

Toward a Net Neutrality Debate that Conforms to the 2010s

Patrick Maillé, Gwendal Simon, Bruno Tuffin

► **To cite this version:**

Patrick Maillé, Gwendal Simon, Bruno Tuffin. Toward a Net Neutrality Debate that Conforms to the 2010s. IEEE Communications Magazine, Institute of Electrical and Electronics Engineers, 2016, 54 (3), pp.94-99. 10.1109/mcom.2016.7432154 . hal-01127958v2

HAL Id: hal-01127958

<https://hal.inria.fr/hal-01127958v2>

Submitted on 7 Dec 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Toward a Net Neutrality Debate that Conforms to the 2010s

Patrick Maillé[†], Gwendal Simon[†], and Bruno Tuffin[‡]

[†] Telecom Bretagne (Institut Mines-Telecom), France

[‡] Inria, France

September 5, 2015

Abstract

Network neutrality has been topic of discussion for the past 25 years, with current legislation/regulation in the US and Europe targeting the ISPs or “common carriers”. But the reality of the Internet in the 2010s is that various actors contribute to the delivery of data, with sometimes contradictory objectives.

In this paper, we highlight the fact that neutrality principles can be bypassed in many ways without violating the rules currently evoked in the debate. For example via Content Delivery Networks (CDNs), which deliver content on behalf of content providers for a fee, or via search engines, which can hinder competition and innovation by affecting the visibility and accessibility of content.

We therefore call for an extension of the net neutrality debate to all the actors involved in the Internet delivery chain. We particularly challenge the definition of net neutrality as it is generally discussed. Our goal is to initiate a relevant debate for net neutrality in an increasingly complex Internet ecosystem, and to provide examples of possible neutrality rules for different levels of the delivery chain, this level separation being inspired by the OSI layer model.

1 Introduction

Net neutrality is “the principle that Internet service providers should enable access to all content and applications equally, regardless of the source, without favoring or blocking particular online services or websites” (e.g., Oxford Dictionary). This *universality* principle for the Internet has been present for a long time in the US. In the 1996 *Telecommunications Act* [3], the Federal Communications Commission (FCC) declared Internet Service Providers (ISPs) as “common carriers” providing a public service. In February 2015, it voted to regulate broadband Internet service as a public utility. The same principle was adopted by the European Parliament in 2014 with a vote that restricts the

ability of ISPs to charge specific Service Providers (SPs) [2]. This principle is often translated as guaranteeing that all packets are treated equally at each intermediary step.

The network neutrality debate started in the 1990s, amongst growing concerns about the business models of network operators, which have to deal with several challenging trends, including the increase of traffic volume, the growing traffic asymmetry between operators, and the fast decrease of transit prices. Recently, spectacular disputes between major Internet actors¹ as well as vehement reactions from user associations and governments² have drawn attention to that debate. Arguments and discussions on possible regulation of the Internet resulted in refined definitions of network neutrality [6, 10, 12].

In this paper, we show that the current focus on ISPs behavior is too restrictive in the era of *cloud*-based content delivery. Indeed, the regulatory bodies in the US and Europe—those most active with regard to neutrality—focus on the ISP/internetwork level, but the value chain in the Internet is not restricted to only ISPs and Content Providers (CPs). In particular, web portals (especially search engines) and Content Delivery Networks (CDNs) have become key components in the delivery chain, while they surprisingly remain absent from the debate. Despite some recurring criticism [11], search engines have not been forced to conform to any universal policy yet; similarly, to the best of our knowledge, the only mention of CDNs in official net neutrality reports is from the Norwegian regulator, according to which “*the ordinary use of CDN servers is not a breach of net neutrality*” [9]. Yet, we show in Section 3 that the presence of a CDN, or biases in search listings, can be exploited against the fundamental principles of neutrality.

More generally, we think the net neutrality debate should be extended to all the actors involved in the Internet delivery chain. Our goal is not to discuss the validity of neutrality proponents and opponents, but rather to make a first step toward a global framework that would be more appropriate for the definition of public regulating rules (if any) in the era of the cloud and more generally of *information-centric networks*. We present our vision of such a framework in Section 4. To foster scientific activities, we highlight in Section 5 some topics, where the rigorous analysis of the supposedly “neutral” network scientists could be especially appreciated by decision makers.

Net neutrality is a highly politicized topic. We have tried to be politically neutral and scientifically objective in this article;³ however, we realize that some readers may interpret some of our statements as being politicized. If so, we assure that any perceived politicization is unintentional.

¹For example, Netflix witnessed a 25% reduction in traffic rates for Comcast users, before a deal was struck [1].

²See for example B. Obama’s declaration on neutrality in Nov 2014.

³The authors receive funding from Orange (a french ISP) for their research activities, but the content presented in this paper is solely the responsibility of the authors.

2 The Internet in the 2010s

2.1 Actors in Today’s Internet Delivery Chain

The Internet model usually considered when discussing the relevance of net neutrality is a chain of three actors:

users – ISPs – Content/service providers.

ISPs appear as the centerpiece of the debate, being seen as the main intermediary in the delivery chain. But many “newcomers” are now in the picture [8], as illustrated in Figure 1: In addition to the three aforementioned actors,




Figure 1: The main categories of actors in the Internet ecosystem; arrows represent a provider-customer (seller-buyer) relationship.

transit networks often act as intermediaries when users and contents are hosted by different ISPs; *CDNs* play a key role by storing content closer to users, thereby reducing transit costs and improving performance for that content⁴; *device builders* may introduce biases through the features of their products (possibly colluding with some other actors); and *search engines* (seen as service providers) directly affect the accessibility (visibility) of content. *Regulatory bodies* therefore face the delicate task of defending fairness and universality principles in this complex ecosystem, where ever-changing technical and business conditions prevent (or considerably complicate) analysis and comparison.

⁴The use of CDNs can be seen as against neutrality principles, since they offer an improved delivery for a fee, something new arrivants in the market may not be able to afford. We do not address this question here.

2.2 An Intricate Delivery Chain

The success of a given CP partly depends on its interactions with other actors. For example, its *visibility* is strongly affected by the behavior of search engines, which leaves space for biases that we point out later. Another example, the *engagement of end-users* for a CP is impacted by the delivery performance [4], which is usually under the responsibility of cloud providers and CDNs. Here again, non-neutral behaviors from those actors impact the CP.

The actors that have direct economic relationship negotiate Service Level Agreements (SLAs). In the case of content delivery, the intricate relationships between actors yield *chains* of SLAs, which are based on network-oriented metrics such as the throughput and the ratio of packet losses. But inferring the Quality of Experience (QoE) from a chain of network metrics, potentially cascading, is a difficult task. Most CPs struggle both to assess the QoE of their customers and to identify the failing actor when the quality of delivery is below expectations.

This complex chain of content delivery and the difficulties to assess the performance of the involved actors potentially leave some room for intermediaries to favor some CPs against others, thereby violating the principles of neutrality. We show later that a packet-level interpretation of neutrality can unfortunately not prevent intermediaries from biasing the competition among CPs.

2.3 Existing Protection Tools

Net neutrality opponents often argue that existing legal frameworks already protect fair competition, and thus no additional regulation is needed. This is the case in the US with the *antitrust law* and the associated Federal Trade Commission (FTC). Such antitrust arsenals do not however address some of the points that are frequently raised by net neutrality proponents.

First, the antitrust framework acts reactively and on a per-case basis although the neutrality regulation looks for pro-active solutions. With the antitrust framework, an actor that reckons to suffer economically from the (non-neutral) behavior of another actor should file a complaint with the FTC to obtain redress. The framework requires lawsuits and judgements once a plaintiff has built a case. On the other hand, neutrality regulation aims at anticipating the possible problems and preventing them by specifying allowed and forbidden behaviors.

Second, focusing on competition corresponds to a partial view—since limited to market aspects—which ignores some other principles raised by net neutrality proponents, such as freedom of speech and innovation fostering (the latter involving equal access to resources regardless of the monetary capacity of CPs).

In our opinion, neutrality rules are essentially *complementary* to antitrust laws, and setting boundaries between the application field of antitrust laws and that of neutrality rules should be part of the debate as well.

3 Examples of Neutrality Breaches

As previously stated, the supply chain between content and users includes multiple intermediary actors, which act in a free market with reciprocal engagements based on SLAs. We show in the following that these intermediaries can generate significant biases in the network and create a breach in what represents network neutrality for most people. We focus on CDNs and search engines, which are two key actors in the supply chain.

3.1 Competition Biases due to CDNs

First, we consider the impact of a profit-driven CDN provider, by developing two scenarios depicted in Figure 2. Both scenarios model the economic reality of the traditional interactions between ISPs, CDNs, and CPs.

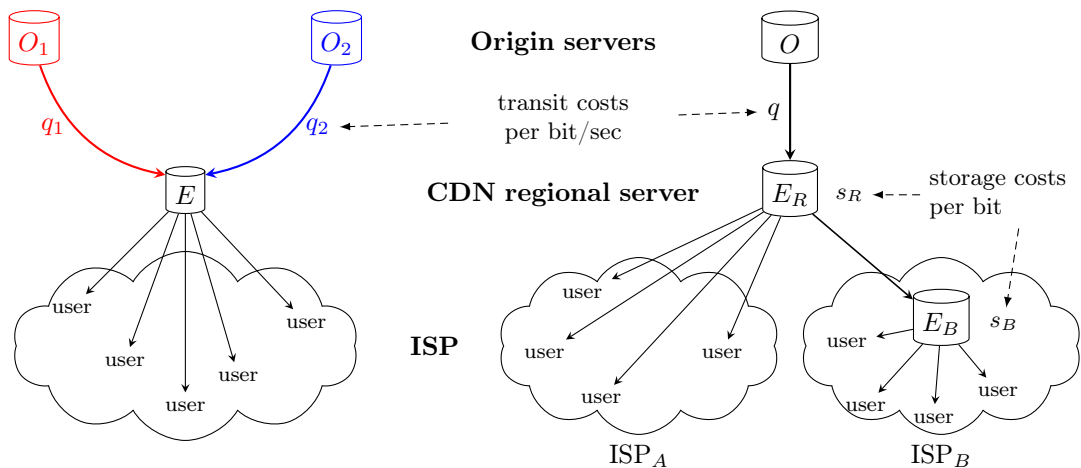


Figure 2: Two cases of neutrality breaches due to the presence of a CDN. On the left, two content providers, each with its own origin server, compete for the storage resources of a regional CDN edge server [7]. On the right, two ISPs are served by a shared CDN regional server, but one ISP (here ISP_B) installs its own edge server.

A CDN serves a large population of end-users by the means of edge servers, which are usually located near the Point of Peering (PoP) between ISPs. When the CDN receives a request, it has two handling options: either (1) fetch the requested content from the origin server, or (2) fetch the requested content from the edge server. Option 1 incurs transit costs (which can be low if the CDN also owns a transit network, but large otherwise), while Option 2 leads to improved user quality of experience (QoE)⁵, and no transit costs are incurred. Thus, in terms of cost and quality, Option 2 should always be chosen; however,

⁵Note that QoE is mainly affected for services requiring high bandwidth or low latency.

storage/cache and processing resources at E are finite. Hence, some requests must be handled by Option 1 (fetch from origin server). A rational CDN will maximize profits by optimizing the balance between Options 1 and 2, playing on caching decisions at E (or possibly on the number/capacity of edge servers).

Note that we focus on static content in these models for the sake of clarity. But the CDN models for interactive services, such as live streaming and cloud gaming, are not essentially different: the edge server should still reserve some hardware resources (mostly computing power here) to serve a subset of users. For example in live rate-adaptive video streaming systems, the edge server should both get multiple video representations from the origin server and manage the multiple concurrent requests for the videos.

3.1.1 Biasing the competition among Content Providers

Let us consider the scenario shown in the left part of Figure 2. That scenario is extensively discussed in [7]. To reach customers in a given ISP, two competing CPs subscribe to the CDN, which has to share the privileged resources in the edge server E among the CPs. The scientific literature comprehensively studies the performances of content replacement strategies in dynamic scenarios, but rarely integrates the economic parameters, namely here: *i*) the *transit costs* q_1 and q_2 , which can differ among CPs due to different paths to reach the ISP of interest; *ii*) the *charging policy* of the CDN, with prices per volume possibly differing depending on whether users are served from an edge or an origin server.

A profit-maximizing CDN stores in the edge server the content that yields the largest revenues. If both transit costs q_1 and q_2 are equal, and if the CDN charges both CPs the same way, then it should store the most popular content. However, if the transit costs of one CP (say CP_1) are higher than the other, then the interest for the CDN of storing CP_1 content in the edge server exceeds that of storing CP_2 content of comparable popularity, which both benefits CP_1 customers and harm CP_2 customers. In other words, **the internal costs of a CDN can distort competition.**

The situation gets even more problematic when the CDN charges CPs differently. In [7] we study a case based on real popularity measures, CP_2 being (on average) approximately five times more popular than CP_1 . Figure 3 displays the “satisfaction” (a normalized customer QoE measure) of both CPs, according to the ratio of content from CP_1 in the edge server.

The “standard scenario” situation corresponds to the “optimal” (for the CDN) sharing of the edge server when both CPs are charged identically. CP_2 being more popular, approximately five-sixths of the edge server is filled with CP_2 content, and the average satisfaction of CP_2 customers exceeds that of CP_1 customers. Hence **the CDN favors the dominant player although both content providers pay the same price.** Note however that the CDN strategy is the most efficient in terms of traffic reduction and overall satisfaction of end-users.

Now, assume that one of the CPs pays more for its requests fulfilled by the edge server. Two of the vertical lines correspond to two extrema where one CP




Figure 3: Satisfaction for both CP_1 and CP_2 customers according to the ratio of the edge server filled with content from CP_1

pays ten times the basic price (see [7] for more details). When CP_2 pays more (typically to reinforce its dominant position), the satisfaction gain of CP_2 is small but competitor satisfaction drops more significantly (from 0.94 to 0.89), possibly leading some CP_1 customers to churn. Alternatively, when CP_1 pays more (typically to increase its audience with better quality), the impact is less spectacular. To summarize, **the dominant player can leverage the CDN to harm its competitors although challengers cannot.**

3.1.2 Biasing the Competition Among ISPs

Consider now the scenario depicted in the right context of Figure 2, where a CDN serves two ISPs. Initially users of both ISPs (labelled A and B) get data either from the origin server O or from a regional “shared” edge server E_R . The size of the cache in E_R is chosen by the CDN according to the transit cost q from O to E_R (based on traffic) and the storage cost s_R at E_R (based on volume).

Now assume one ISP, say ISP_B , installs an edge server E_B within its network⁶ with a storage cost s_B , in order to improve the QoE for its users through an agreement with the CDN. The profit-maximizing CDN now has four ways to store content: in none of the edge servers (so only in the origin server), in E_R only, in E_B only, or in both E_R and E_B .

The appearance of the edge server E_B can create conflict, in particular when content previously stored in E_R gets only stored in E_B (and no longer in E_R): the QoE of ISP_A users is then degraded. Such a scenario is more likely to occur when the storage cost of E_B is significantly lower than E_R , or when ISP_B traffic is significantly bigger ISP_A traffic. In that case, the content stored in the regional “shared” edge server E_R is the “second” most popular in ISP_B , *i.e.*, the most popular among the content not stored in E_B . Thus, **agreements between CDNs and a given ISP can degrade the performances for users in another ISP.**

⁶In most cases, the ISP owns the edge server, which it rents to the CDN.

Figure 4 illustrates the impact of edge-server E_B on the overall satisfaction (still a normalized QoE measure) for users in both networks, using real movie request data over two years obtained from a leading French Video on Demand (VoD) service and including the request origin networks. We set ISP_B as the network with the highest traffic, and affect to ISP_A the average traffic from all other networks. The y -axis shows the ratio between the satisfaction with and without the edge server E_B ; the x -axis corresponds to different storage costs s_B .




Figure 4: Impact on user satisfaction of the creation of an edge server in ISP_B , computed as the ratio $\frac{\text{satisfaction with cache } E_B}{\text{satisfaction without } E_B}$.

We observe that the installation of an edge server in ISP_B has an (expected) positive impact on the satisfaction of ISP_B users, but also an (undesirable) negative impact on the satisfaction of ISP_A users. Observe that the latter exceeds the former for most of the storage cost prices on E_B : a powerful network operator does not need to implement any aggressive pricing policy to degrade the performances of competitors.

3.2 Search engines biasing the access to content

Search engines are the preferred way to discover and access content: according to ComScore, in 2014 about 21 billion searches have been treated each month in the US only. Search engines usually present their suggestions corresponding to keywords into two categories: 1) *sponsored results*, a list of advertisements related to the search and clearly defined as such, from which the engines make money when the links are clicked, and 2) *organic results*, a ranked list of links, believed to associate the most relevant items to the asked keywords since promotional results are assumed to be within sponsored links.

To increase revenues, search engines can be tempted to include in the organic results some results that are not among the most relevant but that can generate (direct and indirect) short-term revenues. A typical case is that of a content provider (called CP_1) vertically integrated with (i.e., owned by) the search engine. The engine is economically incentivized to rank that content higher [5].

Such temptation has raised a *search neutrality* debate, parallel to the one on network neutrality [8]. We do not target here any specific search engine nor claim one is biased: our goal is to show through a model that such biases can be motivated, and to investigate their impact on fairness among CPs.

The model designed in [5] analyzes the temptation to return “non-neutral” organic results. This model includes a long-term effect: end-users who are unsatisfied with the relevance of the organic results may stop using the search engine. The search engine objective is then to maximize its long-term revenue, which corresponds to the trade-off between the short-term revenue per visit on CP₁ and the longer-term number of visits on the search engine. The model helps to understand revenue-maximizing ranking strategies and to anticipate the impact of regulatory interventions for various scenarios. The principle is illustrated in Figure 5.




Figure 5: Search engine whose ranking policy produces an average relevance of results and an average gain. The number of visits (i.e., popularity of the engine) depends on the average relevance.

Table 1 illustrates, for some arbitrarily chosen numerical values, the differences in visit rates to CPs due to the revenue-optimal ranking, the resulting average relevance of the engine’s output, and the revenues it can generate, for different values of a parameter β , which represents the average revenue directly generated per visit to the search engine (via clicks on sponsored links, their assumed single source of revenue); the smaller β , the larger the bias because of larger (relative) incentives to create revenue through integrated CPs. In this scenario, the impact on the average relevance remains small (around 10%), but the impact on the visibility and the revenues of CPs is substantial, possibly threatening the survivability of non-integrated CPs.

4 Toward a Global Framework

Net neutrality proponents often put forward that every packet should be treated equally, which misses important aspects and does not prevent unfair situations. Indeed, that proposal focuses on the regulation of the third layer of the OSI model (the so-called network layer) although the end-users get the results of the

	Relevance	CP ₁ revenue	other CP revenue	CP ₁ visit rate	other CP visit rate
Neutral	0.635	0.028	0.0283	0.057	0.057
Non-neutral (for $\beta = 1$)	0.618 (-3%)	0.066 (+136%)	0.0243 (-14%)	0.112 (+96%)	0.049 (-14%)
Non-neutral (for $\beta = .5$)	0.592 (-7%)	0.084 (+200%)	0.0215 (-24%)	0.140 (+146%)	0.043 (-25%)
Non-neutral (for $\beta = .25$)	0.568 (-11%)	0.093 (+232%)	0.0193 (-32%)	0.158 (+177%)	0.039 (-32%)

Table 1: Impacts of ranking when the search engine owns CP₁.

policies of all intermediaries on the seven layers. As shown in the previous section, intermediaries can introduce significant biases while fulfilling this “neutral” condition at the third layer.

We therefore call for an extension of the debate to any class of actors having an influence in the way end-users consume content. Among them, CDNs, as well as the providers of data-centers, and search engines.

The need is then for a definition of neutrality that would be applicable to all actors, whose enforcement would prevent the behaviors pointed out in Section 3. Such a definition would need to go beyond the packet level, and may involve a classification of all actors into categories, for each of which specific principles should be followed. The regulation would then be specific to each level; below are a few examples:

- at the ISP level, no differentiation among packets according to the origin, destination, service, or price paid (the common suggestion);
- at the CDN level, caching of only the most popular content, independent of the provider, the type of content, or any monetary aspect;
- at the search engine level, ranking results based only on relevance to the query (and possibly the user), independent of the result page owner, the type of query, or possible payment.

Our purpose in the remaining of this section is *not* to put forward a clear framework for a new neutrality definition, ideally robust to network evolution, but rather to launch a discussion in the community towards that goal. To start that discussion, we hereafter pose a few questions and provide limited and probably flawed suggestions, which will (hopefully) ignite the debate.

4.1 Where should neutrality be applied?

The whole topology of the Internet should be abstracted to capture all the relevant actors that may introduce inappropriate biases. This effort should not be restricted to the physical infrastructure, but it should also consider the actors

that contribute to the consumption of content by any means. Typically, search engines are not captured by the current framework targeting packet level and ISPs. This effort should also consider the new actors that are expected to implement and integrate their network management services into the network devices compatible with Network Function Virtualization (NFV) standards. More generally, all actors that transform the content and are part of the delivery should be considered.

A definition of neutrality should study behaviors (actor actions) *at each actor level*: this would imply defining levels for the different types of actions on content, in a way comparable to the OSI reference model. We would for instance need to characterize (in terms of level) the services provided by CDNs and by search engines. Players at a given level/layer should then be neutral in terms of protocols applied at this level, by avoiding any kind of differentiation through the compliance of level-specific rules, extending the current rules defined for ISPs at the network/transport layer.

Note that, unfortunately it is not easy to identify the level at which an actor plays. For example in France, Skype is considered by the regulator as a virtual network operator, which thus competes with ISPs, and not as a service (or application) provider.

4.2 What notion should neutrality be focusing on?

What is the goal of neutrality, what do we want to preserve? The current focus is about equal treatment: no differentiation would be allowed whatever the source, application or service; hence some induced notion of *fairness*.

Among the underlying principles is the universality of access. This is a baseline (as it only focuses on blocking), but we need stronger constraints in a neutrality definition. Slowing down some applications, or favoring others, is generally refused by neutrality proponents and does not conform to the packet-based equal treatment proposal. On the other hand, the supposed equal treatment at the packet level can be criticized since most Internet applications use TCP at the transport layer, which leads to unfair qualities of service offered since the TCP throughput varies with the user/provider distance (it is indeed inversely proportional to the square root of the round trip time). This characteristic of TCP can harm innovation in regions with bad network quality, even with perfectly equal packet treatment.

Defining the “no-differentiation” principle when talking about CDNs for instance seems subtle, because CDNs precisely differentiate services when selecting what content to cache. The notion of neutrality could be in terms of avoiding money-based differentiation: fairness considerations would help decide what to cache, but there can be several potential interpretations:

- Should any packet have the same probability to be cached? We do not see any reason to do that.
- Should it rather be the most popular content which is cached (what is usually assumed natural to limit the load on the network)? But couldn't we

reason differently, by saying that all content providers should be equally cached, or at the end should offer the same quality of experience (something not done for the traditional neutrality principle though)?

It seems reasonable to set a “neutrality rule” requiring CDNs to cache the most popular content, *independently* of the source *identity*. But the focus could also be on reducing the overall network load by favoring the caching of distant content (that contributes more to the load since using more links) in addition to considering popularity.

In order to (partially) solve these questions, we suggest to define fairness, like neutrality, at the layer the actor plays (related to the question “where?”). This type of definition could also encompass search engines and ISPs. The basic idea is to have an “equal treatment”, but one could also look at other notions, such as social welfare, user welfare, or even service welfare (note however that such efficiency measures are in general incompatible with fairness, an argument often used by non-neutrality defendants). The goal is to determine whether a behavior harms the desired goal at the layer it is operating.

4.3 When should it be controlled?

Should we ensure that any newcomer has a particular treatment? Regarding the current strategies of CDNs, the CPs in their initial stages are disadvantaged against the incumbents. However, favoring these companies would be in disagreement with the current neutrality rules of equal treatment. It is therefore a slippery slope, opening the door to “mission creep” situations where new types of biases would defeat the initial goal of achieving neutrality.

With the current focus of the debate, remark that this question is not really a problem, but it is if we try to have a broader view of the Internet network.

5 Potential next steps

To conclude this article, we would like to present a few issues to be tackled by the community.

Taking into account economics in future Internet studies In general the question of actor profit is rarely central in the literature related to network and service management. Computer science and electrical engineering scientists generally aim at maximizing the efficiency of their proposal regarding technical objectives and not economic ones. Typically, to our knowledge the literature related to Content Centric Network (CCN) does not deal with the economics of actors although, in our opinion, the revamping of the Internet cannot be seriously studied without taking economic factors into account.

Evaluating the impact of possible regulatory policies The role of a regulating agency is to recommend policies that guarantee widely accepted principles. As previously mentioned, the stated principles are not (yet?) well defined

when it comes to neutrality. It is up to the scientific community to provide rigorous studies, as unbiased as possible, about potential policies, their impact on the considered actor level, and also the possible impact on the overall delivery chain.

Defining and studying new fairness models Our brief discussion in Section 4 highlights the lack of a better definition of fairness and in particular the subtle equilibrium between the guarantee of fair competition among actors pursuing similar objectives (the notion of *actor level*) and the preservation of the motivation for investment toward satisfying end-users better.

References

- [1] J. Brodtkin. Netflix performance on Verizon and Comcast has been dropping for months. *Ars Technica*, Feb 2014.
- [2] European Commission. Proposal for a regulation of the european parliament and of the council laying down measures concerning the european single market for electronic communications and to achieve a connected continent, and amending directives 2002/20/ec, 2002/21/ec and 2002/22/ec and regulations (ec) no 1211/2009 and (eu) no 531/2012. <http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=COM:2013:0627:FIN:EN:PDF>, 2013.
- [3] Federal Communications Commission. Telecommunications act of 1996. <http://transition.fcc.gov/telecom.html>, 1996.
- [4] S. Shunmuga Krishnan and Ramesh K. Sitaraman. Video stream quality impacts viewer behavior: Inferring causality using quasi-experimental designs. *IEEE/ACM Trans. Netw.*, 21(6):2001–2014, 2013.
- [5] P. L’Ecuyer, P. Maillé, N. Stier-Moses, and B. Tuffin. Revenue-Maximizing Rankings for Online Platforms with Quality-Sensitive Consumers. Technical report, Inria, February 2014. <https://hal.inria.fr/hal-00953790>.
- [6] T.M. Lenard and R.J. May (Eds.). *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [7] P. Maillé, K. Pires, G. Simon, and B. Tuffin. How Neutral is a CDN: An Economic Approach. In *Proceedings of the 10th International Conference on Network and Service Management (CNSM)*. IEEE, November 2014.
- [8] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
- [9] Norwegian Communications Authority. Norwegian report on content delivery networks (cdn). NKOM report, May 2012.

- [10] A. Odlyzko. Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. *Review of Network Economics*, 8(1):1–21, 2009.
- [11] F.A. Pasquale and O. Bracha. Federal Search Commission? Access, Fairness and Accountability in the Law of Search. *Cornell Law Review*, September 2008.
- [12] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecom. and High Tech.*, 2003.

Biographies

PATRICK MAILLÉ (patrick.maille@telecom-bretagne.eu) graduated from Ecole polytechnique and Telecom ParisTech in 2000 and 2002, respectively. Since then, he has been with the Networks, Security, Multimedia department of Telecom Bretagne, obtaining a PhD in 2005. His research interests are in all economic aspects of telecommunication networks, from pricing schemes at the user level, to auctions for spectrum and regulatory issues (net neutrality, search neutrality).

GWENDAL SIMON (gwendal.simon@telecom-bretagne.eu) is graduated from University Rennes 1 (France). During his PhD (obtained in 2004), he worked at Orange Research Labs. From 2004 to 2006 he was a researcher at Orange Labs. Since 2006, he has been Associate Professor at Telecom Bretagne. He was a visiting researcher at University of Waterloo from September 2011 to September 2012. His research interests include multimedia delivery systems (video and gaming) and network management.

BRUNO TUFFIN (bruno.tuffin@inria.fr) received his PhD degree in applied mathematics from the University of Rennes 1 (France) in 1997. He is currently a senior researcher at INRIA (the French National Institute for Research in Computer Science and Control). His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems and telecommunication-related economical models.