Traffic. Washington Post, 19 October, available: <http://www.washingtonpost.com/wp-dyn/content/article/2007/10/19/AR2007101900842.html>, 2007.
- [109] M. Trevisan, D. Giordano, I. Drago, M. Mellia, and M. Munafò. Five Years at the Edge: Watching Internet from the ISP Network. In *ACM CoNEXT*, 2018.
- [110] Mozilla Corporation v. FCC. On Petitions for Review of an Order of the Federal Communications Commission. District of Columbia Circuit, No. 18-1051, 2019.
- [111] N. Vallina-Rodriguez, S. Sundaresan, C. Kreibich, and V. Paxson. Header Enrichment or ISP Enrichment? Emerging Privacy Threats in Mobile Networks. In *ACM HotMiddlebox*, 2015.
- [112] M. Vanberg. *Competition and Cooperation Among Internet Service Providers—A Network Economic Analysis*. Nomos, 2009.
- [113] I. Vogelsang. Net Neutrality Regulation: Much Ado about Nothing? *Review of Network Economics*, 2019.
- [114] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger. Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN. In *ACM SIGCOMM*, 2018.
- [115] T. Wu. A Proposal for Network Neutrality. <http://www.timwu.org/OriginalNNProposal.pdf>, 2002.
- [116] T. Wu. Network Neutrality, Broadband Discrimination. *J. on Telecommunication and High Technology Law*, 2, 2003.
- [117] T. Wu and C. Yoo. Keeping the Internet Neutral?: Tim Wu and Christopher Yoo Debate. *Faculty Scholarship at Penn Law*, 779, 2007.
- [118] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat. Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. In *ACM SIGCOMM*, 2017.
- [119] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger. How Cloud Traffic Goes Hiding: A Study of Amazon's Peering Fabric. In *ACM IMC*, 2019.
- [120] Y. Yiakoumis, S. Katti, and N. McKeown. Neutral Net Neutrality. In *ACM SIGCOMM*, 2016.
- [121] C. S. Yoo. Network neutrality and the economics of congestion. *Georgetown Law J.*, 94(6), 2005.
- [122] C. S. Yoo and J. Lambert. 5G and Net Neutrality. The Future of the Internet – Innovation, Integration, and Sustainability. Nomos, 2019.
- [123] Z. Zhang, O. Mara, and K. Argyraki. Network Neutrality Inference. In *ACM SIGCOMM*, 2014.