Non-Neutrality With Users Deciding Differentiation: A Satisfying Option?

Anne Kieffer[1], Patrick Maillé[2], and Bruno Tuffin[3]

[1] *INSA*
Rennes, France
anne.kieffer@insa-rennes.fr
[2] *IMT Atlantique, IRISA, UMR CNRS 6074*
F-35700 Rennes, France
patrick.maille@imt.fr
[3] *Inria, Univ Rennes, CNRS, IRISA*
Rennes, France
bruno.tuffin@inria.fr

Abstract. The network neutrality debate has been raging worldwide for around fifteen years now. Our goal in this paper is to model and discuss a quite recent option which could be seen as a trade-off between neutrality and differentiation operated by Internet service providers (ISPs), and satisfy both ends of the world: differentiation potentially chosen by end users. By using a model from the literature, we compare the outcomes of three scenarios: neutrality, non-neutrality with differentiation decided by ISPs, and non-neutrality decided by users. We illustrate that, depending on network parameters, letting end users decide may end up as a fair and viable solution, and that non-neutrality imposed by ISPs is not necessarily bad for all actors.

1 Introduction

The success of the Internet is based on the notion of packets treated equally and routed through the network in a "best-effort" way. The *network neutrality debate* [18] started with Internet service providers (ISPs) claiming that they are supporting the infrastructure development and maintenance while content providers (CPs) send an increasing traffic and get an increasing share of the revenue generated by the network activity. ISPs were asking for the possibility to request payments for services and to differentiate traffic. This raised complaints from user associations and CPs arguing among other things that it would prevent innovation [7, 8, 10, 9, 14]. The debate has been raging, with neutrality rules defined and imposed by regulators worldwide (see as examples [4] for the FCC in the USA and [2] for BEREC in Europe). The debate is even exacerbated with the recent decision to repeal neutrality in the USA [5, 11].

There have been many models to analyze the pros and cons of neutrality and/or service differentiation; see among others [1, 3, 10, 13] and the references therein. The goal is in general to discuss (using tools from game theory) whether service differentiation can be beneficial or hurtful to the society, and whether introducing regulation would improve the outcomes.

It would indeed be interesting to find a trade-off satisfying all actors, CPs, users and ISPs. Quoting [16]:

> Rather than focusing on network behavior only, it may be more helpful to consider end-user choice as the principle for deciding whether a particular traffic management or other policy is reasonable if it cannot be readily justified as protecting the network against attack or abuse.

In other words, i) service differentiation is often seen as a restraint to innovation, but users should be able to choose to favor new services if they find them relevant, and innovative CPs could also gain from this; ii) with user-driven differentiation, the ISP stays "neutral" in the sense that it does not choose to differentiate; iii) such an approach would allow ISPs to operate some differentiation, release the necessity to invest on capacity and get a reasonable share of revenue.

The purpose of this paper is to apply a mathematical model to investigate whether the option of allowing differentiation, operated by users, is worthwhile with respect to neutrality or differentiation operated by the ISP. This is studied in the context of actors (here users, an ISP and CPs) making decisions to optimize their own interest. (Therefore the mathematical framework is that of non-cooperative game theory [10, 15].) We propose to use the model and results in [12] where users want to access two types of services, video (it was voice in [12] but we "update" it) and data, through a network modeled by an M/M/1 queue. We assume an infinity of potential *infinitesimal* users sensitive to price and delay and asking for service as soon as they get a non-negative utility: the user equilibrium notion is then called *Wardrop equilibrium* [17]. The three scenarios we compare are: a unique class of service (neutral scenario) for which packets are served in a first-come-first-served manner; a situation with two priority classes but for which the ISP decides which class should get priority, and the same two-priority classes but for which each user decides its service class. On top of that, the ISP optimizes its price(s) for service, anticipating the users' response, for the three scenarios. Our comparison of the outcomes shows that letting users decide the differentiation is a viable and balanced solution, since it allows both CPs to be served, and may also satisfy all actors with respect to a neutral situation. Similarly, letting the ISP differentiate may actually be beneficial for users, a conclusion consistent with other analyses in standard revenue management theory.

The remaining of the paper is organized as follows. Section 2 presents the basic model we are taking from [12], the relevant results therein, and the scenarios with the structure of the game that is played. Section 3 compares the outputs of the three scenarios to discuss if results are surprising, and conclusions and recommendations are made in Section 4.

2 Model

2.1 Basic model

We recall in this subsection the model and results first introduced in [12].

The model considers two classes of (infinitesimal) customers, say video (indexed by v) and data (indexed by d) users, each one generating packets with an average rate λ_d (resp. λ_v) per unit of mass of customers. If p is the per-packet price charged to users, the respective (per-packet) utilities of type-d and type-v users are

$$U_d(\mathbb{E}[D]) = u_d(\mathbb{E}[D]) - p \quad \text{with } u_d(y) = y^{-\alpha_d}$$
$$U_v(\mathbb{E}[D]) = u_v(\mathbb{E}[D]) - p \quad \text{with } u_v(y) = y^{-\alpha_v},$$

where $u_i(\mathbb{E}[D])$ for $i \in \{d, v\}$ is the willingness to pay for a packet transmission in the network if the expected delay is $\mathbb{E}[D]$. It is assumed that $0 < \alpha_d < \alpha_v$, meaning that video users are more sensitive to congestion than data users. The curves intersect to highlight that voice users give more value to small delays.

Let N_d (resp. N_v) be the number (or more precisely, mass since users are assumed infinitesimally small) of data (resp. video) users. It is assumed that there is a potential unlimited number of video and data customers coming in as soon as their utility is positive, or leaving if negative. The equilibrium notion in terms of the actual mass of customers is the Wardrop equilibrium [17], where users of a given type either do not use at all a class of service because of a non-positive utility, or use a class of service but have a zero utility (otherwise new users would enter or leave).

The network is represented by its bottleneck modeled as an M/M/1 queue. In the neutral case without differentiation,

$$\mathbb{E}[D] = \frac{1}{\mu - (\lambda_d N_d + \lambda_v N_v)}$$

where μ is the network service rate. In the case of two priority classes H for high and L for low, with packet rates (to be more clearly defined later) λ_H and λ_L,

$$\mathbb{E}[D_H] = \frac{1}{\mu - \lambda_H}$$
$$\mathbb{E}[D_L] = \frac{1}{(\mu - \lambda_H)(1 - (\lambda_H + \lambda_L)/\mu)}.$$

See [6] if a proof of those formulas is needed.

2.2 Scenarios and goal

We will consider three different scenarios:

- The neutral scenario for which the ISP proposes a single class of service;

- The non-neutral scenario for which the ISP decides which type of users will use which priority class;
- The non-neutral scenario for which each infinitesimal user will selfishly decide its priority class.

In each case, the decisions are taken in the following order:
1) The ISP first determines (revenue-maximizing) price(s).
2) Customers use the service or not: given the price(s) and the service policy, the mass of users of each type asking for service satisfies the Wardrop principle. Even if the ISP plays first, it plays anticipatively, taking into account the reaction of users.

We will compare the output of the game for the three types of actors: users, CPs and the ISP, in order to determine if a strategy should be favored by regulators, in particular the recent proposition to let users decide their service class. This type of comparison was not the purpose in [12]. More precisely, we are going to compare for the three scenarios: i) the ISP revenue to see if and why the ISP is pushing for a solution; ii) the CPs individual revenue $a_j \lambda_j N_j$, $j \in \{v, d\}$ (a_j being a per-unit-of-volume advertisement revenue), and cumulated revenue $a_d \lambda_d N_d + a_v \lambda_v N_v$ to evaluate if differentiation necessarily means a loss for CPs and the type of preferred differentiation; iii) user satisfaction/demand N_j and total demand $N_v + N_d$.

In the next subsections, we describe the user equilibria for our three scenarios that were computed in [12] but called differently and for a different analysis.

2.3 User equilibrium from the literature

No differentiation/neutrality The Wardrop equilibrium in the case with no differentiation is such that:

1. If $p > 1$, only type-v users join, with a mass $N_v(p) = \frac{\mu - \sqrt[\alpha_v]{p}}{\lambda_v}$.
2. If $p < 1$, only type-d users join, $N_d(p) = \frac{\mu - \sqrt[\alpha_d]{p}}{\lambda_d}$ and the ISP revenue is $\Pi(p) = \lambda_d N_d(p)$.
3. If $p = 1$, there is an infinite number of equilibria, type-d and v users having the same sensitivity to price, but we will consider in that case that only type-v users will be present.

The optimal revenue in the neutral case is $\Pi^{(n)} = \max\left(p_d \frac{\mu}{\alpha_d + 1}, p_v \frac{\mu}{\alpha_v + 1}\right)$ with $p_d = \left(\frac{\mu \alpha_d}{\alpha_d + 1}\right)^{\alpha_d}$ and $p_v = \left(\frac{\mu \alpha_v}{\alpha_v + 1}\right)^{\alpha_v}$.

The type of users present at the revenue-maximizing price therefore depends on the value of the service rate μ. More specifically, there is a threshold

$$\mu^* = \left(\left(\frac{\alpha_d}{\alpha_d + 1}\right)^{\alpha_d} \left(\frac{\alpha_v + 1}{\alpha_v}\right)^{\alpha_v} \frac{\alpha_v + 1}{\alpha_d + 1}\right)^{\frac{1}{\alpha_v - \alpha_d}} \quad (1)$$

such that if $\mu < \mu^*$, we are in the case with type-d users only, whereas if $\mu \geq \mu^*$ we are in the scenario with only type-v users.

Dedicated classes/ISP deciding differentiation The case when the ISP decides which type of users gets which priority class is called the *dedicated class scenario* in [12].

Define p_H and p_L as the respective prices for the high and low priority service classes. The output of the case with dedicated classes (with type-v users being assigned the high-priority class) is:

• The mass of class-H is $N_H = N_v(p_H, p_L) = \frac{\mu - \sqrt[\alpha_v]{p_H}}{\lambda_v}$ if $p_H < \mu^{\alpha_v}$ and 0 otherwise.

• The mass of class-L is

$$N_L = N_d(p_H, p_L) = \begin{cases} \lambda_d^{-1}\left(\sqrt[\alpha_v]{p_H} - \mu \frac{\sqrt[\alpha_d]{p_L}}{\sqrt[\alpha_v]{p_H}}\right) & \text{if } p_L \leq \frac{p_H^{2\frac{\alpha_d}{\alpha_v}}}{\mu^{\alpha_d}} \text{ and } p_H \leq \mu^{\alpha_v} \\ \lambda_d^{-1}(\mu - \sqrt[\alpha_d]{p_L}) & \text{if } p_L < \mu^{\alpha_d} \text{ and } p_H > \mu^{\alpha_v} \\ 0 & \text{otherwise.} \end{cases}$$

Price optimization is tricky, with non-informative formulas provided in [12], hence we will perform a numerical optimization.

Open classes/user-defined differentiation Users deciding which priority class to use is called the *open class scenario* in [12]. The user equilibrium (that is, masses of customers) with open classes is:

• If $p_L, p_H > 1$, we only have type-v users, with

$$N_H = \frac{\mu - \sqrt[\alpha_v]{p_H}}{\lambda_v} \text{ and } N_L = \frac{\sqrt[\alpha_v]{p_H} - \mu \sqrt[\alpha_v]{\frac{p_L}{p_H}}}{\lambda_v}$$

• If $p_L, p_H < 1$, we only have type-d users, with

$$N_H = \frac{\mu - \sqrt[\alpha_d]{p_H}}{\lambda_d} \text{ and } N_L = \frac{\sqrt[\alpha_d]{p_H} - \mu \sqrt[\alpha_d]{\frac{p_L}{p_H}}}{\lambda_d};$$

• If $p_L < 1$ and $p_H > 1$, the low-priority queue will be used by data customers, and the high-priority queue by video customers; $N_H = N_v = \left[\frac{\mu - \sqrt[\alpha_v]{p_H}}{\lambda_v}\right]^+$ and

$$N_L = N_d(p_H, p_L) = \begin{cases} \lambda_d^{-1}\left(\sqrt[\alpha_v]{p_H} - \mu \frac{\sqrt[\alpha_d]{p_L}}{\sqrt[\alpha_v]{p_H}}\right) & \text{if } p_L \leq \frac{p_H^{2\frac{\alpha_d}{\alpha_v}}}{\mu^{\alpha_d}} \text{ and } p_H \leq \mu^{\alpha_v} \\ \lambda_d^{-1}(\mu - \sqrt[\alpha_d]{p_L}) & \text{if } p_L < \mu^{\alpha_d} \text{ and } p_H > \mu^{\alpha_v} \\ 0 & \text{otherwise.} \end{cases}$$

Similarly to the case with dedicated classes, choosing the prices optimizing the ISP revenue will be performed numerically.

3 Numerical results

We now compute and compare the revenues and demand in the three scenarios. To start, Figure 1 displays, for the neutral case (only one class), all values on the

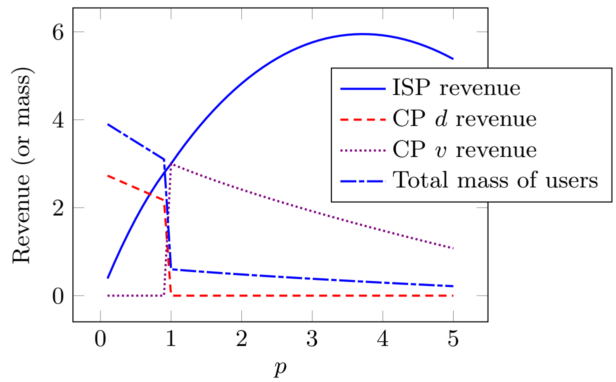

Fig. 1. Revenues and demand in the neutral case when $\mu = 4, a_d = 0.7, a_v = 1, \alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

same graph when varying the price p charged by packet by the ISP. (Note that units are different for demand and revenues.) When $p < 1$ only type-d users are served, while there are only type-v users when $p > 1$. We can check that there is an optimal price to be charged if the ISP wishes to maximize its revenue. As could be expected, that price does not correspond to an optimum for users or CPs.

The outputs at prices optimizing the ISP revenue for the three scenarios are displayed in Figure 2 (optimal price), Figure 3 (corresponding ISP revenue), Figure 4 (CPs revenue), and Figure 5 (demand) when varying the service rate μ of the M/M/1 queue.

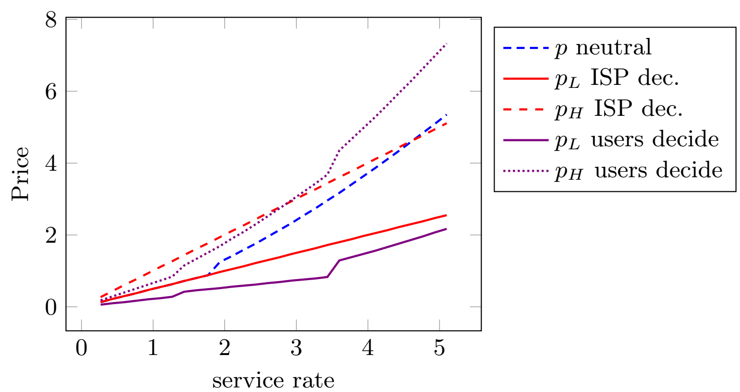

Fig. 2. Optimal prices in terms of μ when $\alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

It can be readily checked in Figure 2 that prices increase with the service rate μ; in other words, a scarce resource does not correspond to a price increase. When μ is small, the price p_L when the ISP decides the classes is equal to the optimal neutral price, but that class is not served. If users decide differentiation, p_L is smaller. The price p_H of the high priority class is larger if users decide than if the ISP does when μ is large, and, it may be counter-intuitive, for the largest values of μ, the neutral price is larger than the high-class price when ISP decides differentiation; it is due to type-d users in service in the former case and type-v in the latter.

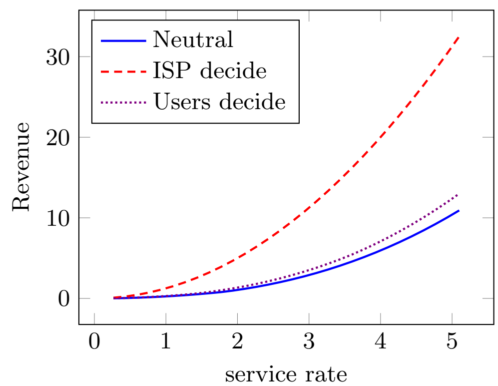

Fig. 3. Optimal ISP revenues in terms of μ when $a_d = 0.7, a_v = 1, \alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

On Figure 3, the ISP revenue is as expected larger if it decides differentiation, but letting users decide is a better solution for the ISP than a fully neutral scenario. The larger the network capacity, the larger the revenue differences.

Now let us look at CP revenues in Figure 4 and demands N_d, N_v in Figure 5: note that they do not always increase with μ. Remark also that letting users decide differentiation is the only case when both types of service will be active in the network for a range of service rates; in this sense it is an interesting and fair scenario.

Looking at total demand in Figure 6, letting users decide is the best option when μ is small, and is always better than the neutral case. As μ increases, differentiation decided by the ISP is better, but it hides that only one type of service is used, which is again not the case when users decide for a range of values.

4 Conclusions

As the neutrality debate is still raging, we have discussed in this paper the option of letting users choose a service class for each application. This way, innovation

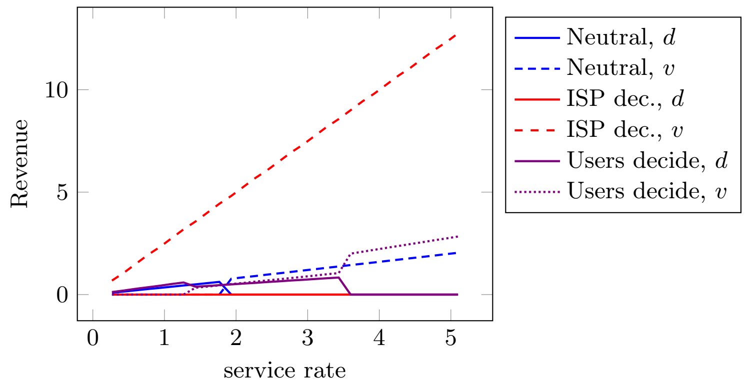

Fig. 4. Optimal CP revenues in terms of μ when $a_d = 0.7, a_v = 1, \alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

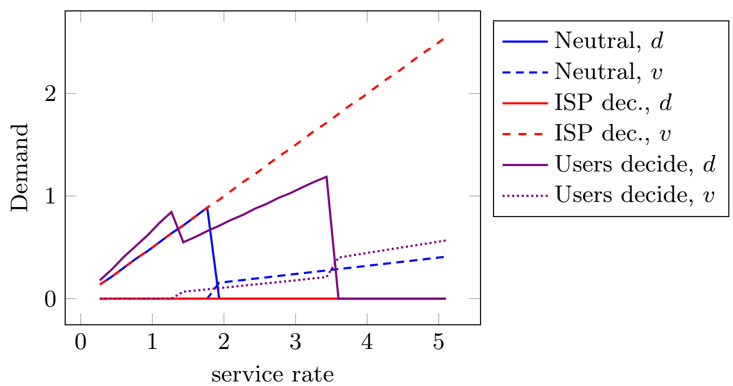

Fig. 5. Demand at optimal prices in terms of μ when $\alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

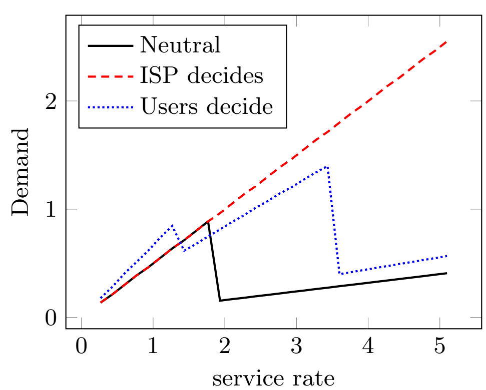

Fig. 6. Total demand at optimal prices in terms of μ when $\alpha_d = 1, \alpha_v = 1.5, \lambda_d = 1, \lambda_v = 5$

would not be slowed down with users asking for good quality for new and good quality applications; it would also reduce the load on ISPs, with no requirement to provide a good quality for all services. The network would stay neutral because not interfering (except for the price decision) on service class choices.

We have used a model from the literature to discuss and compare the output with a neutral situation and a fully non-neutral one where the ISP decides how to differentiate traffic. According to our results, letting users decide seems a nice trade-off in terms of demand and revenues for CPs and ISPs.

As future works, we would like to analyze other types of models and check whether they corroborate the results we have obtained here. Another issue is how to implement this promising option in practice, from a technical point of view.

References

1. Altman, E., Legout, A., Xu, Y.: Network non-neutrality debate: An economic analysis. In: Proc. of Networking. pp. 68–81. LNCS 6641, Springer Verlag, Barcelona, Spain (May 2011)
2. BEREC: Berec guidelines on the implementation by national regulators of european net neutrality rules. Document number: BoR (16) 127 (Aug 2016), available at https://berec.europa.eu/eng/document_register/subject_matter/berec/regulatory_best_practices/guidelines/6160-berec-guidelines-on-the-implementation-by-national-regulators-of-european-net-neutrality-rules
3. Coucheney, P., Maillé, P., Tuffin, B.: Network neutrality debate and ISP inter-relations: Traffic exchange, revenue sharing, and disconnection threat. Netnomics **1**(3), 155–182 (2014)
4. Federal Communications Commission: Protecting and promoting the open internet (April 2015), available at https://www.federalregister.gov/documents/2015/04/13/2015-07841/protecting-and-promoting-the-open-internet
5. Federal Communications Commission: Restoring internet freedom (Jan 2018), available at https://docs.fcc.gov/public/attachments/FCC-17-166A1.pdf
6. Kleinrock, L.: Queueing Systems, Vol. 2. Wiley, New York (1976)
7. Lenard, T., May, R.E.: Net Neutrality or Net Neutering: Should Broadband Internet Services be Regulated. Springer (2006)
8. Maillé, P., Reichl, P., Tuffin, B.: Internet governance and economics of network neutrality. In: Hadjiantonis, A., Stiller, B. (eds.) Telecommunications Economics - Selected Results of the COST Action IS0605 EconTel, pp. 108–116. Lecture Notes in Computer Science 7216, Springer Verlag (2012)
9. Maillé, P., Simon, G., Tuffin, B.: Toward a net neutrality debate that conforms to the 2010s. IEEE Communications Magazine **54**(3), 94–99 (2016)
10. Maillé, P., Tuffin, B.: Telecommunication Network Economics: From Theory to Applications. Cambridge University Press (2014)
11. Maillé, P., Tuffin, B.: Neutral and Non-Neutral Countries in a Global Internet: What Does it Imply? In: Proc. of GECON - 16th International Conference on the Economics of Grids, Clouds, Systems, and Services. Springer, Leeds, United Kingdom (Sep 2019), https://hal.inria.fr/hal-02128437

12. Mandjes, M.: Pricing Strategies under Heterogeneous Service Requirements. In: IEEE INFOCOM (2003)
13. Njoroge, P., Ozdaglar, A., Stier-Moses, N., Weintraub, G.: Investment in two sided markets and the net neutrality debate. Tech. Rep. DRO-2010-05, Columbia University, Decision, Risk and Operations Working Papers Series (2010)
14. Odlyzko, A.: Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets. Review of Network Economics **8**(1), 40–60 (2009)
15. Osborne, M., Rubinstein, A.: A Course in Game theory. MIT Press (1994)
16. Schulzrinne, H.: Network neutrality is about money, not packets. IEEE Internet Computing **22**(6), 8–17 (2018). https://doi.org/10.1109/MIC.2018.2877837, `https://doi.org/10.1109/MIC.2018.2877837`
17. Wardrop, J.: Some theoretical aspects of road traffic research. Proceedings of the Institute of Civil Engineers **1**, 325–378 (1957)
18. Wu, T.: Network neutrality, broadband discrimination. Journal of Telecommunications and High Technology (2003)