Transparency Instead of Neutrality

Christos Pappas*, Katerina Argyraki†, Stefan Bechtold*, Adrian Perrig*

*ETH Zürich
{pappasch, sbechtold, aperrig}@ethz.ch

†EPFL, Switzerland
katerina.argyraki@epfl.ch

Categories and Subject Descriptors

C.2.3 [**Computer-Communication Networks**]: Network Operations

General Terms

Measurement, Security

1. INTRODUCTION

With the proliferation of applications that rely on the Internet, network connectivity has become an elementary requirement in today's societies. Coupled with the emergence of a booming digital economy, the importance of network connectivity rightfully raises concerns about who will shape the rules of the game in the coming years. Search-engine providers have been accused of manipulating results in order to favor their own content services over competitors [1, 2]; Internet service providers (ISPs) have been accused of blocking competing services from their networks [3, 4] and of occasionally hijacking unresolvable DNS requests to display ads [5]. In vertical market structures that are characterized by a lack of competition at certain levels – e.g., at the level of Internet access, content or search engine provision – market failures may prevent a healthy competitive ecosystem from emerging.

One of the most contentious issues in this debate focuses on the question of whether ISPs should be allowed to discriminate against certain kinds of Internet traffic. Regulators in the U.S., Europe, and beyond have started to introduce network-neutrality regimes that require ISPs to treat all network packets equally. While this neutral notion seems very appealing, in particular to individuals untrained in network science, it has created a formidable mess within the Internet community.

With this paper, our goal is to provide arguments to help clean up the mess. We first argue that current technical definitions of neutrality are problematic and even undesirable: they lead to situations where neutrality violations are indistinguishable to end users from – until now – legitimate ISP practices, and a non-neutral ISP can, in fact, be more appealing to end users. We then argue that the technical discussion should refocus: rather than trying to detect neutrality violations or enforce neutrality through technical means, we (the technical community) should focus on increasing network transparency. We discuss transparency solutions that would make trustworthy information about packet loss and latency of an administrative domain (AD) available to only authorized entities. Other researchers have advocated network transparency in the past – albeit not in connection with neutrality – so we frame candidate solutions to certain aspects of the problem and outline the sub-problems that remain open.

Transparency could alleviate some of the concerns raised by network-neutrality proponents. As the exposed information would be detailed and trustworthy, consumers would have much richer information available on how "neutral" their ISP is and could make informed decisions about which ISP fits their needs best. Even if increased transparency does not lead to full competition on Internet access, regulators and courts could combine this verifiable information with higher-level information to make an informed judgment about whether an ISP has violated network-neutrality rules or other duties from antitrust, communications, contract, tort, or criminal law. It is our hope, however, that the increased competition induced by better information will reduce the importance of neutrality regulations.

This paper does not argue that policy attempts to define network neutrality are misguided. Rather, it argues that *technical* attempts to define network neutrality are doomed; that the real contribution of the technical community could lie in the creation of a transparent network layer, which could fundamentally alter many discussions on Internet policy and network management; and that such contribution could also diminish the importance of neutrality regulations.

HotNets '15, November 16–17 2015, Philadelphia, PA, USA
Copyright 2015 ACM. ISBN 978-1-4503-4047-2/15/11..$15.00
DOI: http://dx.doi.org/10.1145/2834050.2834082

2. NEUTRALITY, NOT

> "Network neutrality is the idea that ISPs should treat all data that travels over their networks fairly, without improper discrimination in favor of particular apps, sites, or services." – Electronic Frontier Foundation

This is a typical example of a definition encountered in the neutrality debate: it qualitatively captures the spirit of neutrality, but it cannot be used to detect or enforce it through technical means.

In contrast, the technical networking community has defined neutrality violation in terms of specific mechanisms: blocking drops all target traffic[1] [6]; policing drops enough target traffic to enforce a maximum allowed rate [7]; shaping drains a separate queue with the target traffic at a maximum allowed rate [7].

One problem with this mechanism-based definition is that it leads to scenarios where a non-neutral network leads to a better and more balanced user experience than its neutral counterpart. For example, consider a standard IP network, where each forwarding device subjects each packet to the same algorithm, whether it carries an email or a video frame. This is the best we can ask from a network in terms of neutrality and yet, does this network really treat all traffic "the same"?

It is easy to construct scenarios where the network is neutral by today's standards, yet the user experience is excruciating for some applications. For example, it makes sense that traffic from medical applications be prioritized over BitTorrent traffic. The same is true for any latency-sensitive traffic; it has been shown that a 500ms increase in latency drops users' search queries by 20% [8]. Yet a network that is neutral by today's standards treats packets from medical applications, search queries, and BitTorrent with the same priority.

To argue about neutrality in a meaningful way, we need to involve application semantics. When users complain about neutrality violations, they typically refer to their experience, e.g., voice quality or number of buffering events. If there existed a universal scale for rating user experience, then we could say that a network is neutral when user experience is the same for all network-bottlenecked applications. Lacking such a scale, is there a compelling reason to prefer a neutral network layer over a non-neutral?

Another problem with the mechanism-based definition of neutrality is that the chosen mechanisms can have the same end-to-end effect on target traffic as legitimate ISP practices like traffic engineering or peering. Does it make sense to consider, say, shaping of P2P or video traffic a neutrality violation when a traffic engineering policy or a peering dispute can produce the same end-to-end effect on that traffic? We discuss two illustrating examples.

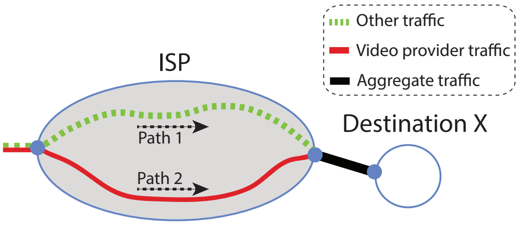

Figure 1: Traffic engineering and neutrality.

Traffic engineering. Consider the scenario in Figure 1, where an ISP (shaded network) receives at the same ingress point traffic from a video provider and other latency-sensitive traffic, both addressed to a destination X; for simplicity, we ignore any other traffic received/forwarded by the ISP. At first, the ISP forwards all this traffic through the same intra-domain path (Path 1). Then it observes that the video provider occasionally produces traffic bursts that consume most of the bottleneck capacity of Path 1, and that the resulting congestion harms the other latency-sensitive traffic.

To protect the rest of the traffic, the ISP starts shaping the video provider's traffic such that it consumes a small fraction of Path-1 capacity. This is typically considered a neutrality violation, and the ISP is accused of discriminating against the video provider.

To avoid the accusation, the ISP stops the shaping and starts forwarding traffic from the video provider through a separate path (Path 2), whose bottleneck link has the same transmission rate and queue size as the previously used shaper. This approach has exactly the same end-to-end effect on the video provider's traffic as shaping, and it is typically not considered a neutrality violation. In fact, such flexible routing decisions are an ever increasing practice in today's networks [9].

Is the ISP violating neutrality by presenting the video provider's traffic with different network conditions (a smaller-capacity pipe) than other traffic observed at the same ingress and egress points?

Peering dispute. Consider the scenario in Figure 2, where a broadband ISP receives all ingress traffic through inter-domain Link 1, except for traffic that originates at a specific content provider, which is received through inter-domain Link 2. The broadband and backbone ISP have a peering agreement to exchange traffic free of charge. At some point, the popularity of the content provider increases such that inter-domain Link 2 becomes regularly congested.

This leads to the following dispute, which is a simplified version of the Netflix disputes[2] with Comcast and

[1] We use the term "target traffic" to refer to the traffic class that is discriminated against, e.g., P2P or traffic originating from specific content providers.

[2] The difference is that, in those cases, the peering link between the backbone (Cogent) and broadband (Comcast or Verizon) ISP may have carried a mix of traffic from various sources, with Netflix contributing the majority of (but not

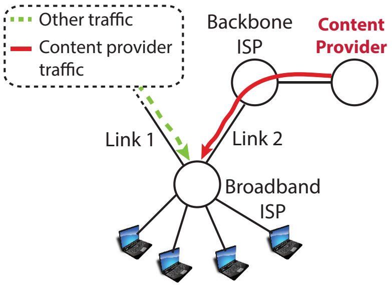

Figure 2: Peering disputes and neutrality.

Verizon [10]: The backbone ISP is happy with the extra traffic (because it can charge the content provider for it) and wants to upgrade the capacity of the peering link. However, the broadband ISP cannot monetize the extra traffic and refuses to upgrade the peering link unless their agreement is renegotiated and it charges the backbone ISP for some of the extra traffic. As a result of the dispute, the quality with which the broadband ISP's customers experience the content deteriorates.

The broadband ISP's refusal to upgrade the congested peering link is another way of rate-limiting the ingress traffic that originates at the content provider. Consider an alternative scenario where the broadband ISP receives all ingress traffic through a single inter-domain link that has the aggregate capacity of Links 1 and 2 from Figure 2, and it shapes traffic from the content provider at the transmission rate of Link 2. The two scenarios have the same end-to-end effect on the content provider's traffic.

Is the broadband ISP violating neutrality by presenting the content provider's traffic with a smaller-capacity pipe than other ingress traffic addressed to the same customers? (Netflix argued that it is.) Differently said, to qualify as neutral, must a network provide either congestion-free or equal-capacity pipes between every possible Internet source and any given destination?

Conclusion. To detect or enforce neutrality through technical means, we need a concrete technical definition. However, the way our community has defined neutrality, non-neutral networks may be preferable to neutral ones. In addition, these definitions rely on specific network mechanisms that can have the same end-to-end effect on the target traffic as – until now – legitimate ISP practices. Should we classify all these as neutrality violations?

In the end, when an ISP receives at an ingress point more traffic than it can deliver congestion-free, it has to use some form of admission control. It could randomly drop a fraction of incoming traffic; or discriminate against traffic that the ISP expects to be less useful to its customers in the short term (e.g., web-crawling traffic); or discriminate against traffic that brings less economic benefit to the ISP itself (e.g., P2P traffic or traffic from certain content providers). A regulator or court may analyze these admission criteria and find them inconsistent with the idea of network neutrality, but it does not make sense to automatically label an ISP non-neutral based on particular network-layer mechanisms used for admission control.

Our position is that our community should not pursue a technical definition of network neutrality at all. The goal of the neutrality debate is to create an open, healthy market, where all applications have equal opportunity to succeed. This is not something that can be captured by analyzing network mechanism or policy in isolation, but requires a holistic analysis of mechanism, policy, and the broadband and backbone markets.

Instead, the network layer should provide transparency: information that (a) can be used as a building block for a holistic analysis of ISP behavior, and (b) is verifiable by private parties, regulators and courts. It will then be the job of regulators to combine the information provided by a transparent network layer with higher-level information (e.g., peering agreements between ISPs, communication between ISP executives) and argue in court whether an ISP's practices are against a healthy market or not.

3. TRANSPARENCY, HOW?

We now discuss how to provide network transparency. We start with an informal problem description (§3.1), summarize interesting aspects of the problem that have already been addressed by prior work (§3.2), and outline those that remain open (§3.3).

3.1 Problem Description

We define an "aggregate" as a finite set of packets that have a subset of their headers in common. The Internet consists of non-overlapping "administrative domains"; each AD has a well-defined set of entry and exit points where it exchanges traffic with other ADs.

Goal. Consider an aggregate that traverses AD X; X should export enough information to any on-path entity Y that observes this aggregate such that Y can estimate, for this aggregate, the packet-loss rate, the packet-modification rate, and the packet-delay distribution experienced by the aggregate within X. We discuss how we think this information can be useful in Section 4.

The exported information should be trustworthy in the face of X's equipment failure, attacks against X, and manipulation by X. Hence, we require that when Y estimates the packet-loss or modification rate experienced by a given aggregate within X, it should be able to compute the relative standard deviation of the estimate; and when Y estimates a packet-latency distribution, it should be able to compute lower and upper bounds for it, as well as a confidence level for the bounds; these accuracy metrics have been specified in detail in the literature [11]. A minimum level of accuracy for each type of disclosed information is required to ensure that

necessarily the only) traffic on the peering link.

trustworthy conclusions can be reached, and this can be specified through regulation; then, ADs can start competing on the quality of the information they export, and the system should converge to the accuracy that matters to the users.

An alternative approach to bootstrapping transparency is to rely on data collected from an active measurement framework [12, 13, 14, 15, 16]. Although we stand in solidarity with such edge-based techniques, we do not expect the information they provide to be useful in court, at least not on its own. First, the information provided by the end points (that participate in the measurement framework) cannot be trusted, thus, feedback from the network is required anyway to catch dishonest reports. Second, ISPs can treat probe traffic preferentially if they know the addresses of the measurement vantage points or other distinguishing properties of the measurement traffic. Third, information inferred from end-to-end measurements can be inconclusive (a packet loss or latency event observed end-to-end can be attributed to multiple possible causes).

Cost and flexibility constraints. We cannot avoid adding new mechanism to the network devices located at AD boundaries, but cost and complexity should stay within reachable bounds.

First, we should not significantly increase the (aggregate) data-to-control-plane bandwidth of networks. To export information on observed traffic, a network device collects it at the data plane, then transfers it to the control plane. So, exporting x bits per packet requires a data-to-control-plane pipe that has x times the aggregate bandwidth of the device in terms of minimum-sized packets. Given today's forwarding capacities of multiple 10 GbE ports per device, even low values of x can result in a tremendous bandwidth overhead.

Second, we should not require more than a few MB of memory per data-path chip.[3] Assuming a device with a line rate of a few tens of Gbps and minimum packet size (the worst case), data-path memory must support a lookup time of a few tens of nsec, which is best achieved with on-chip memory. Existing manufacturing processes do not integrate more than a few MB of memory – based on our discussion with hardware experts, this is because the yield becomes too low for the manufacturing process to be sustainable.

Third, ADs should not export information on a predetermined set of aggregates; we do not see how ADs and regulators would agree on such a set. Rather, we want a flexible approach where each AD exports the same information to different recipients, and each recipient can combine the information exported by different ADs to derive estimates for the aggregates that are relevant to the recipient.

As a result of these constraints, we cannot reuse proposals where network devices export several bytes per observed packet [17], store state per TCP/UDP flow on the data path [18], or export summaries/sketches for specific aggregates [19, 20, 21].

3.2 Sub-problems We Know How to Solve

Performance estimation with honest ADs. If we could assume that ADs always export correct information, then we could satisfy our constraints by building on consistent sampling [22, 23] in the following way.

Each AD entry and exit point exports a cryptographic summary and a timestamp on a small sample of the packets it observes (e.g., 1%). The sampling is such that when a set of entry/exit points (from multiple ADs) observe the same packet p, either they all sample p or none of them does; this can be achieved through hash-based sampling, where each point chooses whether to sample a packet or not by using the same hash function on non-mutable packet fields. Exported information for packet p is accessible only to entities that observed p (or would have observed it if it had not been dropped).

One can then estimate network performance between two points for a given aggregate as follows: First, by comparing the information exported for the same packet p by two points, one can determine whether p was lost or modified between the two points and, if not, what delay it experienced. Second, by combining the conclusions drawn for multiple packets of a given aggregate, one can estimate the packet loss, modification rate, and delay distribution experienced by the aggregate.

This approach provides the flexibility we require. ADs do not track individual aggregates, but sample packets based on the same hash function; at the same time, different entities can access different subsets of the information exported by ADs and combine these to compute estimates for the aggregates that are relevant to them.

Catching lies by dishonest ADs. A dishonest AD can export false information to exaggerate its performance; can we catch such lies and how? A lot of work already exists in the area of fault localization and forwarding accountability that answers this question in different contexts [18, 21, 24, 25, 26, 27]. It shows that we can track down each lie to two consecutive ADs, as long as we make all ADs that observe a given aggregate export the same information about the same traffic units. Consider the case where an AD's exit point claims observing packet p, but the next AD's entry point claims not observing p. Any entity outside the two ADs that receives information about p concludes that either the link between the two ADs dropped p or one of them is lying. The two ADs themselves know more: if one is lying, both know it, and both know the liar since each of them knows whether it observed p or not. So, a lie becomes an externally visible inconsistency between the

[3] In high-end devices, a chip typically corresponds to a linecard with one or a few multi-Gbps ports.

information exported by two consecutive ADs, while the liar is exposed to the neighbor that was implicated in its lie. This result holds even in the case where two or more ADs lie in collusion, in which case the lie is tracked down to the last lying AD and the next neighbor. Note that these proposals provide tools to argue if an event (e.g., packet drop) happened or not, to locate where it happened (e.g., inter-domain link), but not to determine the cause of the event (e.g., congestion or deliberate packet drop).

3.3 Open Sub-problems

Preventing sampling bias by dishonest ADs. Another way an AD can misbehave is through a coward attack [28]: it can treat sampled packets preferentially, in which case the corresponding estimates computed based on these packets would be biased.

Can we leverage sampling for its low cost and high flexibility while preventing the ADs from treating sampled packets preferentially?

In theory, this is possible through delayed disclosure of the sampling function [27, 29]: Each AD entry/exit point computes and temporarily stores on the data path state for every single observed packet, while it periodically receives an in-band "disclosure signal" that tells it which of the temporarily stored state to export. The key is that when an information-exporting point observes and makes a forwarding decision for packet p, it does not yet know whether it will have to export information on p, hence it cannot treat p preferentially.

The challenge is how to implement this idea in a real network device. The related proposals [27, 29] suggest implementations with extraordinary requirements on network devices, like per-packet timers (for discarding temporary state that does not get sampled) and/or tens of MB of data-path memory per Gbps port (for storing the temporary state until it is safe to discard it). Can we implement delayed disclosure without imposing infrastructure requirements that outweigh the benefits of sampling as a lightweight solution?

Transparency & sensitive AD information. Is it possible for ADs to export the information we want while keeping their internal topology, infrastructure properties, and business agreements secret?

A world in which all such information is publicly available is neither in the interest of the individual ADs, nor necessarily in the interest of society at large. First, revealing information about topology and link capacities/latencies would make it easier to launch flooding attacks against AD infrastructure. More interestingly, it has been shown that keeping business information secret can be beneficial, as it can provide incentives to firms to invest in innovative infrastructure and stay ahead of their competitors and contracting partners [30, 31].

We do not propose that ADs publicly disclose sensitive information. However, if evidence shows that an AD is involved in obscure network practices, it will have to provide such information to the regulator, so that the latter can argue about the AD's practices.

But will a transparency mechanism *implicitly* disclose sensitive AD information? We have proposed above (§3.2) that each AD exports enough information to on-path entities so that they can estimate network performance between distinct AD entry/exit point pairs for particular aggregates. If these entities combine their estimates for a given AD, will they be able to reverse-engineer sensitive AD information? In theory, such information can already be inferred by network-tomography techniques that take as input network-path measurements and infer network topology [32, 33] and link latencies [34]. In practice, tomography requires an impractically large number of vantage points for performing the network-path measurements. In our context, however, an AD will by itself provide the necessary input. Will transparency make it easier for tomography techniques to violate AD secrecy?

The answer depends on how ADs aggregate the exported information. If the exported information enables accurate estimation of network performance between distinct AD entry/exit point pairs, then we indeed expect AD secrecy to suffer. But if the exported information only enables accurate estimation of *aggregate* network performance between *sets* of AD entry/exit point pairs, then we expect AD secrecy to be preservable. So, how exactly should an AD aggregate the information it exports such that the relevant entities can compute useful performance estimates but not reverse-engineer the AD's internal topology and infrastructure properties?

Transparency & anonymity networks. Is it possible to have network transparency without compromising the functionality of anonymity networks?

Anonymity networks aim to hide user identities from curious/untrusted destinations and observers. The typical approach, introduced by Tor [35], is to have each client communicate with each destination through a circuit, constructed in such a way that no on-path entity – other than the client itself – can reconstruct the entire path from client to destination.

An adversary can deanonymize clients by correlating packet-timing information collected at circuit entry and exit points (and inferring which client is talking to each destination) [36, 37, 38]. In practice, this is only achievable by powerful adversaries, who either have presence at multiple network locations or can launch routing attacks to attract anonymized traffic [39]. In our context, however, there will be no need for such strong presence or routing attacks, as an AD will by itself provide the necessary timing information. Will transparency make life easier for eavesdroppers and censors?

Once again, the answer depends on how ADs aggregate the information they export. If the exported information reveals the presence of specific traffic at specific AD entry/exit points and at specific points in time, then we indeed expect anonymity to be compromised. But if the exported information only reveals the presence of each specific traffic unit at *any one of a set* of AD entry/exit points at a specific point in time, then we expect anonymity to be preservable. How exactly to aggregate exported information to prevent specific deanonymization attacks is an open question.

4. WILL TRANSPARENCY HELP?

We have argued that the technical community should focus on network transparency instead of network neutrality. But will network transparency really help promote a healthy ISP market? And how does a network-transparency mechanism relate to current regulations? We address both points in this concluding discussion.

Effectiveness of transparency. In theory, information disclosure is an easy means of overcoming market failures created by information asymmetries [40]. In practice, consumers often pay less attention to disclosed information than theory predicts [41, 42]. We believe that the effectiveness of our transparency mechanism will depend on how information is presented to end users [43]. We expect that network transparency will be most effective when the information exported by ADs is first processed by specialized intermediaries – regulators, neutrality watchdogs, news sites – which then present the information to end users in an aggregate comparison. Such a system could benefit from the insights of a booming literature at the intersection of economics, psychology, and law, which analyzes how simplifying, standardizing and personalizing information, as well as introducing information intermediaries can increase the effectiveness of information-disclosure regimes [41].

One additional strength of a network-transparency mechanism is that it would benefit not only individual end users, but also commercial users (e.g., content providers) and network providers who are interested in whether their peering partners adhere to their contractual promises. Such entities will have both stronger incentives to analyze the exported information and identify indications of neutrality violation, and easier access to regulators and courts to make their case.

Transparency and neutrality regulation. The idea of introducing network-neutrality regulations has been in the air for about 15 years [44, 45]. After long political and legal battles, the U.S. Federal Communications Commission (FCC) imposed its latest version of network-neutrality obligations on ISPs in June 2015 [46], while European rules on network neutrality are still under development. The obligations imposed by the FCC include prohibitions of blocking or throttling traffic, paid traffic prioritization, and interference of communication between edges of the network.

Given the importance current Internet regulation puts on network neutrality, how does a call for network transparency fit into this picture? In our opinion, reshifting the debate from network neutrality to network transparency could not only improve the technical discussions; it could also focus the policy debate on an often-overlooked aspect of network-neutrality regulations: transparency rules. In this regard, the FCC's 2015 order mandates ISPs (for both wired and wireless access) to disclose all fees, data caps and data allowances, as well as to disclose packet loss information as a measure of network performance. In the European Union, Articles 20 and 21 of the Universal Service Directive include similar provisions [47].

Our call for increased transparency is therefore supported by the current network-neutrality regulations of the FCC and the European Union. However, our proposal goes further than transparency rules in the U.S. and Europe. First, under our proposal, information on network state and management practices would be provided in a trustworthy manner. Consumers would not have to rely on cheap talk by ISPs. Second, under our proposal, the information provided would be more fine-grained. The higher quantity and better quality of information could increase competition between ISPs, thereby reducing the importance of network-neutrality regulations.

Third, our approach would not only benefit consumers; it could also assist courts, regulators and policy makers who would receive verifiable information from a transparent network layer. The increased data availability could hopefully lead to better-informed legal decisions and policy judgments in antitrust, communications, contract, tort, and criminal law. Finally, not only ISPs would disclose information about their network operations; rather, backbone providers would ideally also disclose information (subject to the secrecy considerations discussed in Section 3.3). This could enable novel service level agreements – e.g., between different ADs – to be contractible and enforceable.

As this discussion demonstrates, our proposal is in line with the emerging regulatory framework on network neutrality. Under both our proposal and current network-neutrality rules, a detailed technical definition of what constitutes a network-neutrality violation is not necessary. Rather, our proposal supports the regulatory push on both sides of the Atlantic for increased transparency, and argues that the solution to network neutrality concerns lies in technical and legal means to increase transparency in the marketplace for Internet services.

5. REFERENCES

[1] "EU Files Formal Antitrust Charges Against Google," http://on.wsj.com/1LZyqLO, Apr. 2015.

[2] "Inside the U.S. Antitrust Probe of Google," http://on.wsj.com/1GfEcSs, Mar. 2015.

[3] "AT&T Faces Formal FCC Complaint for Blocking Cellular FaceTime Use," http://bit.ly/1JYNxpt, Sep. 2012.

[4] "FCC Fines Telecom that Blocked Vonage VoIP Calls," http://bit.ly/1MokIA4, Mar. 2005.

[5] "Netalyzr Reveals ISPs Hijacking Users' Web Search Queries," http://bit.ly/1K2dSA3, Aug. 2011.

[6] N. Freed, "Behavior of and Requirements for Internet Firewalls," RFC 2979, IETF, 2000.

[7] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," RFC 2475, IETF, 1998.

[8] W. M. Petullo, X. Zhang, J. A. Solworth, D. J. Bernstein, and T. Lange, "MinimaLT: Minimal-latency Networking Through Better Security," in *Proc. of ACM CCS*, 2013.

[9] "Deutsche Telekom: A Software-Defined Operator," http://ubm.io/1RzUFfZ, Jun. 2013.

[10] "Netflix Performance on Verizon and Comcast Has Been Dropping for Months," http://bit.ly/1URc8zR, Feb. 2014.

[11] J. Sommers, P. Barford, N. Duffield, and A. Ron, "Accurate and Efficient SLA Compliance Monitoring," in *Proc. of ACM SIGCOMM*, 2007.

[12] Z. Zhang, O. Mara, and K. Argyraki, "Network Neutrality Inference," in *Proc. of ACM SIGCOMM*, 2014.

[13] M. Dischinger, M. Marcon, S. Guha, K. P. Gummadi, R. Mahajan, and S. Saroiu, "Glasnost: Enabling End Users to Detect Traffic Differentiation," in *Proc. of USENIX NSDI*, 2010.

[14] P. Kanuparthy and C. Dovrolis, "Diffprobe: Detecting ISP Service Discrimination," in *Proc. of IEEE INFOCOM*, 2010.

[15] Y. Zhang, Z. M. Mao, and M. Zhang, "Detecting Traffic Differentiation in Backbone ISPs with NetPolice," in *Proc. of ACM SIGCOMM*, 2009.

[16] G. Lu, Y. Chen, S. Birrer, F. Bustamante, C. Y. Cheung, and X. Li, "End-to-End Inference of Router Packet Forwarding Priority," in *Proc. of IEEE INFOCOM*, 2007.

[17] N. Handigol, B. Heller, V. Jeyakumar, D. Mazières, and N. McKeown, "I Know What Your Packet Did Last Hop: Using Packet Histories to Troubleshoot Networks," in *Proc. of USENIX NSDI*, 2014.

[18] K. Argyraki, P. Maniatis, O. Irzak, S. Ashish, and S. Shenker, "Loss and Delay Accountability for the Internet," in *Proc. of IEEE ICNP*, 2007.

[19] J. Wang, S. Lian, W. Dong, Y. Liu, and X.-Y. Li, "Every Packet Counts: Fine-Grained Delay and Loss Measurement with Reordering," in *Proc. of IEEE ICNP*, 2014.

[20] R. R. Kompella, K. Levchenko, A. C. Snoeren, and G. Varghese, "Every Microsecond Counts: Tracking Fine-grain Latencies with a Lossy Difference Aggregator," in *Proc. of ACM SIGCOMM*, 2009.

[21] B. Barak, S. Goldberg, and D. Xiao, "Protocols and Lower Bounds for Failure Localization in the Internet," in *Proc. of EUROCRYPT*, 2008.

[22] T. Zseby, M. Molina, N. Duffield, S. Niccolini, and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection," RFC 5475, IETF, 2009.

[23] N. G. Duffield and M. Grossglauser, "Trajectory Sampling for Direct Traffic Observation," *IEEE/ACM Trans. Netw.*, 2001.

[24] C. Pappas, R. M. Reischuk, and A. Perrig, "FAIR: Forwarding Accountability for Internet Reputability," in *Proc. of IEEE ICNP*, 2015.

[25] K. Argyraki, P. Maniatis, D. Cheriton, and S. Shenker, "Providing Packet Obituaries," in *Proc. of ACM HotNets*, 2004.

[26] A. Mizrak, Y.-C. Cheng, K. Marzullo, and S. Savage, "Fatih: Detecting and Isolating Malicious Routers," in *Proc. of IEEE DSN*, 2005.

[27] X. Zhang, C. Lan, and A. Perrig, "Secure and Scalable Network Fault Localization under Dynamic Traffic Patterns," in *Proc. of IEEE Security and Privacy*, 2012.

[28] B. Liu, J. T. Chiang, J. J. Haas, and Y.-C. Hu, "Coward Attacks in Vehicular Networks," *Mobile Computing and Communications Review*, 2010.

[29] K. Argyraki, P. Maniatis, and A. Singla, "Verifiable Network-performance Measurements," in *Proc. of ACM CoNEXT*, 2010.

[30] D. D. Friedman, W. M. Landes, and R. A. Posner, "Some Economics of Trade Secret Law," *Journal of Economic Perspectives*, 1991.

[31] S. Bechtold and F. Höffler, "An Economic Analysis of Trade-secret Protection in Buyer-Seller Relationships," *Journal of Law, Economics & Organization*, 2011.

[32] B. Eriksson, G. Dasarathy, P. Barford, and R. Nowak, "Efficient Network Tomography for Internet Topology Discovery," *IEEE/ACM Trans. Netw.*, 2012.

[33] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson, "Measuring ISP Topologies with Rocketfuel," *IEEE/ACM Trans. Netw.*, 2004.

[34] H. H. Song, L. Qiu, and Y. Zhang, "NetQuest: A Flexible Framework for Large-scale Network Measurement," in *Proc. of ACM SIGMETRICS*, 2006.

[35] R. Dingledine, N. Mathewson, and P. Syverson, "Tor: The Second-generation Onion Router," in *Proc. of USENIX Security*, 2004.

[36] N. Hopper, E. Y. Vasserman, and E. Chan-TIN, "How Much Anonymity Does Network Latency Leak?" *ACM Trans. Inf. Syst. Secur.*, 2010.

[37] P. Mittal, A. Khurshid, J. Juen, M. Caesar, and N. Borisov, "Stealthy Traffic Analysis of Low-latency Anonymous Communication Using Throughput Fingerprinting," in *Proc. of ACM CCS*, 2011.

[38] V. Shmatikov and M.-H. Wang, "Timing Analysis in Low-Latency Mix Networks: Attacks and Defenses," in *ESORICS*, 2006.

[39] Y. Sun, A. Edmundson, L. Vanbever, O. Li, J. Rexford, M. Chiang, and P. Mittal, "RAPTOR: Routing Attacks on Privacy in Tor," in *Proc. of USENIX Security*, 2015.

[40] G. A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970.

[41] G. Loewenstein, C. R. Sunstein, and R. Golman, "Disclosure: Psychology Changes Everything," *Annual Review of Economics*, 2014.

[42] B. van Schewick, "Network Neutrality and Quality of Service: What a Nondiscrimination Rule Should Look Like," *Stanford Law Review*, 2015.

[43] W. Lehr, E. Kenneally, and S. Bauer, "The Road to an Open Internet is Paved with Pragmatic Disclosure & Transparency Policies," http://ssrn.com/abstract=2587718, 2015.

[44] T. Wu, "Network Neutrality, Broadband Discrimination," *Journal of Telecommunications and High Technology*, 2003.

[45] M. A. Lemley and L. Lessig, "The End of End-to-End: Preserving the Architecture of the Internet in the Broadband Era," *UCLA Law Review*, 2001.

[46] Federal Communications Commission, "Protecting and Promoting the Open Internet," 80 Fed. Reg. 19738 2015.

[47] "Directive 2009/136/EC of 25 November 2009, amending Directive 2002/22/EC (Universal Service Directive)," O.J. EU L 337, Dec. 18, 2009, p. 11.