REDE
DE PESQUISA
EM GOVERNANÇA
DA INTERNET

Anais do Encontro Virtual da Rede de Pesquisa em Governança da Internet

Novembro de 2022

V ENCONTRO DA REDE DE PESQUISA EM GOVERNANÇA DA INTERNET - REDE 2022

FICHA TÉCNICA

ANAIS DA REDE DE PESQUISA EM GOVERNANÇA DA INTERNET - VOL. 5, ISSN 2675-1690

Editores
Diego Vicentin
Maria Vitoria Pereira de Jesus

Autores
Adriana Veloso
Cristina Carvalho
Hermann Bergmann
Javier de la Cruz Martínez
Julio Marinho
Maiara Stefany
Manuel Alberto
Mariana Gomes
Marina Gonçalves
Paula Guedes
Raimundo Miguel
Rebeca Souza

COMITÊ ORGANIZADOR

Núcleo de Coordenação da Rede de Pesquisa em Governança da Internet
Alexandre Arns Gonzales
Carolina Batista Israel
Diego Vicentin
Fernanda R. Rosa
Gustavo Ramos Rodrigues
Hemanuel Jhosé Alves Veras
Kimberly de Aguiar Anastácio
Maria Vitoria Pereira de Jesus
Nathan Paschoalini Ribeiro Batista

Comitê Científico
Carolina Batista Israel
Daniela Araújo
Diego Vicentin
Fernanda R. Rosa
Flávio Rech Wagner
Gills Vilar-Lopes
Jean Santos
Leonardo Ribeiro da Cruz
Maria Mónica Arroyo
Marta Mourão Kanashiro
Raquel Gatto
Rodolfo Avelino
Rosemary Segurado

Moderadores
Ana Claudia Farranha
Bianca Kremer
Carolina Batista Israel
Diego Vicentin
Fernanda R. Rosa
Flávio Rech Wagner
Jess Reia
Leonardo Ribeiro da Cruz
Marisa von Bulow

Debatedores
Adriana Veloso
Alessandra Ogassavara
Cristina Carvalho
Hermann Bergmann
Javier de la Cruz Martínez
Julio Marinho
Luis Henrique Antunes
Maiara Stefany
Manuel Alberto
Mariana Gomes
Marina Gonçalves
Paula Guedes
Raimundo Miguel
Rebeca Souza
Rosemary Segurado

O V Encontro da Rede de Pesquisa em Governança da Internet e a publicação dos Anais resultantes deste encontro receberam apoio de:

Internet Society - Capítulo Brasil

Programa de Pós-Graduação em Ciências Sociais do Instituto de Filosofia e Ciências Humanas (IFCH) da Universidade Estadual de Campinas (UNICAMP) por meio de recursos da Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) - convênio n. 0524/2021.

Internet Society Capítulo Brasil

CAPES

ON THE COEXISTENCE OF 5G NETWORK SLICING AND NET NEUTRALITY

HERMANN BERGMANN GARCIA E SILVA

MANUEL ALBERTO PEREIRA RICARDO

SUMÁRIO

ON THE COEXISTENCE OF 5G NETWORK SLICING AND NET NEUTRALITY

Hermann Bergmann Garcia e Silva
Manuel Alberto Pereira Ricardo

ABSTRACT
Since John Perry Barlow's cyberspace declaration of independence more than two decades ago, many transformations have occurred in the digital ecosystem. The utopian vision of free space was replaced by a reality in which informational flows became conditioned by the Internet architecture, its communication protocols, and the intervention of network operators, which influence the users' behavior and autonomy. Currently, the fifth generation of mobile communications systems (5G) standardizes a programmable network capable of supporting the exponential growth of Internet traffic and handling a heterogeneity of use cases. Central to the 5G technology, network slicing introduces new traffic differentiation paradigms that segment users, applications, and services into customized logical domains having dedicated radio and computational resources. This logic potentially conflicts with the net neutrality principle, which aims to ensure that all Internet communications should be treated in the same manner, without discrimination. This paper addresses the contradictory nature of these two concepts and proposes a strategy based on standardized functions already available in the 5G core network to increase network transparency and subsidize the regulatory monitoring of Internet traffic management practices, which are essential elements to enable the coexistence of network slicing with net neutrality.

KEYWORDS
Internet governance; Telecommunications policy; Net neutrality; 5G; Network slicing.

RESUMO
Desde a declaração de independência do ciberespaço de John Perry Barlow há mais de duas décadas, muitas transformações ocorreram no ecossistema digital. A visão utópica de um espaço livre foi substituída por uma realidade em que os fluxos informacionais se tornaram condicionados pela arquitetura da Internet, seus protocolos de comunicação e pela intervenção dos operadores de rede, que influenciam o com-

portamento e a autonomia dos usuários. Atualmente, a quinta geração de sistemas de comunicações móveis (5G) padroniza uma rede programável capaz de suportar o crescimento exponencial do tráfego de Internet e lidar com uma heterogeneidade de casos de uso. Central à tecnologia 5G, o fatiamento de rede introduz novos paradigmas de diferenciação de tráfego que segmentam usuários, aplicações e serviços em domínios lógicos personalizados com recursos computacionais e de rádio dedicados. Esta lógica entra em potencial conflito com o princípio da neutralidade de rede, que visa garantir que todas as comunicações da Internet sejam tratadas da mesma forma, sem discriminação. Este artigo aborda a natureza contraditória desses dois conceitos e propõe uma estratégia baseada em funções padronizadas já disponíveis no núcleo da rede 5G para ampliar a transparência da rede e subsidiar o monitoramento regulatório das práticas de gerenciamento do tráfego de Internet, que são elementos essenciais para possibilitar a coexistência do fatiamento de rede com a neutralidade da rede.

PALAVRAS-CHAVE
Governança da Internet; Política de telecomunicações; Neutralidade de rede; 5G; Fatiamento de rede.

INTRODUCTION

The evolution of telecommunications systems, the expansion of the Internet on a global scale, and the proliferation of Web 2.0 technologies have transformed the way humans access, use, and disseminate information in contemporary society, becoming both producers and consumers of digital content (prosumers).

The magnitude of the Internet can be scaled by the growth in the number of users. According to the International Telecommunications Union (ITU) estimates, in 2021, approximately 4.9 billion people were using the Internet, representing 63 percent of the world's population. The first year of the COVID-19 pandemic, 2020, evidenced the most significant increase in Internet users in a decade. The number raised 10.2 percent due to the restrictions on working, learning, accessing essential services and living with social distance (ITU, 2021).

The total number of connected devices is expected to surpass 150 billion by 2025, increasing the data traffic from 33 zettabytes in 2018 to 175 zettabytes in 2025, most of which will be real-time traffic, such as video streaming. Every connected person is expected to have over 4.900 daily digital interactions by 2025, representing one digital interaction every eighteen seconds (RYDNING; REINSEL; GANTZ, 2018).

Obiodu and Sastry (2020) consider that the remarkable expansion of the Internet is associated with the various technical standards developed in the past fifty years that

have steadily improved the quality and reliability of data networks. Lemley and Lessig (2001) argue that Internet's success can be attributed to the end-to-end principle of its architecture. According to this principle structured by the innovative work of Saltzer, Reed, and Clark (1984), the technical elements of a communications system should be implanted at its edges to ensure its simplicity, being the core of the network in charge of transmitting the information in the most efficient way possible. Cerf (2006, p. 9) asserts that the foundations of Internet success "can be traced to a few simple network principles - end-to-end design, layered architecture, and open standards - which together give consumers choice and control over their online activities."

However, network traffic management policies have changed significantly in the last two decades. The prospect of a best-effort-based routing has been replaced by a growing need for enhanced quality of service (QoS) to sustain the exponential growth in Internet traffic due to bandwidth-hungry multimedia applications and the widespread adoption of intelligent devices, which have led Internet service providers (ISPs) to adopt strategies to differentiate the traffic on the Internet.

Molavi Kakhki *et al.* (2015) argue that traffic differentiation is related to managing scarce network resources to provide better or worse performance to certain classes of Internet traffic. Finnie (2009) emphasizes that ISPs' traffic management technologies have improved to allow fine-grained distinction among the digital content transported over IP networks, enabling discrimination to reach the end-user level. This is one of the reasons why Jordan and Ghost (2010) assert that ISPs' traffic management practices have become an issue of public concern.

Currently, the fifth generation of mobile communications systems (5G) establishes a new programmable network standard which makes it possible to respond quickly to changes in market structure, business models, and users' needs. 5G is not limited to an incremental improvement of previous generations; it constitutes a technological leap that enables the creation of virtual networks capable of supporting QoS-demanding services and applications.

This versatile architecture points in a direction that differs from the original Internet design philosophy, which was founded on the end-to-end principle and the best-effort service. 5G network slicing adopts new traffic differentiation paradigms that reserve critical resources of the physical infrastructure, including radio and computational resources, for certain types of traffic, creating customized logical domains in potential

conflict with the isonomic and non-discriminatory treatment established by a net neutrality regime. In this context, feasible alternatives should be found to allow network slicing and net neutrality to coexist.

The Body of European Regulators for Electronic Communications (BEREC) indicates that transparency is a fundamental pre-condition for achieving net neutrality (BEREC, 2011), notably to evaluate how Internet traffic management practices employed by broadband ISPs affect end users` access to content, applications, and services. Papas *et al.* (2015) take this approach further and consider that rather than trying to detect net neutrality violations, the focus should be on increasing network transparency so that regulators can use qualified information for a holistic analysis of the ISP behavior. In this way, the Transatlantic Consumer Dialogue (TACD), a network of consumer organizations that develop policy recommendations for the US government and the European Union, endorses that regulators should continuously monitor networks to evaluate whether ISPs and network providers discriminate between content, services, applications, or devices on their networks (TACD, 2015).

In this paper, we discuss the contradictory nature of network slicing and net neutrality and show that a transparency-enabling function, allowing the controlled exposure of information to the regulatory authority, could be defined by taking advantage of standardized functions already available in the 5G core network.

After this introduction, the paper is structured into five sections. The first section addresses the net neutrality principle. The second section presents the 5G architecture and defines the network slicing concept. The third section discusses the potential conflict between network slicing and net neutrality. The fourth section proposes a strategy for network transparency and regulatory monitoring in 5G networks. Finally, the last section presents concluding remarks.

NET NEUTRALITY

One of the most controversial topics in Internet governance is the concept and extent of the net neutrality principle, given the multiple and often conflicting interests involved in the Internet arena, which include aspects of technical, economic, political, and human rights nature. Depending on the analysis perspective, different interpretations make its definition an important research object for understanding the dynamics of contemporary relations in the digital ecosystem.

Net neutrality reignites the discussion about the obligations of companies providing infrastructure services, which are essential to the economy, and how these companies relate to the public interest. In the 19th century, means of navigation and railways had a central role in the debate; in the 20th century, electricity and telephony were the protagonists; in the 21st century, the Internet assumed the focus of the discussion (LEE; WU, 2009).

The debate on net neutrality originated in the United States in the early 2000s. It was guided by the establishment of limits for network operators regarding their interference in Internet traffic, free flow of information, and users' freedom of choice. Tim Wu introduced the term in 2003, defining it as a principle in which all Internet traffic should be treated equally, with network intermediaries refraining from making decisions in the place of end-users (WU, 2003). That is to say, the agents responsible for the telecommunications infrastructure should not discriminate against any traffic, whether due to the nature of its content, the equipment's location, the software or hardware used, or the origin or destination of packets transmitted.

Barratt and Shade (2007) incorporate into the idea of net neutrality the perspective that the Internet does not have a centralized control mechanism, which implies that network operators should not interfere with the content transported through their networks. Belli (2016) argues that the rationality of net neutrality is to preserve the open and decentralized character of the Internet architecture, intending to empower end-users and protect their fundamental rights. In this sense, Kaye (2018) affirms that the ISPs and network operators should not be the ones deciding what kind of information or content individuals should access online, concluding that net neutrality and democracy are inextricably linked to one another.

Andreoletti *et al.* (2018) consider that net neutrality represents an ideal goal of a platform that enables open and non-discriminatory access to information in an era where traffic differentiation is mandatory. The authors highlight that the net neutrality principle should keep track of the transformations on the Internet and not be tied to the notion of the network as a mere dumb pipe. Therefore, a balance of trade-offs between net neutrality and QoS should be assumed, allowing Internet traffic differentiation through categorizing traffic into classes, as long as it is not discriminatory.

The battle over net neutrality is essentially a dispute over the conception of a regulatory framework that would protect the Internet from unreasonable network mana-

gement practices, such as but not limited to blocking, throttling, and paid prioritization (STOCKER; SMARAGDAKI; LEHR, 2020).

Blocking is a practice that prevents the user from accessing certain lawful content, application, or service made available on the Internet, which happens through software or hardware located at a specific control point of data traffic in the network (HALL *et al.*, 2022). It is not a recent practice since it started from the commercial use of the Internet in the 1990s to prevent the proliferation of spam through e-mail (CALLANAN *et al.*, 2009). Blocking methods can be differentiated according to the amplitude at which they are implemented. From a narrower perspective, the end-user can establish blocking rules on the personal computer level; as the range is expanded, other techniques can be employed by a company, an ISP, or even the State itself at the national level.

Throttling is the intentional degradation or delay of Internet traffic that reduces the transmission rate in the telecommunications network. Originally conceived as a traffic differentiation measure to handle network congestion, throttling represents a limitation on users´ informational access when it ceases to be agnostic and is directed to a specific content provider, a particular class of applications, or users with a given data consumption profile (e.g., heavy users). Garret *et al.* (2018) highlight that the motivation of ISPs for this differentiated treatment is associated with a demand for competitive advantage or an increase in profitability by using the network.

Paid prioritization consists of a practice that provides preferential treatment to certain types of data packets through financial compensation. In essence, it is related to creating fast lanes in the network infrastructure and instituting a two-tiered Internet through traffic management techniques such as resource reservation, traffic shaping, or packet scheduling schemes. According to the Center for Democracy and Technology (CDT), in a letter addressed to the FCC, paid prioritization can be defined as the ISP practice of charging a fee to deliver the traffic of CAPs in an enhanced fashion over subscribers' last-mile facilities (COOPER; MORRIS; SOHN, 2010). ISPs employ three arguments to defend the practice of paid prioritization on the Internet. The first refers to the need to manage data traffic to avoid congestion and enable efficient network use. As the second argument, ISPs argue that, without traffic differentiation, they would not have incentives to expand network capacity and provide a better quality of service, which could hinder the development of the Internet infrastructure. The third argument is related to the idea that the market remuneration model of the Internet is not balanced since the profitability

of ISPs is maintained, as a rule, through revenue from the use of the network, while content and application providers receive both for the provision of services and online advertising (MA; WANG; CHIU, 2017).

Blocking, throttling, and paid prioritization interfere with the free flow of information and compromise the users' broad and unrestricted access to content in the digital ecosystem. Under the aegis of a regime in which net neutrality prevails, ISP discriminatory practices can only occur in exceptional circumstances, such as preserving the integrity of the network or controlling its congestion and should be done in a reasonable manner for the benefit of the network and its users as a whole.

The arguments in favor of net neutrality are based on preserving an open and decentralized communications platform that supports the exercise of fundamental human rights, fostering at the same time an innovative and competitive environment (LESSIG, 2006; VAN SCHEWICK, 2012; BELLI, 2016). On the other hand, the discouragement of investment in network infrastructure, the need for the efficient use of a scarce resource (e.g., radio spectrum), and the provision of new applications and services requiring customized QoS are some of the reasons against the net neutrality principle (YOO, 2005; SIDAK, 2006; YOO, 2010).

The absence of net neutrality facilitates the practice of traffic discrimination by network operators through commercial agreements with content and application providers, which impacts the free flow of information and puts at risk the freedom of users to choose the content they want to explore, create or disseminate online (MARQUES; PINHEIRO, 2014).

In this sense, it is possible to assert that two distinct perspectives of network architecture polarize the net neutrality debate. The first perspective seeks to preserve the foundations of the original Internet design model and maintain the decisions in the application layer, providing greater autonomy to end-users. The second perspective considers that the Internet architecture needs to evolve and support new remuneration models, enabling traffic differentiation for specific services, applications, or customers, prioritizing data traffic through fast lanes, and implementing charges according to the user profile.

5G NETWORK

In the ITU´s vision, materialized through Recommendation M.2083-0 (ITU, 2015b), also known as IMT-2020, the 5G network must be able to support at least

three usage scenarios: enhanced mobile broadband (eMBB), ultra-reliable and low-latency communications (URLLC), and massive machine type communications (mMTC). The eMBB scenario should be able to provide an increase of 10 to 100 times in data transmission rates to enable a variety of high-resolution multimedia services and applications. The URLCC should support a set of critical and time-sensitive services and applications, including wireless control of industrial processes, cloud gaming, and autonomous vehicle, that require an end-to-end latency of less than 1 ms and an aggregate reliability and availability service of 99.999% for packet transmission. The mMTC should provide the capacity to connect a large number of IoT devices (1.000.000/km²) with low cost and longer battery life (BARAKABITZE *et al.*, 2020).

In order to meet these specific technical requirements, the 5G architecture logically partitions the physical infrastructure into slices through the virtualization of access, transport, and core networks, operating through the support of software-defined networking (SDN) and network function virtualization (NFV) technologies.

The term SDN was originally introduced at Stanford University to represent ideas and studies that were being developed on the OpenFlow protocol and refers to a network architecture in which the data plan is decoupled and managed by the control plan (KREUTZ *et al.*, 2014). The ITU defines SDN as a set of techniques that enable programming, orchestrating, controlling, and directly managing network resources through software, allowing the provision of services in a deterministic, dynamic, and scalable manner (ITU, 2014). The SDN approach provides logical control of the elements of a communications network responsible for forwarding and routing datagrams. It is an emerging technology where the primary concept is to remove the intelligence from network devices and manage all their functionality through a centralized controller (SDN controller). In essence, SDN materializes the vision of a programmable network, which enables dynamic control and management through homogeneous interfaces (HALEPLIDIS *et al.*, 2015). Xia *et al.* (2014) highlight that the decoupling of control and data plans not only allows a simpler programmable environment but also provides greater flexibility for software to define the network behavior.

While SDN aims to establish centralized control of the network, NFV focus on virtualizing network functions and deploying them into generic devices. Virtualization was a widely studied theme in the 1960s and 1970s. It emerged with the purpose of creating a logical layer between the hardware and software so that more than one

operating system could run in parallel on the same equipment in a homogeneous and flexible manner (PEREIRA; ALBERTI, 2013). According to the European Telecommunications Standards Institute (ETSI), NFV can be conceptualized as the principle of separating network functions from the equipment in which they operate through virtual hardware abstraction (ESTI, 2021).

Mijumbi *et al.* (2016) consider that SDN and NFV technologies are complementary and can provide an integrated solution of open-source software and non-proprietary equipment, enabling the network programming capability necessary to provide services adapted to the performance requirements of heterogeneous use cases. These are inputs for customizing virtual networks through one of the main 5G features: network slicing.

The first conception of network slicing was introduced in 2002 in a research effort to develop new protocols, services, and applications that could optimize the capabilities of the Internet, from the idea of sharing network resources through virtualization. An example of these initiatives was the Planet-Lab project (NAKAO *et al.*, 2017).

The Internet Engineering Task Force (IETF) establishes the concept of network slicing as an approach to provide an end-to-end independent logical network, from the user equipment to multiple mobile applications, where each network slice has its service level agreement (SLA), with the arrangement of strict technical parameters according to the 5G use case (ROUKI *et al.*, 2020).

To ITU, network slicing is a management mechanism that enables the deployment of multiple logical and independent shared network resources and service functions on a common infrastructure platform (ITU, 2020).

Nakao *et al.* (2017) highlight that the main characteristics of network slicing are the individualization of network resources and the programmability of these resources that guarantee an end-to-end QoS for the demands of the telecommunications infrastructure, which is why Zhu *et al.* (2019) consider that the 5G slice can provide the functionalities of a complete network, from the multiple technologies of the access network to the core network functions.

Creating different virtual instances makes the users' experience transparent as if they were enjoying a dedicated physical infrastructure. Chatras, Kwong, and Bihannic (2017, p. 219) say that this customization of the 5G network allows recognizing network slicing as "a new sophisticated form of Virtual Private Network (VPN) technology."

Therefore, each slice has the ability to refine and isolate performance indicators regardless of the congestion and performance level of other virtual instances, which means that network slicing can implement the differentiation of traffic at the user level, individually, including the reservation of network resources. Consequently, a slice of the 5G network could be considered an autonomous logical network with its own architecture, protocols, and provisioning rules.

The 5G network architecture (figure 1) is composed of a control plane and a data plane, which in 5G is referred to as the user plane. The control plane is responsible for coordinating the overall dynamics of the network, including routing functions, the decision of packet forwarding, and the application of QoS. The user plane is in charge of forwarding data packets from source to destination according to the control plane's established policies.

Figure 1: 5G architecture

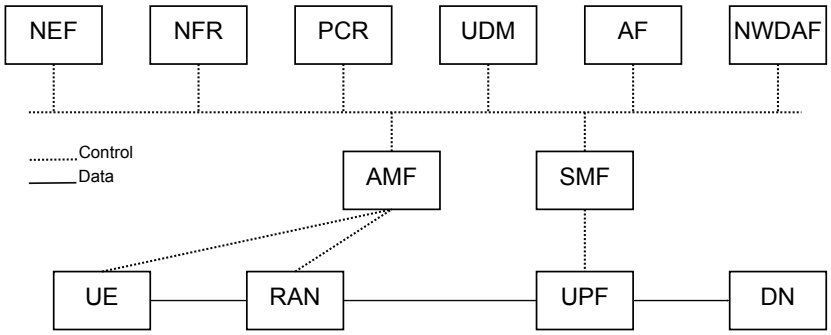

Source: elaborated by the authors.

The user plane function (UPF) is an essential component of the 5G core, providing the interconnection interface between the radio access network (RAN) and the data network (DN). It is also responsible for packet routing and forwarding, packet inspection, buffering of downlink data, QoS handling, and policy rule enforcement. The RAN is in charge of providing radio access to the user equipment (UE) and is composed of a node connected to the core network (gNodeB). Beyond managing radio resources and mobility during the session, its functions include QoS flow control and data routing toward the UPF. The session management function (SMF) is responsible for controlling the UPF and managing the entire lifecycle of a protocol data unit (PDU) session . The access and mobility management function (AMF) is in charge of access authori-

zation and authentication, UE location, and connection management. The other control plane functions in the 5G core are virtualized and executed as web services.

The network data analytics function (NWDAF) is of particular interest to this study. As part of the 5G core architecture, the NWDAF is a 3GPP standardized function introduced in release 15 and improved in releases 16 and 17 (3GPP, 2022a), which can collect data in a centralized manner from other 5G network functions (NFs) and operations, administration, and management (OAM) systems, providing statistical analysis and prediction capabilities to oversee the user experience and network performance.

NWDAF will allow the automation of the network with minimum human intervention, taking advantage of artificial intelligence and machine learning technologies to cope with the increasing complexity of communications systems (GHOSH *et al.*, 2019).

NWDAF has a consumer-producer architecture that is established through standard interfaces to obtain periodic information on the network status and resource usage on a per-slice basis, gathering data to provide statistics of past events or predictive insights by request or subscription from other NFs.

The data analytics function of the 5G system offers five services, detailed in the technical specification TS 23.288 (3GPP, 2022b), which enable the consumer NF to request and receive different types of analytics information, including network performance, QoS sustainability, user data congestion, and observed service experience derived from a particular application, set of applications, individual user equipment, or group of subscriber's mobile devices. Services provided by NWDAF can also be consumed by external entities through the network exposure function (NEF) that enables communication with the 5G core.

RELATED WORK

The potential conflict between net neutrality and network slicing is one of the main policy challenges for implementing 5G networks (ALEXIADIS; SHORTALL, 2016). This cutting-edge discussion is an evolving area of research that has attracted increased interest in the scientific community.

Frias and Martínez (2018) analyze the tensions between the current network neutrality regime in the European Union (EU) and the combination of technologies that characterize 5G, particularly network slicing. The authors conclude that much of the potential envisaged through 5G is not feasible under a vision that imposes a paradigm

of neutrality, mainly because 5G networks provide anything-as-a-service through the offer of customized network capabilities at the infrastructure layer, which conflicts with the net neutrality logic that assumes that the concept of service is attached to the application layer.

Alén-Savikko (2019) argues that discrimination rather than equality runs in the core of 5G networks, which indicates that net neutrality and network slicing are contradictory concepts. Although the conflict seems evident at first glance, the compatibility in their relationship may rest in the employment of the end-to-end principle in 5G architecture. The author acknowledges that assessing net neutrality in the 5G era should not be made exclusively with a technological approach. The Internet has significant economic, legal, socio-cultural, and political implications, demanding a holistic strategy to include all dimensions of this complex issue in the discussion. Therefore, technological determinism must be avoided, and net neutrality as a design principle should be perceived as a matter of choice in preventing the centralization tendency of the Internet.

To Smirnova *et al.* (2019), tailoring network slices to specific vertical applications or services with different performance requirements cannot be executed using a best-effort logic, which raises compatibility concerns with net neutrality regulation. The authors propose a measurement methodology for assessing traffic management practices performed via network slices that can be used to monitor the compliance of customized QoS parameters with the non-discrimination policy established in the EU. This measurement methodology utilizes the national Internet exchange points as a standard reference where quality indicators can be obtained and compared.

Kantola (2019) analyses the impact of the controversy between the new paradigm of virtualized networking and the EU net neutrality regulation on network providers. The author considers that the normative requirement of equal treatment, except for specialized services, is creating uncertainties for the 5G investment and an unlevel playing field for the European cloud service providers, especially with regard to edge computing, concluding that the net neutrality regulation established in 2015 is out of date since the best-effort service model will not be able to integrate the digital and the physical worlds effectively.

Given this scenario, it is possible to identify research opportunities for pursuing a convergent pathway for these two concepts. We envisage that the 5G data analytics function could play a crucial role in amplifying network transparency, reducing informa-

tion asymmetry, and improving the regulatory monitoring of ISPs' discriminatory practices, balancing the contradictory nature of network slicing and net neutrality.

MANAGING THE POTENTIAL CONFLICT

The detection of traffic differentiation by ISPs has drawn significant attention for over ten years. One of the first studies was conducted by Beverly, Bauer, and Berger (2007), which measured the blocking of traffic belonging to a particular application or class of applications in their assigned TCP or UDP port. Zhang, Mao, and Zhang (2009) developed NetPolice as a system to verify traffic differentiation in backbone ISPs by taking loss measurements from end hosts. Lu *et al.* (2009) presented a user-level tool named POPI to detect packet forwarding prioritization in routers through an end-to-end measurement. Dischinger *et al.* (2010) proposed Glasnost to identify whether the ISP was performing application-specific traffic shaping of Internet users. Through DiffProbe, Kanuparthy and Dovrolis (2010) focus on delay discrimination and loss discrimination by comparing latency and packet loss between exposed and controlled traffic. Li *et al.* (2019) address the impact of content-based traffic differentiation practices deployed in operational mobile networks by conducting a large-scale study that detected fixed-rate bandwidth limits as the most observed conduct. Khandkar and Hanawal (2021) propose a measurement framework named FairNet to detect any deliberate traffic discrimination over the Internet using audio and video media streaming services.

Therefore, it is possible to systemize methods to identify traffic differentiation into three categories: active probing, passive measurement, and through a VPN connection. Active probing refers to the technique of generating packet flows into the network to observe its response. Passive measurement is based on observation and collection of information on undisturbed and unmodified packet flows of interest that pass through one or more monitoring points. The third method is using a VPN connection to record and replay the traffic generated by applications and compare the performance experienced with and without the VPN.

Evaluating traffic differentiation in a model of dynamically assembled logical networks is a complex task since it involves a variety of possible configurations that influence how the network is managed. In this way, it would be unrealistic to assume the prospect of a 5G network that treats in the same manner slices customized for time-critical communications (e.g., industrial automation, autonomous driving, remote surgery),

slices customized to maximize data rate (e.g., high-definition video streaming), and slices customized to transport the information of a massive number of sensors (e.g., smart agriculture), which imposes a pragmatic implementation of the net neutrality principle.

We propose that the isonomic and non-discriminatory treatment of Internet traffic in 5G networks should be considered on a per-slice basis and for slices that are or will be made publicly accessible to users or content and applications providers (CAPs). Therefore, net neutrality should be assessed per slice and not by comparing different slices.

To that extent, the measures that differentiate data traffic should be evaluated within the boundaries of each 5G public slice. This strategy enables the analysis of traffic differentiation practices against a specific application, content, or user, providing the information needed to determine if some regulatory action is required. For private slices, that is, those created in a specific geographic area for the usage of a private entity, we assume they are exempt from net neutrality regulation and thus out of the scope of this paper.

The assessment of net neutrality can be realized by taking advantage of the 5G major usage scenarios: eMBB, URLLC, and mMTC, as illustrated in figure 2. Each slice type is expected to value performance requirements differently: for an eMBB public slice, the throughput provided to traffic flows is relevant since it is designed for multimedia applications; for a URLLC public slice, low latency and high reliability are the key requirements; for an mMTC public slice, the connection density is the most important metric.

Figure 2: Net neutrality assessment in 5G

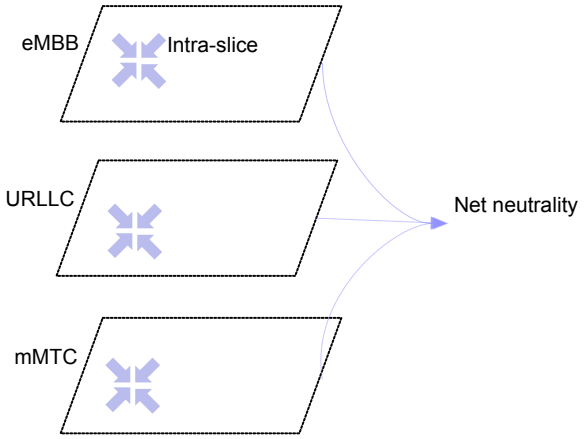

Source: elaborated by the authors.

Therefore, the most relevant performance parameters of each 5G public slice constitute the inputs required to evaluate whether measures applied by mobile network operators could characterize a discriminatory practice.

For instance, in an eMBB slice, NWDAF can provide statistics related to the level of congestion or loss experienced by data packets in the 5G user plane that can affect applications demanding high data rates. In a URLCC scenario, NWDAF can provide delay statistics related to the communications between the user equipment and time--sensitive application servers.

In this manner, the regulator can request the mobile network operator to collect information regarding the relevant metrics for each 5G public slice. In this way, it is possible to establish a regulatory monitoring procedure based on NWDAF analytics outputs, as shown in figure 3. In this figure, the NEF is in charge of providing secure exposure to 5G data analytics by request or subscription of an external party.

Figure 3: Regulatory monitoring using the NWDAF

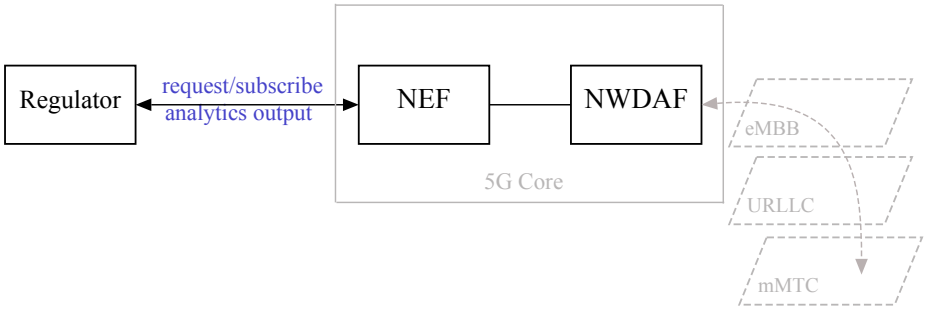

Source: elaborated by the authors.

Two approaches can be used to monitor the differentiation of traffic in 5G slices. The first approach is preventive and performs monitoring as a regular and continuous activity, searching for patterns that may indicate net neutrality violations and producing data to increase transparency on traffic management policies. The second is a reactive approach that initiates the monitoring on a case-by-case basis, triggered when a stakeholder reports a complaint.

Nevertheless, pinpointing the root cause of traffic differentiation is a complex task that involves characterizing multiples variables related to congestion levels experienced by the service, legitimate traffic differentiation (e.g., related to specialized services, illegal content, or cyber-attacks), and the emergence of content delivery networks (CDNs).

A promising strategy to address traffic differentiation in 5G networks consists in comparing data packets from different flows sharing common QoS characteristics and evaluating if they received similar network performance. This comparative assessment requires the selection of a target content and applications provider (tCAP) and a baseline content and applications provider (bCAP).

The analysis should focus on packet loss and delay statistical distributions, which are metrics that are directly impacted by traffic differentiation mechanisms. For instance, packet loss can indicate a variety of problems in cellular networks, including interference, inadequate signal strength, hardware failure, or network congestion, but it could also provide visibility to the effects of traffic degradation related to different packet scheduling and buffer management policies. Packet delay depends on variables such as packet length, the distance separating the communication peers, the number of nodes that data packets need to travel through, and the network congestion level. It could also provide important information in evaluating the dynamics of internal network queuing delays that could lead to traffic discrimination.

The proposed strategy consists in evaluating traffic differentiation by observing and collecting information on packet loss and packet delay measurements along the same network section. This strategy demands monitoring points in different 5G network boundaries. The NWDAF can obtain these measurements in the UPF and RAN for both the downlink and uplink, isolating the effects of user mobility and transmission over the air interface, as illustrated in figure 4.

Figure 4: Comparing *t*CAP and *b*CAP

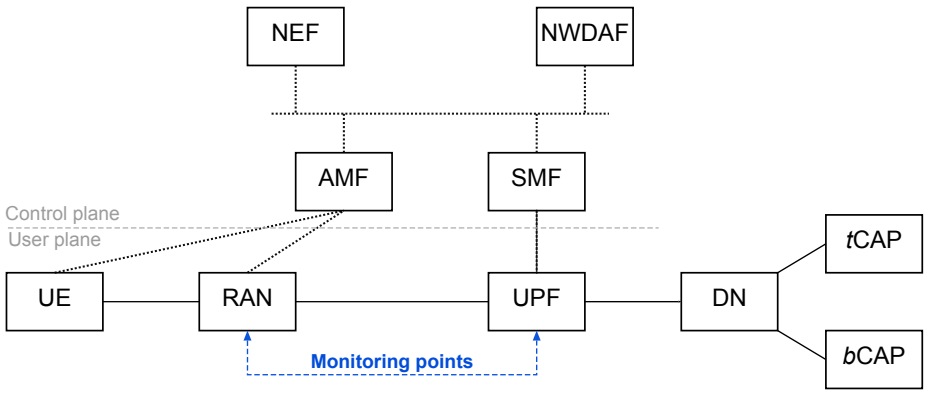

Source: elaborated by the authors.

If no traffic differentiation is implemented by the ISP, loss and delay results are expected to be similar for both CAPs, without significant differences, reflecting how the corresponding type of traffic is treated. The QoS classification and priority level should be the same for tCAP and bCAP packets and the packet scheduling policy should be neutral (null hypothesis). On the other hand, statistically significant differences in measurements are likely to be caused by traffic differentiation along the path. For instance, the ISP can classify tCAP packets with lower priority than bCAP packets, applying different packet scheduling and buffer management policies to each traffic class. In this case, higher loss rates may indicate a buffer management policy that would drop incoming tCAP packets with a higher probability than bCAP packets (e.g., weighted random early detection - WRED). In the same sense, longer delays may indicate a packet scheduling policy that provides higher priority to tCAP packets (e.g., strict priority - SP) or allocates a larger fraction of the bandwidth to bCAP packets (e.g., weighted fair queueing - WFQ). Another form of differentiation is through rate-limiting the transmission of tCAP packets using traffic shaping mechanisms in relevant nodes such as the UPF.

It should be emphasized that traffic differentiation does not necessarily imply a violation of the net neutrality principle, as it could be employed, for example, to ensure network integrity and security or to handle network congestion, which requires a holistic analysis by the regulator. Nevertheless, the proposed approach shows how real-time 5G data analytics could be used as an instrument of transparency on traffic management policies, enabling the regulator to monitor the ISPs' practices in a net neutrality context.

CONCLUSION

The fifth generation of mobile communications systems sets new traffic differentiation paradigms based on multiple programable and independent virtual networks whereby performance parameters are customized to support heterogeneous QoS requirements.

This paper presents a new use case for the standardized NWDAF, capable of supporting the assessment of net neutrality in 5G networks, increasing network transparency, and subsidizing the regulatory monitoring of Internet traffic management practices.

Future studies will focus on specifying an interface for the regulator to request and receive net neutrality-related information from the NWDAF. From a quantitative

perspective, these studies can explore key performance indicators (KPIs) of 5G slices to identify patterns that may indicate discriminatory practices.

REFERENCES

ALÉN-SAVIKKO, Anette. Network neutrality in the era of 5G–a matter of faith, hope, and design? **Information & Communications Technology Law**, v. 28, n. 2, p.115-130, 2019.

ALEXIADIS, Peter; SHORTALL, Tony. The advent of 5G: Should technological evolution lead to regulatory revolution? **CPI Antitrust Chronicle November**, v. 3, p. 1-13, 2016.

ANDREOLETTI, Davide *et al.* To be neutral or not neutral? The in-network caching dilemma. **IEEE Internet Computing**, v. 22, n. 6, p. 18-26, 2018.

BARAKABITZE, Alcardo Alex *et al*. 5G network slicing using SDN and NFV: A survey of taxonomy, architectures and future challenges. **Computer Networks**, v. 167: 106984, 2020.

BARRATT, Neil; SHADE, Leslie Regan. Net neutrality: Telecom policy and the public interest. Canadian Journal of Communication, v. 32, n. 2, p. 295, 2007.

BELLI, Luca. End-to-end, net neutrality and human rights. In: **Net neutrality compendium**. Springer, Cham, 2016. p. 13-29.

BEREC. Guidelines on transparency in the scope of net neutrality: Best practices and recommended approaches. BoR (11)6. 2011.

BEVERLY, Robert; BAUER, Steven; BERGER, Arthur. The Internet is not a big truck: Toward quantifying network neutrality. In: INTERNATIONAL CONFERENCE ON PASSIVE AND ACTIVE NETWORK MEASUREMENT. Springer, Berlin, Heidelberg, 2007, p. 135-144.

CALLANAN, Cormac *et al*. **Internet blocking: balancing cybercrime responses in democratic societies**. Aconite Internet Solutions, 2009.

CERF, Vincent. Net neutrality. Hearing before the US Senate Committee on Commerce, Science, and Transportation. U.S Government Printing Office: Washington, DC, 2006. p. 7-14.

CHATRAS, Bruno; KWONG, U. Steve Tsang; BIHANNIC, Nicolas. NFV enabling network slicing for 5G. In: 20TH CONFERENCE ON INNOVATIONS IN CLOUDS, INTERNET AND NETWORKS (ICIN). IEEE, 2017. p. 219-225.

COOPER, Alissa; MORRIS, John; SOHN, David. Re: Preserving the Open Internet, GN Docket No. 09-191; Framework for Broadband Internet Service, GN Docket No. 10-127. Washington, DC. 2010. Available at: https://cdt.org/wp-content/uploads/pdfs/CDT_FCC%20Letter_9_8_10.pdf. Access in: 30 July 2022.

DISCHINGER, Marcel *et al*. Glasnost: Enabling End Users to Detect Traffic Differentiation. In: NSDI. 2010.

ETSI. Network functions virtualization (NFV); terminology for main concepts in NFV. GR NFV 003, V1.6.1, 2021.

FINNIE, Graham. Isp traffic management technologies: The state of the art. **Heavy Reading.** Report for the CRTC, 2009.

FRIAS, Zoraida; MARTÍNEZ, Jorge Pérez. 5G networks: Will technology and policy collide? **Telecommunications Policy**, v. 42, n. 8, p. 612-621, 2018.

GARRETT, Thiago *et al*. Monitoring network neutrality: A survey on traffic differentiation detection. **IEEE Communications Surveys & Tutorials**, v. 20, n. 3, p. 2486-2517, 2018.

GHOSH, Amitabha *et al*. 5G evolution: A view on 5G cellular technology beyond 3GPP release 15. **IEEE Access**, v. 7, p. 127639-127651, 2019.

GREENSTEIN, Shane; PEITZ, Martin; VALLETTI, Tommaso. Net neutrality: A fast lane to understanding the trade-offs. **Journal of Economic Perspectives**, v. 30, n. 2, p. 127-50, 2016.

HALEPLIDIS, Evangelos *et al*. Software-defined networking (SDN): Layers and architecture terminology. In: RFC 7426. IETF. 2015.

HALL, Joseph *et al*. A survey of worldwide censorship techniques. Internet-Draft, work in progress. IETF. 2022. Available at: https://datatracker.ietf.org/doc/draft-irtf-pearg-censorship/. Access in: 15 July 2022.

ITU. Framework of software-defined networking. Recommendation ITU-T Y.3300. 2014.

ITU. IMT traffic estimates for the years 2020 to 2030. Report ITU-R M.2370-0. 2015a.

ITU. IMT Vision–Framework and overall objectives of the future development of IMT for 2020 and beyond. Recommendation ITU-R M.2083-0. 2015b.

ITU. Terms and Definitions for Network 2030. FG-NET2030 – Focus Group on Technologies for Network 2030. ITU-T. Technical Report. 2020.

ITU. Measuring digital development Facts and figures. 2021. Available at: https://www.itu.int/en/ITU-D/Statistics/Documents/facts/FactsFigures2021.pdf. Access in: 22 July 2022.

JORDAN, Scott; GHOSH, Arijit. A framework for classification of traffic management practices as reasonable or unreasonable. **ACM Transactions on Internet Technology (TOIT)**, v. 10, n. 3, p. 1-23, 2010.

KANTOLA, Raimo. Net Neutrality Under EU Law–a Hindrance to 5G Success. In: PROCEEDINGS OF THE 30TH EUROPEAN CONFERENCE OF THE INTERNATIONAL TELECOMMUNICATIONS SOCIETY (ITS). Helsinki, Finland, 2019.

KANUPARTHY, Partha; DOVROLIS, Constantine. DiffProbe: Detecting ISP service discrimination. In: 2010 PROCEEDINGS IEEE INFOCOM. IEEE, 2010. p. 1-9.

KAYE, David. A Neutralidade de Rede faz parte da luta internacional pelos direitos humanos na era digital. [entrevista concedida a] Oliver Hudson. **SUR – Revista Internacional de Direitos Humanos**, v.15, n. 27, p. 61-67, 2018. Available at: https://sur.conectas.org/wp-content/uploads/2018/07/sur-27-portugues-entrevista-com-david-kaye.pdf. Access in: 01 Out 2022.

KHANDKAR, Vinod S.; HANAWAL, Manjesh K. Detection of traffic discrimination in the internet. In: *2020* INTERNATIONAL CONFERENCE ON COMMUNICATION SYSTEMS & NETWORKS (COMSNETS). IEEE, 2020. p. 677-679.

KREUTZ, Diego *et al*. Software-defined networking: A comprehensive survey. **Proceedings of the IEEE**, v. 103, n.1, p. 14-76, 2014.

LEE, Robin S.; WU, Tim. Subsidizing creativity through network design: Zero-pricing and net neutrality. **Journal of Economic Perspectives**, v. 23, n.3, p. 61-76, 2009.

LEMLEY, Mark A.; LESSIG, Lawrence. The end of end-to-end: Preserving the architecture of the Internet in the broadband era. **UCLA Law Review**, v. 48, p. 925, 2001.

LESSIG, Lawrence. **Code Version 2.0**. New York: Basic Books, 2006.

LI, Fangfan, *et al*. A large-scale analysis of deployed traffic differentiation practices. In: PROCEEDINGS OF THE ACM SPECIAL INTEREST GROUP ON DATA COMMUNICATION. 2019. p. 130-144.

LU, Guohan *et al*. POPI: a user-level tool for inferring router packet forwarding priority. **IEEE/ACM Transactions on Networking**, v. 18, n. 1, p. 1-14, 2009.

MA, Richard TB; WANG, Jingjing; CHIU, Dah Ming. Paid prioritization and its impact on net neutrality. **IEEE Journal on Selected Areas in Communications**, v. 35, n. 2, p. 367-379, 2017.

MARQUES, Rodrigo Moreno; PINHEIRO, Marta Macedo. Informacão e poder na arena da Internet. **Informação & Sociedade**, v. 24, n. 1, 2014.

MIJUMBI, Rashid *et al*. Network function virtualization: State-of-the-art and research challenges. **IEEE Communications surveys & tutorials**, v. 18, n. 1, p. 236-262, 2015.

MOLAVI KAKHKI, Arash *et al.* Identifying traffic differentiation in mobile networks. In: PROCEEDINGS OF THE 2015 INTERNET MEASUREMENT CONFERENCE. 2015. p. 239-251.

NAKAO, Akihiro *et al*. End-to-end network slicing for 5G mobile networks. **Journal of Information Processing**, v. 25, p. 153-163, 2017.

OBIODU, Emeka; SASTRY, Nishanth. From ATM to MPLS and QCI: The evolution of differentiated QoS standards and implications for 5G network slicing. **IEEE Communications Standards Magazine**, v. 4, n. 2, p. 14-21, 2020.

PAPPAS, Christos *et al*. Transparency instead of neutrality. In: PROCEEDINGS OF THE 14TH ACM WORKSHOP ON HOT TOPICS IN NETWORKS. 2015. p. 1-7.

PEREIRA, Getúlio Emílio; ALBERTI, Antônio Marcos. Análise de Tecnologias Emergentes para a Nova Geração de Redes sem Fio. **Revista Telecomunicações**, v. 15, n. 2, p. 82-92, 2013.

RYDNING, David; REINSEL, John; GANTZ, John. The digitization of the world from edge to core. Framingham: International Data Corporation, 2018. Available at: https://www.seagate.com/files/www-content/our-story/trends/files/idc-seagate-dataage-whitepaper.pdf. Access in: 30 June 2022.

SALTZER, Jerome H.; REED, David P.; CLARK, David D. End-to-end arguments in system design. **ACM Transactions on Computer Systems (TOCS)**, v. 2, n. 4, p. 277-288, 1984.

SMIRNOVA, Inga *et al*. Network Slicing in the Scope of Net Neutrality Rules. In: 2019 PHOTONICS & ELECTROMAGNETICS RESEARCH SYMPOSIUM-SPRING (PIERS-SPRING). IEEE, 2019. p. 1516-1521.

STOCKER, Volker; SMARAGDAKIS, Georgios; LEHR, William. The state of network neutrality regulation. **ACM SIGCOMM Computer Communication Review**, v. 50, n. 1, p. 45-59, 2020.

TACD. Resolution on the open and neutral Internet. DOC NO: INFOSOC 53/15. Available at: https://tacd.org/wp-content/uploads/2015/06/TACD-INFOSOC-Resolution-on--Net-Neutrality-2015-GREEN.pdf. Access in: 30 July 2022.

VAN SCHEWICK, Barbara. **Internet architecture and innovation**. MIT Press, 2012.

VAN SCHEWICK, Barbara. **T-Mobile's Binge On violates key net neutrality principles**. Stanford Law School, The Center for Internet and Society, 2016.

WU, Tim. Network neutrality, broadband discrimination. **J. on Telecomm. & High Tech. L.**, v. 2, p. 141, 2003.

XIA, Wenfeng *et al*. A survey on software-defined networking. **IEEE Communications Surveys & Tutorials**, v. 17, n. 1, p. 27-51, 2014.

YOO, Christopher S. Beyond network neutrality. Harvard Journal of Law & Technology, v. 19, n. 1, p. 1-77, 2005.

YOO, Christopher S. Innovations in the Internet's architecture that challenge the status quo. **Journal on Telecomm. & High Technology Law**, v. 8, p. 79-99, 2010.

ZHANG, Ying; MAO, Zhuoqing Morley; ZHANG, Ming. Detecting traffic differentiation in backbone ISPs with NetPolice. In: PROCEEDINGS OF THE 9TH ACM SIGCOMM CONFERENCE ON INTERNET MEASUREMENT. 2009. p. 103-115.