

# Ensuring Network Neutrality for Future Distributed Systems

Thiago Garrett\*, Schahram Dustdar†, Luis C. E. Bona\* and Elias P. Duarte Jr.\*

\*Federal University of Paraná, Brazil

Emails: {tgarrett,bona,elias}@inf.ufpr.br

†TU Wien, Austria

Email: dustdar@infosys.tuwien.ac.at

**Abstract**—Network Neutrality is essential for ensuring a level playing field for the development of new applications and services on the Internet. Laws and rules alone might not be enough to protect innovation, fair competition and consumer’s freedom of choice online. The research community has the responsibility to propose solutions that reveal discriminatory traffic management mechanisms on the Internet. We present the potential risks of a non-neutral Internet, identify several open challenges for designing solutions that detect traffic differentiation, and propose a model that addresses such challenges by taking advantage of distributed systems technologies.

## I. INTRODUCTION

Network Neutrality (NN) has been discussed for more than a decade. However, it was just in recent years that laws and rules regarding NN began to be implemented in several places around the globe [1]. Examples include the USA, the European Union, Japan, and some countries in South America.

There is no precise definition for NN. However, a common definition that can be extracted from current regulations states that NN is the principle by which all traffic on the Internet must be treated equally. Therefore, an Internet Service Provider (ISP) cannot slow down, prioritize or block any type of specific traffic, regardless of its origin, destination and/or content, i.e., traffic differentiation (TD) practices are not allowed [2].

ISPs may employ TD to deal with congestion, or because of commercial agreements or even to benefit their own services. By slowing the traffic of bandwidth-hungry applications, such as video streaming or P2P file sharing, ISPs reduce congestion in their networks, postponing the need for upgrading the infrastructure (which is expensive) [3]. ISPs may also prioritize the traffic from providers that are willing to pay for it, the so-called fast-lanes [1]. An ISP may even prioritize its own services or degrade competitor’s services in order to attract more users, and thus increasing the revenue [4].

The discrimination between different types of traffic threatens innovation, fair competition and the consumer’s freedom of choice in the Internet [5]. In a world without NN, it would be possible for an ISP and its partners to control which services consumers would most likely use (no real freedom of choice) and which services would most likely succeed (unfair competition). This control may give the big corporations the power to greatly influence people’s online behavior and which services they consume.

Innovation in a non-neutral Internet can be hindered mostly by the ISPs and the big corporations, which can afford to have their data prioritized. Startup companies and independent innovators might not be able to fairly compete with more established services, since they do not have the same amount of resources [6], [7]. Innovative services may struggle to succeed or they may even not even see the light of day due to poor performance including higher response times caused by ISPs discriminating their traffic. NN is essential for ensuring a level playing field for the development of new applications and services on the Internet. Cloud services, Over-the-top (OTT) services, and Internet of Things (IoT) devices and applications are examples of such applications and services that should care about ensuring NN in the future.

The Internet has become a critical infrastructure which supports all kinds of businesses, from entertainment and social interaction to financial transactions. This only became possible due to the open nature of the Internet: an equal opportunity environment for innovation and freedom of expression [8]. It is thus important to protect the Internet from becoming a toll-based highway controlled by a few.

Nevertheless, laws and rules alone might not be enough to ensure a neutral Internet. ISPs may employ surreptitious TD even if it is illegal, and big companies might find loopholes in the regulations that allow for discriminatory practices [9]. It is necessary to have effective solutions to check whether ISPs are complying with the regulations, and even “legal” discriminatory practices should be transparent.

In this work, we first briefly describe the worldwide NN debate and give a strong motivation for the importance of ensuring NN. We then identify several open challenges for designing solutions that increase the transparency on TD practices, and propose a model which takes advantage of the technologies and infrastructure of current and future distributed systems in order to address such challenges.

The rest of the paper is organized as follows. Section II presents an overview of the different parties interested in the NN context. In Section III, we describe the risks of a non-neutral Internet. We identify several open challenges for detecting TD in Section IV, followed by the proposed model in Section V. We conclude the paper in Section VI.

## II. ACTORS

The context around the NN worldwide debate [10] includes several different interested parties or stakeholders. We call them actors, each having its own interests and expectations regarding the outcome of the debate. Figure 1 gives a high-level overview of these actors, how they interact and whether they are in favor or against NN.

**ISPs** provide access to the Internet to individuals and organizations. They are usually against NN since it prevents them from freely managing the traffic on their networks. Some ISPs argue that in order to support the fast growing traffic on the Internet, it is necessary to maximize their revenue for further expanding the infrastructure. To achieve such extra revenue, ISPs may employ discriminatory practices such as fast-lanes, charging extra fees from heavy traffic producers, or prioritizing their own services to attract more consumers. Therefore, ISPs expect that their businesses continue to be viable in the future, so they can keep expanding their network to accommodate the growing traffic on the Internet.

**Consumers** are the end-users that access various types of content and services offered through the Internet. They pay to the ISPs to have access to the network and thus expect to receive from them exactly the service they are paying for. Furthermore, consumers expect being able to access any content or service they choose without any interference from ISPs. Therefore, they are in favor of NN since it protects their freedom of choice.

We call **content/service providers** any entity that provides services or content through the Internet. Examples include OTT content providers (such as video and audio streaming) and Cloud services. Similarly to consumers, they pay ISPs in order to have access to the network. Content/service providers expect ISPs to just transport data packets as fast as possible, with no prioritization or any interference on their traffic. Moreover, they expect there will not be any extra fees for ensuring that their traffic is not going to be degraded in any situation, since they are already paying for access to the network. Therefore, most content/service providers are in favor of NN. However, some of them may be owned by ISPs or **big corporations**, which may have their own interests not always aligned with the NN principle. For instance, a big corporation might want to have their services prioritized in order to attract more consumers.

**Innovators** are individuals or startup companies creating new Internet-based devices, applications and services which will compete with established content/service providers. Innovators expect to be able to fairly compete without the need for special treatment from ISPs. Innovators are thus in favor of NN, since it guarantees they will be able to compete in equal conditions, at least regarding the access to the network and traffic management policies.

**Regulators** consist of governments, regulatory authorities, lawmakers and law enforcement agencies. They create and enforce laws and rules regarding the Internet, expecting to ensure the interests of all actors. Several aspects should be

considered by regulators for conceiving future-proof laws and rules regarding NN, such as: fair competition between both existing and new applications and services; Future Internet requirements; fostering innovation; consumers' freedom of choice; consumers' rights; and expanding the Internet infrastructure. Aware of all these aspects, **researchers** have the responsibility to propose solutions that help regulators to ensure a neutral Internet.

Big corporations, ISPs and some content/service providers are part of what we call the **corporate-industrial complex**. This group of big and influential companies and individuals often want to have control over the market and its consumers. They influence regulators, expecting laws and regulations to favor their own interests.

## III. THE RISKS FOR THE FUTURE

A possible picture of a non-neutral Internet does not look nice. First, ISPs would be able to charge consumers differently depending on which type of content they want to access. For instance, if the consumer wants video streaming, they should pay for the extra "video streaming" package, otherwise they will have this type of service degraded or even blocked. Alternatively, a consumer might just use the video streaming service from his/her own ISP, which would be free of extra charges. ISPs would be also able to sell fast-lanes or charge content/service providers in order to not degrade their traffic. Content/service providers which cannot afford those fees would not be able to compete with the others. Moreover, consumers could be given faster and/or cheaper access to specific services from ISPs that are partners of the service providers. In the context of IoT, an ISP could argue that their network is better suited or more secure for IoT traffic, employing discriminatory practices in order to hinder competitors or to prioritize traffic from specific device vendors.

It is important to understand that this scenario can become reality if NN is not ensured. As the network traffic increases, ISPs become more likely to adopt discriminatory policies in order to reduce the pressure on their infrastructure or to increase their revenue [5]. Similarly, as the amount of consumers on the Internet grows, the big corporations and well established content/service providers might want to get special treatment in the network in order to capture more consumers or to keep those they already have. In order to further elucidate these behaviors, some real cases are described below.

Between 2007 and 2008, several ISPs from the USA and Canada started degrading traffic generated by P2P applications (such as BitTorrent and Gnutella) [3]. They employed techniques for limiting the amount of bandwidth used by these applications. Moreover, it was observed in 2014 that Netflix performance was being degraded by two major ISPs in the USA [11]. The reason for ISPs to slow such bandwidth-hungry services is to reduce the amount of traffic in their networks, thus reducing the probability of congestion. Using this strategy, ISPs need less investments in their infrastructure, which would have otherwise to be expanded more often in order to support the much larger traffic.




Figure 1. Overview of actors in the NN context.

In 2013, Facebook launched a program called *Internet.org*, which was later renamed as *Free Basics*. Free Basics aims at providing free access to “basic” Internet services to people who can’t afford to have proper Internet access [12], i.e., it gives poor people in developing countries free access to a set of services selected by Facebook. However, if they want to access content outside what Facebook considers “basic”, they have to pay. Despite their “humanitarian” speech, one can argue that the real interest of Facebook and its partners in Free Basics seems to be all about fueling their business model. It gives them millions of new consumers who will most likely not switch to competitors, since this means changing from a free service to a paid service.

In February 2016, Free Basics and all other zero-rating practices were prohibited in India, due to violations of NN laws, since they prioritize a selected set of services. Zero-rating consists in giving free access to selected services [13]. It is a common practice among mobile carriers, which allow access to popular online services (such as social networks) without accounting for them on monthly data caps.

In March 2016, Netflix admitted limiting the bandwidth of its own video streaming to consumers of most mobile carriers across the world [14]. According to Netflix, the reason for this is “to protect users from exceeding data caps”. However, this practice was not employed for consumers of at least 2 ISPs in the USA, since “historically those two companies have

had more consumer-friendly policies”. Therefore, Netflix, a leading proponent of NN, was prioritizing a few ISPs which have policies more aligned to its interests. While this is not against any current NN law or regulation, it certainly shows that not only ISPs can make use of discriminatory practices in order to gain a questionable control of the market.

These examples show market strategies that can pose a huge threat to innovation, fair competition and consumers’ freedom of choice. If NN is not ensured, then corporate-industrial forces will have the power to control what consumers will buy and from whom. Innovative services could just disappear or not even see the light of day. Discriminatory practices could be employed to reflect corporate alliances or commercial/political agreements, favoring selected companies or even influencing people to vote on certain candidates. Note that there is a thin line between these issues and censorship [15]: influencing people by favoring one side is not so different from blocking the other. The Internet would cease to be an equal opportunity environment for innovation and freedom of expression.

Regulators are creating laws and rules implementing NN all around the world [1]. However, this might not be enough to ensure NN. Even being illegal, those interested in a non-neutral Internet may still employ discriminatory practices, or even do it legally through loopholes on the regulations [9]. Furthermore, regulators themselves are also subject to corporate-industrial influences.

#### IV. OPEN CHALLENGES

The debate regarding NN [10] has mainly focused on political, economical, social, ethical, and legal aspects so far. Technical issues, such as how to detect TD, have been underexplored in the past decade. Since regulations by themselves are not enough, ensuring a neutral Internet may well depend on further exploring technical issues. The research community has the responsibility to propose new solutions to help regulators, consumers and innovators be aware of what is happening in the network.

Regardless of being legal or not, we argue that TD practices should be transparent. In order to achieve such transparency, solutions for detecting TD are necessary. For instance, these solutions may help regulators to enforce their laws, and consumers to be more aware about the services provided by their ISPs and whether or not they want to switch to a competitor. Innovators may be able to tweak their new applications and services based on how their traffic is being treated.

Since the last decade some proposals for TD detection have been published [16]–[24]. These proposals are based on network measurements and statistical inference. They take measurements from one or several hosts and for different types of traffic. The obtained measurements are compared in order to infer whether there is a significant discrepancy among different sets of measurements. Good statistical models are necessary in order to distinguish variations caused by TD from those caused by other phenomena such as congestion or load balancing – the confounds. However, there are still several open challenges for designing more capable and future-proof solutions. We identify below several of these challenges.

Further investigations on metrics and measurement strategies are necessary. Most existing solutions cannot infer which ISP is practicing TD. Those solutions that address this issue rely on TTL-based probes – which are not universally supported by routers – or require prior knowledge of the network topology. Furthermore, most current measurement strategies result in high network overhead, since they generate a large amount of traffic in order to saturate the network and force TD to occur. Another issue is that several existing solutions require control of all the end-hosts involved in the measurement, which might not be a realistic assumption.

Very few existing solutions address TD in mobile networks, in which different confounds and constraints are present. Measurements in mobile networks may be affected by fluctuations in channel quality or mobility, for example. It is also not feasible to perform measurements which generate large amounts of traffic, since mobile devices are usually subject to data caps.

The presence of TD and how it affects the traffic may change over time or depend on network conditions. An ISP might employ TD only on periods of the day during which the network is under heavy usage, for example, or change which TD mechanism to employ depending on the user location. Existing solutions are not designed to detect such dynamic behaviors, since they usually consist of one-shot analysis,

thus can only detect TD being employed at the time of their execution.

Another challenge is to make any Internet-based application, service or device aware of whether its traffic is being discriminated. There is no proposal currently that enables an arbitrary application to monitor how their traffic is performing compared to others without having to implement TD detection on their own. Such feature would allow applications not only to benefit from this type of monitoring but also to contribute increasing the precision of the system.

The Internet infrastructure and TD mechanisms employed by ISPs are ever-changing. The design of TD detection solutions should be extensible, enabling them to keep up with network evolution. However, existing solutions are designed assuming a set of network features, as well as specific TD mechanisms and a set of applications that may be discriminated.

#### V. NN MONITORING MODEL

We propose a NN monitoring model, shown in Figure 2, which addresses the technical challenges presented above and also enables innovators and consumers to benefit from and contribute to the system. The model allows any kind of device (mobile or wired) or third-party applications to join the system, contributing with measurements and/or checking how their traffic is being treated.

The idea is to continuously monitor measurements obtained by a plethora of agents (crowdsensing) in a hybrid active/passive approach, while making all acquired data available in an Open Data paradigm for further analysis by the system itself or by any third party system. This strategy takes advantage of several features of distributed systems, thus enabling such systems to incorporate the proposed model. We argue that the NN-related issues discussed in this paper should be taken into account when designing distributed systems or any other Internet-based application.

The authors of [25] advocate the use of a similar approach to build a “citizen observatory” of NN in the context of mobile Internet. In this paper, we build on that idea targeting a NN monitoring system that can gather data from any kind of source (not only mobile devices) for better assessing the behavior of the network. Furthermore, we also propose an actual model and possible directions for implementing such ideas in a real system. To the best of our knowledge, there is no solution currently that employs hybrid active/passive measurements and crowdsensing for detecting TD.

We describe the proposed model in Subsection V-A. In Subsection V-B we give a discussion of the model, followed by some possible implementation scenarios in Subsection V-C. Finally, we present the challenges for implementing the proposed model in Subsection V-D.

##### A. The Model: How It Works

The model is divided in four components, as shown in Figure 2: measurement agents, continuous monitoring, storage,




Figure 2. NN monitoring model.

and data analyses. Each component has a specific purpose, described below.

The measurement agents can be embedded in virtually any device capable of making and reporting measurements. Examples include smartphones, tablets, personal computers, IoT sensors, middlewares, Cloud services or any other third-party application. During the normal operation of such devices or applications, the agents make passive measurements of the traffic and collect confounds. They may also perform active measurements when requested. Examples of traffic-related metrics are delay, loss rate, and bandwidth. Confounds are the factors that may affect the measurements (other than TD) and/or help characterizing them – such as location, type of network (mobile or wired), ISP, application protocol, and time of day.

Different agents may have different sensing capabilities, due to differences in hardware or operating system features. If a measurement agent is embedded in a third-party application, for example, it may be able to report only measurements regarding the traffic of the application itself, since the application might not have enough permissions to measure all traffic that goes through the device on top of which it is running. On the other hand, if an agent is running in a personal computer with enough privileges, it may be able to make passive measurements regarding all the traffic in the device.

As measurements and confounds are reported by agents, they are aggregated and real-time TD inference is performed. If a potential presence of TD is detected, active measurements

can be promptly requested to the corresponding agents for further investigating the suspicious case. Furthermore, there can also be measurement campaigns, in which active measurements are issued regardless of suspicions, configuring a more proactive approach instead of just reacting to potential cases of TD.

All data (measurements, confounds and inferences) should be anonymized and stored in a database. This database should be publicly accessible through an Open Data API. Furthermore, data should be distributed and replicated, in order to both increase its availability and protect it from any potential attack coming from those that might be compromised by the information. From the obtained data, the system can make more detailed and complex analyses. Participating applications and devices may benefit from these analyses at runtime, changing their behavior depending on how their traffic is being treated, for example. It is also possible for third parties to access the data and make their own use of it, expanding the capabilities of the model. Examples of analyses that may be performed include: warnings regarding NN violations; a “neutrality index”, which indicates how neutral each ISP is; historical behavior of different ISPs or applications; and data mining, which can identify patterns regarding the deployment of traffic management techniques.

#### B. Discussion: Thoughts on the Model

The crowdsensing approach takes advantage of the large amount of devices already deployed in the wild. The model also allows for any third-party application or service to benefit

from the TD inference provided by the system without having to implement it. It is similar to a service-oriented approach, in which any agent can make use of the “TD inference service”. Moreover, the data each agent contributes is used to improve the overall effectiveness of the system for all participants. This aggregation of measurements from multiple vantage points may help not only detecting TD, but also identifying which ISP is employing it. Furthermore, this approach does not require control of end-hosts, since edge-devices and applications are external entities contributing willingly to the system.

The hybrid active/passive measurement strategy has a much lower overhead when compared to purely active strategies. By passively capturing measurements, it is possible to detect when TD might be occurring and then trigger active measurements. Thus there is no need for generating a large amount of artificial traffic in order to saturate the network before taking measurements. Continuous monitoring also enables the detection of dynamic behaviors – such as an ISP employing TD only on specific periods of the day. The historical data obtained allows for deeper analyses regarding traffic management policies and TD patterns.

The proposed model makes no assumption regarding the network, TD mechanisms, applications being discriminated, or characteristics of the participating devices and applications. The measurement agents may be embedded in anything, such as edge-devices or even another system, which may be connected to any type of network (mobile or wired). This agnostic approach makes the model more future-proof to the evolution of the networks, devices and protocols. Specific characteristics of the agents and the network are reflected by the confounds during aggregation and real-time inference, and later analyses may also be performed considering such specific properties.

By adopting the Open Data paradigm, this model not only helps ensuring innovation by monitoring NN compliance, but also creates new possibilities on its own for new innovative solutions. Third parties can create applications and services that make unforeseen uses of the data obtained by the system. Therefore, the crowdsensing approach allied with Open Data enables any consumer and/or innovator to contribute with a more transparent and innovative Internet.

### *C. Implementation: Possible Scenarios*

A scenario for implementing the proposed NN monitoring model is the Smart City. A Smart City can take advantage of its infrastructure to provide its citizens and other stakeholders a NN monitoring service. Citizens would be more aware of traffic management practices employed by ISPs, empowering them to demand their rights as consumers. This transparency may also enable other stakeholders, such as public authorities, to make more informed decisions or even take law enforcement actions. A possible implementation of the proposed model in this scenario is embedding it as part of the Smart City middleware. In this approach, the measurements can be collected as sensors and other Smart City devices communicate through the middleware.

Another possibility is to deploy measurement agents on large-scale testbeds, such as PlanetLab, and/or on measurement platforms, such as M-Lab. Devices in these platforms normally generate a large amount of traffic, making them a prolific environment for making measurements. This would allow for a deeper study of traffic management policies around the globe. These platforms also provide a real-world environment for evaluating the proposed model. Furthermore, there are several platforms with different properties, such as sensor networks and mobile testbeds, forming a diverse set of networks and devices.

### *D. Implementation challenges*

The proposed model addresses several open challenges regarding the TD detection problem. However, implementing it comes with a whole new set of challenges. We identify several open issues for implementing the proposed model and also provide some possible directions.

The main implementation challenge is arguably how to aggregate all the measurements from different sources and compare them to make NN-related inferences. Different sources may have different confounds and thus their measurements might not be comparable. Therefore, a good statistical model should be employed to properly adjust for confounds, achieving a more reliable inference [21]. For instance, comparing measurements from a smartphone (connected to a mobile network) with measurements from a personal computer (connected to a broadband network) might not be a good approach. Measurements in a mobile network can be affected by several different phenomena from those in a broadband network, such as fluctuations in channel quality and mobility.

The proposed model depends on devices and applications voluntarily joining the system, in a community sensing approach. Creating incentives and favorable conditions for the massive adoption of the system is certainly a big challenge for deploying the proposed model in the wild [25]. Consumers might need motivations for installing a monitoring application on their devices, for example. Easy to use software frameworks, libraries and APIs might help increasing the adoption of the system in applications and services. Another issue that may be of concern is privacy. Sensible data may be collected, such as location of devices and specific applications being used, which can be intimidating. Ensuring proper anonymization of the data is thus particularly important.

The hybrid measurement approach also poses a new challenge: deciding when and on which agents active measurements should be performed. Active measurements could be required immediately after a TD suspicion is detected, or when the involved agents are available or devices are idle. It is even possible to schedule measurement campaigns. Several potential constraints on the agents must be taken into account. Mobile devices, for example, have energy and processing limitations, and due to mobility and wireless signal reception issues might not always be available. Moreover, mobile devices are usually subject to data caps, thus active

measurements should not generate too much traffic in such devices.

Real-time inference as proposed in the model presents a couple of challenges that must be addressed as well. Real-time inference consists basically of statistical analyses of measurement distributions. While these analyses might not be a problem in terms of computational complexity, they may become a problem when it comes to scalability. All measurements received should be dealt with relatively fast, since longer delays may cause active measurements to be triggered too late, when TD might be no longer occurring. Furthermore, there will be potentially a huge amount of agents continuously reporting measurements. Add the fact that the real-time inference will not be based only on the measurements just received, but also on previous collected/inferred data. Stream processing technologies and/or distributed approaches are possible directions to address this scalability issue.

Finally, defining which later analyses should be performed with the acquired data is also an open issue. It may require a deeper understanding of the economical, social and political aspects involved in the NN context. Further investigation is thus necessary to assess what is possible to achieve with the data available and what is truly relevant.

## VI. CONCLUSION

Ensuring NN goes beyond just regulating ISPs. It is necessary to have tools for checking the compliance of ISPs and other corporate-industrial forces according to the regulations. Even in a non-regulated environment, transparency is important and can lead to a more competitive market.

Nevertheless, detecting discriminatory practices is a non-trivial task. There are still several open challenges that must be addressed by the research community. The model presented in this paper is a first step towards a more capable and future-proof solution that takes advantage of other emerging technologies. Further investigation in multiple topics is still necessary. Statistical models, streaming analytics, data anonymization, scalability, data mining, just to name a few. We are in a critical moment and a joint effort of researchers from different areas is important for securing the Internet as an open and innovative environment.

## REFERENCES

- [1] H. Habibi Gharakheili, A. Vishwanath, and V. Sivaraman, "Perspectives on Net Neutrality and Internet Fast-Lanes," *SIGCOMM Computer Communication Review*, vol. 46, no. 1, pp. 64–69, January 2016.
- [2] J. Crowcroft, "Net Neutrality: The Technical Side of the Debate: a White Paper," *SIGCOMM Computer Communication Review*, vol. 37, no. 1, 2007.
- [3] M. L. Mueller and H. Asghari, "Deep packet inspection and bandwidth management: Battles over BitTorrent in Canada and the United States," *Telecommunications Policy*, vol. 36, no. 6, pp. 462–475, 2012.
- [4] J. Kendrick, "T-Mobile Germany Blocks iPhone Skype Over 3G and WiFi," 2009, accessed in January 25, 2017. [Online]. Available: <https://gigaom.com/2009/04/06/t-mobile-germany-blocks-iphone-skype-over-3g-too>
- [5] B. van Schewick and D. Farber, "Point/Counterpoint: Network Neutrality Nuances," *Communications of the ACM*, vol. 52, no. 2, pp. 31–37, February 2009.
- [6] H. Guo and R. F. Easley, "Network Neutrality Versus Paid Prioritization: Analyzing the Impact on Content Innovation," *Production and Operations Management*, vol. 25, no. 7, pp. 1261–1273, 2016.
- [7] A. Cooper and I. Brown, "Net Neutrality: Discrimination, Competition, and Innovation in the UK and US," *ACM Transactions on Internet Technology*, vol. 15, no. 1, February 2015.
- [8] T. Berners-Lee, "Long Live the Web," *Scientific American*, vol. 303, no. 6, pp. 80–85, 2010.
- [9] P. Maille, G. Simon, and B. Tuffin, "Toward a net neutrality debate that conforms to the 2010s," *IEEE Communications Magazine*, vol. 54, no. 3, pp. 94–99, March 2016.
- [10] A. Joch, "Debating net neutrality," *Communications of the ACM*, vol. 52, no. 10, pp. 14–15, October 2009.
- [11] J. Brodtkin, "Netflix performance on Verizon and Comcast has been dropping for months," <http://arstechnica.com/information-technology/2014/02/netflix-performance-on-verizon-and-comcast-has-been-dropping-for-months>, October 2014, accessed in October 19, 2016.
- [12] M. L. Best, "The Internet That Facebook Built," *Communications of the ACM*, vol. 57, no. 12, pp. 21–23, December 2014.
- [13] L. Taylor, "From Zero to Hero: How Zero-Rating Became a Debate about Human Rights," *IEEE Internet Computing*, vol. 20, no. 4, pp. 79–83, July 2016.
- [14] R. Knutson and S. Ramachandran, "Netflix Throttles Its Videos on AT&T, Verizon Networks," <http://www.wsj.com/articles/netflix-throttles-its-videos-on-at-t-verizon-phones-1458857424>, March 2016, accessed in October 19, 2016.
- [15] G. Aceto and A. Pescapé, "Internet Censorship detection: A survey," *Computer Networks*, vol. 83, pp. 381–421, 2015.
- [16] M. Dischinger, M. Marcon, S. Guha, K. P. Gummadi, R. Mahajan, and S. Saroiu, "Glasnost: Enabling End Users to Detect Traffic Differentiation," in *USENIX Conference on Networked Systems Design and Implementation (NSDI)*, 2010, pp. 405–418.
- [17] P. Kanuparth and C. Dovrolis, "DiffProbe: Detecting ISP Service Discrimination," in *IEEE INFOCOM*, March 2010, pp. 1–9.
- [18] G. Lu, Y. Chen, S. Birrer, F. E. Bustamante, and X. Li, "POPI: A User-Level Tool for Inferring Router Packet Forwarding Priority," *IEEE/ACM Transactions on Networking*, vol. 18, no. 1, pp. 1–14, February 2010.
- [19] A. Molavi Kakhki, A. Razaghpahan, A. Li, H. Koo, R. Golani, D. Choffnes, P. Gill, and A. Mislove, "Identifying Traffic Differentiation in Mobile Networks," in *Proceedings of the 2015 ACM Conference on Internet Measurement Conference*, ser. IMC '15. ACM, 2015, pp. 239–251.
- [20] R. Ravaoli, G. Urvoy-Keller, and C. Barakat, "Towards a General Solution for Detecting Traffic Differentiation at the Internet Access," in *International Teletraffic Congress (ITC)*, September 2015, pp. 1–9.
- [21] M. B. Tariq, M. Motiwala, N. Feamster, and M. Ammar, "Detecting Network Neutrality Violations with Causal Inference," in *International Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '09. ACM, 2009, pp. 289–300.
- [22] U. Weinsberg, A. Soule, and L. Massoulié, "Inferring traffic shaping and policy parameters using end host measurements," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 151–155.
- [23] Y. Zhang, Z. M. Mao, and M. Zhang, "Detecting Traffic Differentiation in Backbone ISPs with NetPolice," in *ACM SIGCOMM Conference on Internet Measurement Conference*, ser. IMC '09. ACM, 2009, pp. 103–115.
- [24] Z. Zhang, O. Mara, and K. Argyraki, "Network Neutrality Inference," *SIGCOMM Computer Communication Review*, vol. 44, no. 4, pp. 63–74, October 2014.
- [25] D. Miorandi, I. Carreras, E. Gregori, I. Graham, and J. Stewart, "Measuring net neutrality in mobile Internet: Towards a crowdsensing-based citizen observatory," in *IEEE International Conference on Communications Workshops (ICC)*, June 2013, pp. 199–203.