



HAL
open science

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, 2015. hal-01127958v1

HAL Id: hal-01127958

<https://inria.hal.science/hal-01127958v1>

Submitted on 9 Mar 2015 (v1), last revised 7 Dec 2015 (v2)

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

March 8, 2015

Abstract

Network neutrality has been topic of discussion for the past 25 years. In this position paper, we aim at highlighting that the issue is somewhat bypassed with the use of often forgotten actors such as Content Delivery Network (CDN) providers. Our goal is then to ignite a discussion on whether an extension of the debate is required. We particularly start to extend the usually discussed neutrality definitions, in order to keep them relevant in an increasingly complexifying Internet ecosystem.

1 Introduction

The *network neutrality debate* has started in the 1990s. The background context is the growth of concerns about the business models of network operators, which have to deal with a set of trends (including the increase of traffic volume, the growing traffic asymmetry between operators, and the fast decrease of transit prices). In the recent years, a series of spectacular disputes between major Internet actors¹ as well as vehement reactions from user associations and governments² have drawn attention on the net neutrality debate. An increasing number of arguments and discussions on what the rules of the Internet should be have resulted in the introduction of definitions of network neutrality [5, 11]. A comprehensive historical review³ is available at [8].

A widely accepted definition for 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”. This principle of ensuring the *universality* of the Internet was already

¹For example, Comcast reduced the Netflix speed by 25% during their 2014 dispute, before a deal was struck.

²see for example <http://is.gd/qhxwV3> for B. Obama’s declaration on neutrality in Nov 2014, or <http://is.gd/Ci8vH5> for a European political decision of April 2014

³See also <http://whatisnetneutrality.org/timeline>

present in the 1990s, when the US regulator, the Federal Communications Commission (FCC), declared Internet Service Providers (ISPs) as “common carriers” providing a public service in its 1996 *Telecommunications Act* [3], and voted in February 2015 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 some specific Content Providers (CPs) [2]. This principle is often translated as guaranteeing that all packets are treated equally at each intermediary network equipment.

In this paper, we show that the current focus on ISPs behavior with respect to services and even packets is too restrictive in the era of *cloud*-based delivery of content. As we will show in Section 2, the value chain in the Internet is not restricted to only ISPs and CPs. In particular, web portals (in particular search engines) and Content Delivery Networks (CDNs) have become key components in the chain of delivery. However they are surprisingly still absent from the net neutrality debate. Despite some recurring critics [9], 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*”. Yet, we will show in Section 3 that the presence of a CDN or biases in search listings can be exploited against the fundamental principles of net neutrality.

More generally, we call for an extension of the net neutrality debate 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 give our vision of such a framework in Section 4. And to foster scientific activities, we highlight in Section 5 three topics, where the rigorous analysis of the supposedly “neutral” network scientists could be especially appreciated by decision makers.

2 Actors in Today’s Internet Delivery Chain

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

users – ISPs – Content/service providers.

This is the reason why ISPs are the center piece of the debate, as they are seen as the only intermediary in the delivery chain. But many “newcomers” are now in the picture [7], as illustrated in Figure 1.

In addition to the three types of actors above, *transit networks* often act as intermediaries when users and contents are not hosted by the same ISP; CDN play a key role by storing some content closer to users, thereby reducing transit costs and improving performance for that content; device builders may introduce biases through the features of their products (possibly colluding with

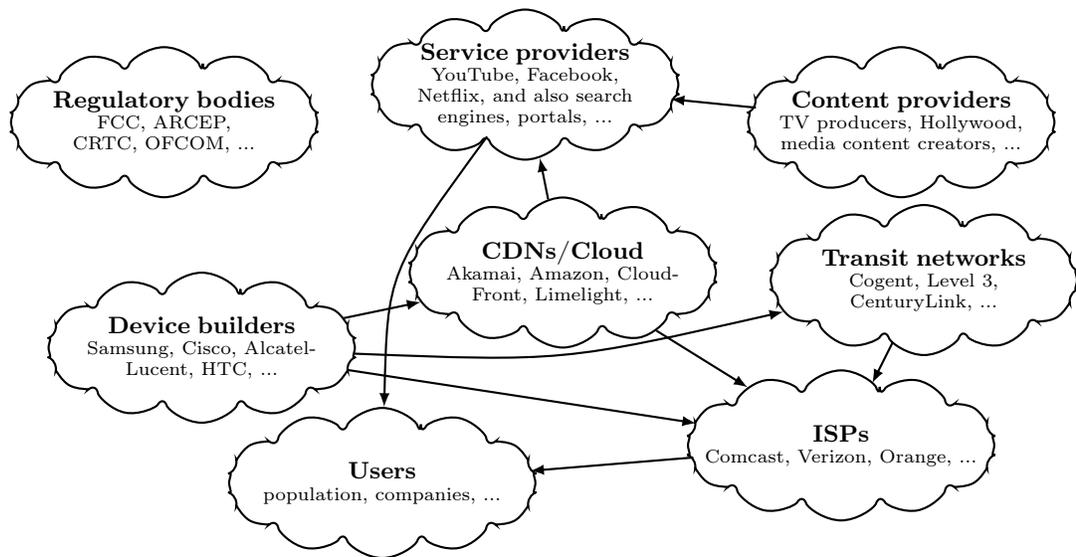


Figure 1: The main categories of actors in the Internet ecosystem; arrows represent a provider-customer (seller-buyer) relationship.

some other actors); and search engines (seen as service providers) directly affect the accessibility (visibility) of content to users. Regulatory bodies therefore face the delicate task of defending fairness and universality principles in this complex ecosystem.

3 Examples of Net Neutrality Breaches

In the following, we show that not only ISPs but most, if not all, of other types of intermediate between content and users can generate significant biases in the network and create a breach in what represents the network neutrality principle for most people. We particularly focus on two types of intermediate (excluding ISPs who are the traditional focus of the debate): CDNs and search engines, who are among the most prominent other actors of the supply chain.

We first focus on a rational (i.e., revenue-driven) CDN provider and its impact on the network. The two scenarios we develop hereafter are depicted in Figure 2. These are traditional models of the economic reality of the interactions between ISPs, CDNs, and CPs.

We then discuss the bias that results from search engines' revenue-maximizing strategies and show that the search neutrality debate is, or should be, integrated into the network neutrality one. Note that, in both cases, the intermediary network equipments treat all packets the same way, hence respecting the most widely accepted net neutrality rule.

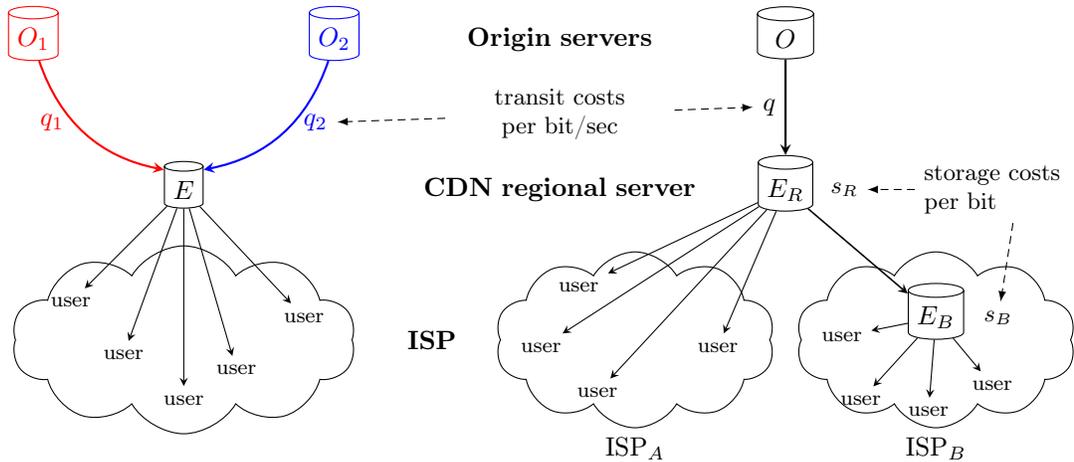


Figure 2: Two cases of neutrality breaches due to the presence of a CDN. On the left, two content providers, each of them using its own origin server, compete for the storage resources of a regional CDN edge server [6]. On the right, two ISPs are served by a shared CDN regional server, but one ISP (here ISP_B) installs its own edge server for its users.

3.1 Biasing the Competition among Content Providers

We consider the following scenario, discussed in [6], where two CPs compete to deliver content to their customers in a given ISP as shown in the left part of Figure 2.

In this simplified scenario, both CPs subscribe to the CDN service to reach their customers. The CDN has two options when it receives a request. If the data should be fetched from the origin server O_i , the cost is the transit cost q_i per unit of volume (which can be low if the CDN owns the transit network, but large otherwise). The second case is when the data are delivered by the CDN edge server E : the quality experienced by users is then better, and no transit costs are incurred. Unfortunately, resources are scarce at the edge server.

The CDN decides how to share the privileged resources of the edge server among both CPs. The scientific literature has abundantly studied the performances of content replacement strategies in dynamic scenarios, for example the Least Recently Used (LRU) replacement policy. However the economic parameters, which are the most significant for the CDN, have rarely been integrated in previous works. In particular, the revenues of the CDN depend on two parameters: *i*) the *transit costs*, which can differ among CPs (i.e., $q_1 \neq q_2$) due to different paths to reach the ISP of interest; *ii*) the *charging policy* of the CDN, which can charge the CPs a different price per unit of data volume according to whether users are served from the edge server or from the origin server.

A rational CDN stores in the edge server the content of each provider 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 the most profitable strategy is to store the content that is the most requested. In other words, the more popular a content is, the more likely it is to be stored in the edge server, and the better the Quality of Experience (QoE) for the end-users consuming it. 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 is both beneficial to CP_1 customers and detrimental to CP_2 customers. In other words, **the internal costs of a CDN can produce competition distortion**.

The situation gets even more problematic when the CDN charges CPs differently. In [6] we have studied a use case with real popularity measures for two different CPs, CP_2 being (on average) around five times more popular than CP_1 overall. To highlight the possible breach of net neutrality, we represent in Figure 3 the evolution of the “satisfaction” (a normalized customer QoE measure) of both CPs, according to the ratio of content from CP_1 in the edge server.

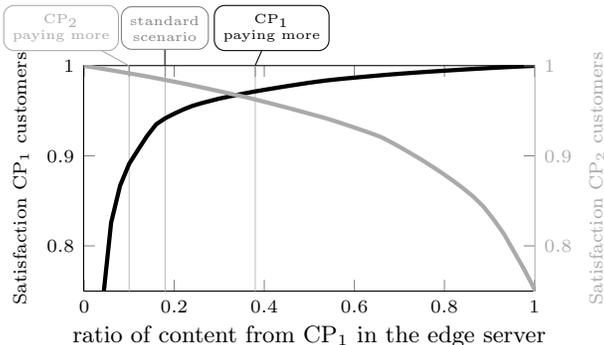


Figure 3: Satisfaction for both CP_1 and CP_2 customers according to the ratio of the edge server filled with content from CP_1

The vertical line labeled with the “standard scenario” corresponds to the “optimal” (most profitable for the CDN) sharing of the edge server when both CPs are charged the same way. As can be seen, since CP_2 is more popular, more than four-fifths of the edge server are filled with CP_2 content. The consequence is that the average satisfaction of CP_2 customers is better than for CP_1 customers. That is, **the CDN favors the dominant player although both content providers pay the same price**. Note however that the action of the CDN is the most efficient regarding both the traffic reduction and the overall satisfaction of end-users.

We explored scenarios when one of the CPs decides to pay more for the requests that are fulfilled by the edge server. The two vertical lines correspond to two extrema where one of the CPs pays ten times the basic price (see [6] for more details). Two observations can be made. First, when CP_2 pays more (typically to reinforce its dominant position), the CDN stores even more content from CP_2 in the edge server (around 90% of the storage capacity). The gain of satisfaction

for CP_2 customers is not significant; however the customer satisfaction of the competitor CP_1 drops more significantly (from 0.94 to 0.89). Such a drop can lead to churn for customers from CP_1 . Second, when CP_1 pays more (typically to increase its audience with better quality), the impact is less spectacular. Even in an aggressive monetary campaign (see the right vertical line), the gain in satisfaction for CP_1 customers and the drop in satisfaction for CP_2 customers are equivalent (both 0.03 in absolute value). These low satisfaction variations are less susceptible to generate drawing power at CP_1 and churn from CP_2 . To summarize, **the dominant player can leverage the CDN to harm its competitors although challengers cannot.**

3.2 Biasing the Competition Among ISPs

We consider now another scenario where a CDN serves two ISPs: ISP_A and ISP_B . The users of both ISPs get data either from the origin server O or from a regional “shared” edge server E_R . In this configuration, a rational CDN stores in E_R the content that is the most popular based on the traffic generated by users from both ISPs. Deciding the size of the cache E_R is among the core business choice of the CDN provider, with regard to both the transit cost q from O to E_R (based on the traffic) and the storage cost s_R at E_R (based on the volume).

The scenario becomes trickier when one of the ISPs, say ISP_B without loss of generality, decides to install an edge server E_B within its network with a storage cost s_B .⁴ This is depicted in the right context of Figure 2. The motivation behind installing an edge server in ISP_B is the improvement of the QoE of users of ISP_B . This installation corresponds to an agreement between the CDN and ISP_B . Overall, the CDN has now four ways to store a given 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 . A rational CDN again decides with regard to economic objectives, so both edge servers E_R and E_B are filled in order to maximize profit.

The switch from the previous configuration with no edge server E_B to the configuration with E_B can be a source of conflicts, in particular when a piece of content previously stored in E_R gets only stored in E_B (and no longer in E_R). The QoE of users from ISP_A is degraded due to the edge server E_B . Such a scenario is more likely to occur when the storage cost of the cache E_B is significantly lower than E_R or when the overall traffic generated by ISP_B is significantly bigger than the traffic generated by ISP_A . In this case, whatever the popularity of content for the users of ISP_A , the content that are stored in the regional “shared” edge server E_R are the “second” most popular in ISP_B , *i.e.*, the most popular among the content that are not stored in E_B . In other words, **agreements between CDN and a given ISP can degrade the performances for users in another ISP.**

To illustrate this claim, we show in Figure 4 the impact of the installation of edge-server E_B on the overall satisfaction (still a normalized QoE measure)

⁴In most cases, the ISP owns the edge server, which it rents to the CDN. It aims at both improving the QoE for its users and getting some revenues from the exploitation of a privileged location in the network.

for users in two networks, ISP_A and ISP_B . We use a real dataset, which we got from a leading French Video on Demand (VoD) service, where all the requests for movies have been recorded during two years, including the network origin of the requesters. We set that ISP_B is the network that generated the highest traffic, and affect to ISP_A the average traffic from every other network. On the y -axis, we compute the ratio between the satisfaction with and without the edge server E_B . On the x -axis, we set different storage costs for s_B , from an aggressive policy where the cost of storing content in E_B is only half the price the CDN pays for storing content in E_R , to a benefit-oriented policy where s_B is twice s_R .

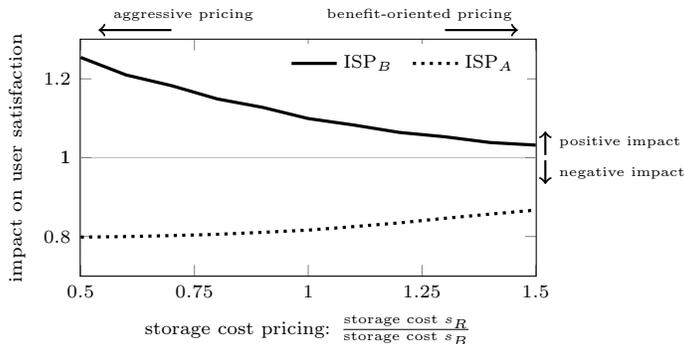


Figure 4: The impact on the satisfaction for users for the configuration with and without an edge server in ISP_B . The impact on satisfaction is computed as the ratio $\frac{\text{satisfaction with cache } E_B}{\text{satisfaction without } E_B}$. On the x -axis, the storage cost pricing at s_B as a function of the storage cost of s_R .

As can be seen in Figure 4, the installation of an edge server in ISP_B has *both* a positive impact on the satisfaction of users in ISP_B (which is expected), but also a negative impact on the satisfaction of users in ISP_A (which is not suitable). It is noteworthy to observe that the negative impact on ISP_A is larger than the positive impact on ISP_B for most of the storage cost prices on E_B . It means that a powerful network operator does not need to implement any aggressive pricing policy to degrade the performances of competitors.

3.3 Search engines biasing the access to content

Search engines have a major role in the Internet, as in many cases the way users choose what content to access. In the US only, about 20 billion searches are treated each month by those actors.⁵ Search engines usually decompose their propositions corresponding to keywords into two categories: 1) *sponsored results* corresponding to a list of advertisements related to the search, and clearly defined as such, from which the engines make money when the links are clicked,

⁵www.comscore.com (data from 2013)

and 2) *organic results*, corresponding to a ranked list of links, whose ranking is obtained from an algorithm believed to associate the most relevant documents to the asked keywords, since the promotional results are assumed to be part of the sponsored links.

The ranking of the organic links returned by the main search engines is now questioned by actors of the Internet and regulators who claim that economic arguments are incorporated into the algorithms. This gave rise to a debate (search neutrality) parallel to the one on network neutrality [7], Google being the centerpiece of the debate, as the most powerful search engine whose decisions impact the whole network economy. Google (and also Microsoft with Bing) is claimed to bias search results in order to favor its own other services, among which e-mail, maps, calendar, video, from which it can generate additional revenue.⁶ Several studies have illustrated this bias in search results toward own-content of search engines, see [1, 10].⁷ The experiments consist in comparing the results of search engines to the same request, and measuring the extent by which the positions of the contents affiliated to one of them differ. A representative illustration is an experiment realized in 2011, where 100 search queries are randomly selected from 20 millions of AOL users [10]. It shows that Google and Bing have statistically a clear tendency to bias toward their own content: for example, when Bing ranks some Microsoft content among its top 3 results, then in 99.2% of the cases neither of the other considered search engines ranks Microsoft content in their top 3 results. The same behavior is symmetrically observed for Google.

It is possible to design as in [4] a model where, corresponding to a request, a search engine applies an optimal ranking rule for the organic links taking into account the direct and indirect revenue of any click on the link, but also integrating the fact that at a longer time-scale, users will come to this search engine depending on their belief that the ranking is based on relevance. This type of model analyses the tradeoff between short-term revenue and longer-term number of visits, with sometimes opposite effects, when trying to maximize the long-term revenue as the product of the average number of visits and a (short-term) revenue per request. That principle is illustrated in Figure 5.

The model, detailed in [4], allows to understand the behavior of revenue-oriented search engines and to anticipate the impact of regulatory interventions in the case of various scenarios. A typical illustration is when a content provider (called CP_1) is vertically integrated with (that is, owned by) the search engine, such that the engine is economically incentivized to put it higher in the ranking. This scenario is studied in [4]. 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 representing the average revenue directly generated per visit to the search engine (thanks to clicks on sponsored links, what is thought to be their single source of revenue); the smaller β , the larger the bias because of larger (relative) incentives to create revenue

⁶<http://www.guardian.co.uk/technology/blog/2011/sep/21/eric-schmidt-google-senate-hearing>

⁷see also <http://www.benedelman.org/searchbias/>

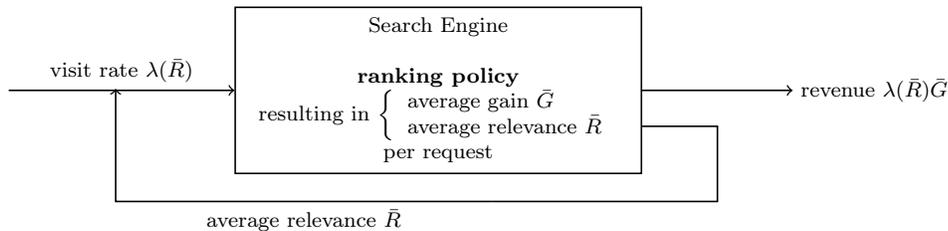


Figure 5: Search engine whose ranking policy produces and average relevance of results and an average gain. The number of visits (i.e., popularity of the engine) depends on the average relevance.

	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 a non-neutral ranking when the search engine is proposing its own content

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, which is likely to threaten the survivability of non-integrated CPs.

4 Toward a Global Framework

The underlying principle of neutrality is that of *universality*, meaning that everyone could reach any type of content using any type of application on any type of device. To reach this objective, net neutrality proponents often propose that every packet should be treated equally, which, as shown in the previous section, misses important aspects and does not prevent unfair situations from occurring. We therefore call for an extension of the debate to other actors than only ISPs and CPs, where any class of actors having an influence on the way end-users consume content would be considered. Among them, CDNs, as well as the providers of data-centers, and search engines.

The need is then for a new definition of neutrality that would be applicable

to all actors involved in the delivery chain, in particular preventing the behaviors pointed out in Section 3. Such a definition would need to go beyond the packet level, and may have to involve a classification of all actors into categories, for each of which specific principles should be followed. Our purpose in the remaining of this section is *not* to propose a clear framework for a new definition of neutrality, 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, hoping to ignite the debate. We will consider the impact of CDNs as an example on each of our assertions.

4.1 Who should be concerned by neutrality?

As argued from the beginning, the neutrality debate keeps focusing on ISPs because they were seen as the single intermediary, but much more exist now as highlighted in Section 2, such as CDNs or search engines whose potential impact has been illustrated. A general notion of neutrality should be addressed to all intermediaries, existing or future ones.

4.2 Where should neutrality be applied?

Neutrality has been analyzed at the packet level and at each node of the network. But this might miss important aspects:

- It misses parts of the “topology”, since the Internet cannot be seen as only a physical infrastructure. This will probably be even more true if Content Centric Networks (CCNs) take off, and with the Internet of Things. A currently touted example is search engines that are not captured by the current framework; nor are CDNs, at the edges of the network. The topology should be better abstracted to capture all elements, including the physical network but also the “transformation” of information.
- While relevant when dealing with ISPs at the transport level, talking about packets for some actors does not make real sense. If we wanted to extrapolate, this would mean that each small piece of data for CDNs should have the same chance to be cached. In order to encompass all activities, the packet level should not be the only level to be considered. In a definition of neutrality, behaviors (actions of one of the actors) should be studied *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; also, the recent efforts of the SDN/NFV (Software-Defined Networking/Network Function Virtualization) toward the abstraction of network topologies may give insights for the definition of new levels. We would for instance need a characterization (in terms of level) of the services provided by CDNs and by search engines. In this way players at a given level/layer should be neutral in terms of protocols applied at this level, by avoiding any kind of differentiation through the compliance of

level-specific rules. This typically corresponds to the concept applied to ISPs, at the network/transport layer in terms of packet differentiation.

Of course, defining the level at which an actor plays is not always easy: for example in France, Skype is considered by the regulator as a virtual network operator, competing with ISPs, not as a service or application provider; this may be subject to interpretation.

4.3 What notion should neutrality be focusing on?

What is the goal of neutrality, what do we want to preserve? The current principle is about equal treatment (as we pointed out at the packet level and at each node of the network), so that no differentiation is made whatever the source, application or service; hence some induced notion of *fairness*.

One of the underlying principles was to ensure universality, meaning that everyone could reach any type of content using any type of application on any type of device. This is a baseline (since forbidding blocking), but we need stronger constraints in the neutrality definition since slowing down applications, or differentiating between them, is not accepted by neutrality proponents and does not conform to the packet-based equal treatment proposal. On another hand, this latter principle to represent fairness could also be criticized, because most applications use TCP at the transport layer, a protocol leading to unfair offered qualities of service: the throughput is not the same if one provider is remote from users (it is indeed inversely proportional to the square root of the round trip time). This issue is often forgotten in the debate and may be at the detriment of innovation in poor countries with bad network quality. Proponents of neutrality seem to accommodate with that though; it seems to be seen as a “negligible” side effect.

Defining the “no-differentiation” principle when talking about CDNs for instance is particularly subtle, because the basic goal of CDNs is to differentiate service by choosing some content that will be cached while other will not be. The notion of neutrality has to be in terms of avoiding money-based differentiation. Fairness considerations could help decide what should be cached, 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 service 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” imposing CDNs to cache the most popular content, *independently* of the source *identity*. But it could also be allowed to focus 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 those 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 not compatible with fairness, an argument often used by non-neutrality defendants). The goal is to look at whether or not a behavior harms the searched goal at the layer it is operating.

Focusing on fairness, in terms of equal treatment as the required or expected goal of the actors and the associated rules/protocols, some example rules could be:

- 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, independently 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), independently of the owner of the result page, the type of query, or possible payments.

4.4 When should it be controlled?

Should we ensure that any newcomer has a particular treatment? With regard to CDNs, starting businesses will not have their content cached (assuming most popular content is cached), hence a disadvantage. Shall we prevent that? But that would be in disagreement with the current vision of neutrality even if favoring innovation. 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 other issues to be tackled by the community.

Taking into account economics in the 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 CCN does not deal with the economic of actors although, in our opinion, the revamping of the Internet cannot be seriously studied without taking into account the economic actors.

Evaluating the impact of possible regulator policies The role of a regulating agency is to propose policies that guarantee widely accepted principles. As previously stated, the said principles are not (yet?) well defined when it comes to net neutrality. It is up to the scientific community to provide some rigorous, and as unbiased as possible, studies about some possible 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 a fair competition among actors pursuing similar objectives (the notion of *actor level*) and the preservation of the motivation for investment toward a better satisfaction of the end-users.

References

- [1] B. Edelman and B. Lockwood. Measuring bias in “organic” web search. <http://www.benedelman.org/searchbias>, Jan 2011.
- [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] 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>.
- [5] T.M. Lenard and R.J. May (Eds.). *Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated*. Springer, 2006.
- [6] 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.
- [7] P. Maillé and B. Tuffin. *Telecommunication Network Economics: From Theory to Applications*. Cambridge University Press, 2014.
- [8] 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.

- [9] F.A. Pasquale and O. Bracha. Federal Search Commission? Access, Fairness and Accountability in the Law of Search. *Cornell Law Review*, September 2008.
- [10] J. D. Wright. Defining and measuring search bias: Some preliminary evidence. George Mason Law & Economics Research Paper 12-14, George Mason University School of Law, 2012.
- [11] T. Wu. Network neutrality, broadband discrimination. *Journal of Telecom. and High Tech.*, 2003.