

Impact of the Backbone Network Market Structure on the ISP Competition

Pierre Coucheney*, Patrick Maillé[†], Bruno Tuffin*, and Luis Guijarro[‡]

*Inria Rennes Bretagne Atlantique

Campus Universitaire de Beaulieu, 35042 Rennes Cedex, France

Email: {pierre.coucheney},{bruno.tuffin}@inria.fr

[†]Institut Telecom/Telecom Bretagne

2, rue de la Châtaigneraie CS17607, 35576 Cesson-Sévigné Cedex, France

Email: patrick.maille@telecom-bretagne.eu

[‡]Universitat Politècnica de Valencia

Camino de Vera s/n, 46022 Valencia, Spain

Email: lguijar@dcom.upv.es

Abstract—The network market shows a hierarchy of providers, with backbone providers and access providers. We aim in this paper at describing the economic relationships between those network providers, where content providers are associated to backbone providers while access providers compete to attract end users. We study those interactions thanks to game theory, using a multi-level game model. We analyze in addition the vertically integrated case, i.e., the situation when the backbone provider is also an access provider.

We then discuss the impact of this analysis on the network neutrality debate, highlighting the importance of a regulation on prices charged by backbone to access providers.

I. INTRODUCTION

The Internet network is made of an inter-connection of Internet service providers (ISPs) and displays a particular type of topology. There are indeed two kinds of ISPs: transit ISPs -also called Internet Backbone Providers (IBPs)- and local ISPs, or simply ISPs in this paper. The latter provide the last-mile link to end-users (the residential consumers). In contrast, IBPs operate portions of the Internet backbone, whereby transit services are offered to local ISPs. The local ISPs are charged for the transit service and this service entitles local ISPs to reach not only the IBP network, but also the IBP peers' networks and any other network served by the latter [1]. On the other hand, IBPs have no residential customers and serve only content providers, though there are notable examples of companies operating both transit and local networks.

The core economic distinction between the two ISP classes is that local ISPs are thought to have more market power. This market power is a result of substantial economy of scale in providing mass-market (residential) broadband access, which limits (today at least) the number of local ISPs in a specific residential area. In particular, today's the end-user choice for a local ISP is limited to one to three at most. While (local) ISPs compete for end users thanks to (in general quite long-term) contracts, IBPs compete for transit traffic and for the access of major content providers. This competition is known to be fierce because individual content providers (*i*) generate a lot of traffic and therefore can economically support dedicated links

from several suppliers, and (*ii*) can pick to locate themselves at the points where the network access is the cheapest—that is, to attach directly to the backbone [2]. Additionally, both the easy supply of raw transmission capacity and the open standard nature of the Internet keep the barrier to entry at a negligible level [3]. We will see that, as a consequence, the competition between IBPs can be modelled as a Bertrand competition.

Our aim in this paper is to describe and analyze, thanks to game theory, the economic relationships between the actors in the Internet, with a particular focus on the impact of the Internet network structure on the ISP competition. This analysis will be helpful for regulatory bodies in particular to better understand which are the parties suffering from competition, and possibly leading to rules for a “better” Internet. The paper can be viewed as an extension of [4], where the economic transit agreements between ISPs have been reviewed in order to determine their best strategy. The network separation between IBPs and (local) ISPs was not considered though, while it could be a key characteristic. As in [4], end users ISP selection is performed thanks to a discrete choice model.

Note that there are other papers analyzing the economic interactions between ISPs. As an example, [5] uses a game-theoretic model of two ISPs, drawn from a larger set consisting of transit and local ISPs, who choose between peering and transit agreements. The study focuses on the costs of inter-connection taking into account traffic imbalances. But it does not consider like here an explicit model of the user behavior. We will also model an alternative scenario where the backbone service is provided by a unique IBP. This monopolistic scenario has also received support in the literature. Actually, [6] argues that the excess capacity which currently characterizes the Internet makes IBP's long-term profit expectation very modest. Taking this into account and the fact that IBPs have to bear substantial fixed costs, it concludes that some players will unavoidably exit the backbone service market. If some players leave, this will shift market power to the remaining networks.

Thus, a scenario where a monopolistic IBPs is operating is also a feasible one.

The paper is organized as follows. Section II describes the model: it explains the separation of providers into IBPs and (local) ISPs, their economic relations and profit functions, as well as the users' preference model determining the ISP they will subscribe to. The decisions being played at different time scales, the multiple stages will be described. We will also explain why the IBP competition for content providers and ISPs' association leads to a Bertrand competition, so that IBPs can be treated as a single backbone network. Section III describes and analyzes the game at the shortest time scale: the competition between ISPs for end users, making use of the users' preference model. It separates two cases: when there is no transit price from the IBP (resuming to an analysis carried out in [4]), and when there is a transit price t per unit of volume. Section IV then looks at the transit price determination which, even if made before the ISP competition, is performed *anticipating* the subsequent equilibrium. Three scenarios are considered: several IBPs in competition (hence as a consequence of Bertrand competition, a transit price $t = 0$), a monopolistic IBP seeking to maximize his profit, and a transit price chosen by a regulator to minimize *social welfare*. Section V discusses a specific scenario where the IBP happens to also be an ISP (a reality in some countries) and investigates the consequences of that vertical integration on the competition among ISPs. Finally Section VI concludes by providing the trends that can be extracted from the analysis. Mathematical proofs are provided in appendix.

II. MODEL

A. System Description

As described in the previous section, our aim is to understand the economic relations between ISPs (and their competition for end users, while the competition for content providers is addressed at a lesser degree), when there is a hierarchy among those ISPs. Indeed, some are local ISPs (just noted ISPs to simplify the notations), to which end-users are directly connected, while other so-called Internet Backbone Providers (IBPs) -also called transit ISPs-, represent providers operating on the backbone, but the interesting situation when an IBP is also a local ISP will be treated in Section V.

We therefore describe and model here the interactions of all those providers in the context of end users choosing their own ISP depending on price. The relations between end users, ISPs, and IBPs are summarized in Figure 1.

A few remarks must be made at this point.

- Note first that CPs are usually connected to an IBP (instead of an ISP) because of the amount of traffic they generate and the possibility to dedicate a link to them [2].
- Moreover, it is well known that competition at the IBP level is rather fierce [7]. In that market, CPs and ISPs simply choose to connect to the *cheapest* IBP (as opposed to end-users who generally have additional considerations,




Figure 1. Representation of relations between users, ISPs and the backbone network. Arrows are labeled with the unit prices (if any), with t per volume and p_A, p_B per subscriber.

such as reputation or loyalty). When several IBPs are present, because of this choice of the cheapest IBP by ISPs and CPs, IBPs are engaged in a classical Bertrand competition [8] leading at equilibrium to set a zero price charged to ISPs and CPs. This is because, roughly, if the smallest price of the opponent is positive, setting your own price slightly less will attract all CPs and ISPs (hence all profit), but this can then be done also by an opponent, so the prices decrease to the marginal costs, that we assume to be zero here. In game-theoretic terms, such a situation is called a *prisoner's dilemma*: all IBPs would be better off in a situation with positive prices, but selfishness leads to an outcome that is bad for all of them. As a result, when several IBPs are present, the set of IBPs will be treated as a (single) backbone network with zero prices.

- Also, since we therefore consider a single IBP with CPs connected to it, the relations between ISPs, the IBP, and end users, are not affected by the price imposed to CPs; we therefore can just consider them as part of the IBP.

Let us now define the economic relations between the actors displayed on Figure 1. To simplify the analysis, we limit ourselves to two ISPs, denoted by A and B . Those ISPs charge *subscription fees*, respectively p_A and p_B , to the end users attached to them. The IBP (possibly representing the set of competing IBPs) charges a *per-unit-of-volume* price t to the ISPs for the traffic they request, since traffic downloaded from CPs to end users constitutes most of the Internet traffic. Note, again, that we take $t = 0$ when several IBPs coexist.

Users are represented as a continuum of total mass assumed to be 1 without loss of generality. We denote by σ_A and σ_B the proportion of users with ISP A and ISP B , respectively. Of course, $\sigma_A, \sigma_B \geq 0$ and $\sigma_A + \sigma_B \leq 1$. We also assume that there is a proportionality factor d_0 relating the amount of traffic downloaded by users to the mass of users (in other words, a mass σ of users has a total demand of traffic $d_0\sigma$).

B. Multi-stage Decision Problem

We can remark that each stakeholder takes its decision (price setting for ISPs and IBP, ISP choice for users) at a specific and a different time scale. The decision taken at a given time of course influences the later choices, but we

also assume that they are played strategically, that means anticipating the actions at smaller time scales. Such interaction problems correspond to *leader-follower* situations [9], where the leader(s) -here, the actors playing at a large time scale- anticipate the reaction of the follower(s) to make the best strategic choice. For our specific situation, the decision order of the multi-stage decision problem is the following:

- 1) At the highest time scale, the transit price t is determined. We will consider three scenarios:
 - The case of several IBPs, then engaged in a Bertrand competition for (local) ISPs, leading as explained above to zero transit prices, and therefore to considering the set of IBPs as a single backbone at no cost ($t = 0$ in Figure 1).
 - The case of a monopolistic IBP that can set the transit price t to maximize its profit.
 - The situation where the transit price t is determined by a regulator seeking to maximize social welfare (explicitly defined later).
- 2) Then the ISPs compete on the subscription prices.
- 3) Finally, each user chooses either one of the two ISPs to get Internet access, or prefers not to get any access.

Multi-level decision problems can be solved by *backward induction*, consisting in determining the followers' reactions, and using that knowledge to compute the leaders' actions. We will therefore successively analyze the behavior of the users, the ISPs, and the IBP.

C. Users preferences

The model we consider to describe the users choice of an ISP is directly taken from [4]. We summarize it in this subsection: choices are taken based on the ISP subscription prices, but also on other considerations (e.g., reputation) that are sometimes difficult to rationalize. Basically, applying traditional discrete choice models (see [4] for details), a user is assumed to have a valuation (utility) for ISP $i \in \{A, B\}$ of the form $V_i = \alpha \log(1/p_i) + \kappa_i$ where κ_i is a random variable following a Gumbel distribution and $\alpha \log(1/p_i)$ expresses the dissatisfaction for large prices. The parameter $\alpha > 0$ models the user sensitivity to price, and the logarithm form comes from studies on human perception of stimuli (here, price) [10].

Users choose the provider giving the largest valuation for them, but there is also the possibility that they prefer not to connect to any ISP, this being represented by the (random) valuation $V_0 = \alpha \log(1/p_0) + \kappa_0$ with κ_0 still a Gumbel-distributed random variable, and $p_0 > 0$ being here the perceived cost of not having access to the Internet. Under those rules, the mass (or equivalently proportion) of users on each ISP is (for $i \in \{A, B\}$):

$$\sigma_i = \frac{p_i^{-\alpha}}{p_A^{-\alpha} + p_B^{-\alpha} + p_0^{-\alpha}}. \quad (1)$$

The limiting cases of one ISP setting a zero price lead to $\sigma_i = 1$ and $\sigma_j = 0$ ($i \neq j$) if $p_i = 0$ and $p_j > 0$. When both prices are null ($p_A = p_B = 0$), we moreover assume that $\sigma_A = \sigma_B = 1/2$.

D. Providers' utilities, user welfare, and social welfare

As a consequence of our model definition, the profits of ISPs are

$$U_i = \sigma_i(p_i - d_0 t) \quad \text{for } i \in \{A, B\} \quad (2)$$

and that of the IBP (indexed by I) is

$$U_I = d_0 t(\sigma_A + \sigma_B), \quad (3)$$

taking into account the subscription from end users to the ISPs, but also the volume-based payment from the ISPs to the IBP.

Another value of interest is the *user welfare* UW, representing the total utility surplus that users get from the internet access (compared to the situation with no access at all, *i.e.* just the outside option). It is directly derived from the user choice model of Section II-C, and has been proved from simple computations in [4] to be

$$\text{UW} = \log \left(1 + \left(\frac{p_0}{p_A} \right)^\alpha + \left(\frac{p_0}{p_B} \right)^\alpha \right). \quad (4)$$

Notice that the user welfare is equal to $-\log(1 - \sigma_A - \sigma_B)$. Hence the maximum is achieved when the mass of users having an internet access, *i.e.* $\sigma_A + \sigma_B$, is maximal, therefore when one ISP offers the access for free ($p_A = 0$ or $p_B = 0$).

A last notion of interest is the *social welfare* SW. This is the aggregated value (or utility) of *all* stakeholders: providers plus end users. The utility of each provider is its profits, while for the (aggregated) set of users it consists of the user welfare. In order to deal with those non homogeneous values, we introduce a factor λ expressing a unit of user utility in a monetary unit. This factor also allows to weigh user welfare against provider profits. As a result, social welfare SW is

$$\text{SW} = U_A + U_B + U_I + \lambda \text{UW} = \sigma_A p_A + \sigma_B p_B + \lambda \text{UW}. \quad (5)$$

III. ISPS PRICE COMPETITION

We now turn our attention to the second decision level highlighted in Section II-B, that corresponds to ISPs playing a game on subscription prices, given the transit price t , and anticipating the reaction of users to their price profile (p_A, p_B) . Section III-A presents the case when $t = 0$, corresponding particularly to the case when there are several IBPs in (Bertrand) competition. Since we do not consider any cost for the IBPs, it corresponds to a net benefit equal to zero. Remark that if costs were taken into account, then the Bertrand competition leads to two possibilities: (i) if there is a unique IBP with the smallest transit cost, its transit price will be just below the second smallest cost among IBPs (and we end up with a monopoly), and (ii) if there are several IBPs with the smallest transit cost, the transit price will be that transit cost (hence again a zero net benefit). The case $t = 0$ actually resumes exactly to the results obtained in [4]. The mathematical contribution of our paper starts in Section III-B (if we except the model definition with hierarchical providers), corresponding to the case of a positive transit price.

A. No transit price ($t = 0$)

The case when there is no transit price has already been analyzed and proved in [4]:

Proposition 1. *The Nash equilibria on the price competition between ISPs depend on the sensitivity to prices parameter α :*

- If $\alpha \leq 1$, the Nash equilibria are
 - $p_A = p_B = 0$, which is not likely to be played since setting a null access price is a weakly dominated strategy for each ISP.
 - $p_A = p_B = \infty$, resulting in infinite profit for ISPs.
- If $1 < \alpha < 2$, the Nash equilibria are
 - $p_A = p_B = 0$, which is not likely to be played since setting a null access price is a weakly dominated strategy for each ISP.
 - $p_A = p_B = \left(\frac{2-\alpha}{\alpha-1}\right)^{1/\alpha} p_0$, which results in the profits $U_A = U_B = \frac{\alpha-1}{\alpha} \left(\frac{2-\alpha}{\alpha-1}\right)^{1/\alpha} p_0$.
- If $2 \leq \alpha$, the unique Nash equilibrium is $p_A = p_B = 0$, hence ISPs get no profit.

Prices observed in practice are finite, suggesting that users are quite sensitive to prices, i.e. $\alpha > 1$. However, the fact that ISPs continue to reduce their prices in several countries does not rule out the possibility of high price sensitivities $\alpha \geq 2$.

B. With transit price ($t > 0$)

With a positive transit price, we are no longer able to derive, in general, the analytical expression for the Nash equilibria. However, we manage to provide the existence and some useful characterizations.

Before that, let us give some properties of the best response function, detailed in the following lemma, that will be exploited. We denote by $BR_A^t(p_B)$ and $BR_B^t(p_A)$ the best response of ISP A and B respectively, when the transit price is t : it corresponds to the set of prices maximizing the profit of the ISP given the action of the other ISP. In this scenario, ISPs are fully symmetric (in terms of profit and set of actions), hence we have $BR_A^t(p) = BR_B^t(p)$ for every $p > 0$, and we will omit the ISP subscript.

Lemma 1. *Assume $\alpha > 1$ and $t > 0$. The best response of each ISP is (i)*

- 1) single valued,
- 2) continuous,
- 3) uniformly upper- and lower-bounded by strictly positive bounds, with in particular the lower-bound $\frac{\alpha}{\alpha-1}d_0t < BR^t(p)$,
- 4) strictly increasing (in p),
- 5) strictly increasing with the transit price t : $BR^t(p) > BR^{t'}(p)$ if and only if $t > t'$.

The proof is given in Appendix A. We now enunciate the main result of the section, regarding the Nash equilibria of the

game between ISPs. Similarly to Proposition 1, the equilibria depend on the price sensitivity of users.

Proposition 2. *Assume $t > 0$, then*

- if $0 < \alpha \leq 1$, then $p_A = p_B = \infty$ is the unique Nash equilibrium, leading to infinite profits for ISPs;
- if $1 < \alpha$, there exists a Nash equilibrium, and every Nash equilibrium is such that $p_A = p_B \stackrel{\text{def}}{=} p^{NE} > 0$, and results in a strictly positive profit for each ISP, equal to $\frac{\alpha-1}{\alpha}p^{NE} - d_0t$. Moreover, if the Nash equilibrium is unique, p^{NE} strictly increases with the transit price.

The proof is given in Appendix A.

Figure 2 shows on one example the symmetric best response of ISPs. One can notice the properties given in Lemma 1: the functions are increasing, and greater than $\frac{\alpha}{\alpha-1}d_0t = 3$, with $\alpha = 1.5$, $d_0 = 1$, and $t = 1$. It also appears that the Nash equilibrium is unique, which we have observed in all the scenarios we analyzed numerically.

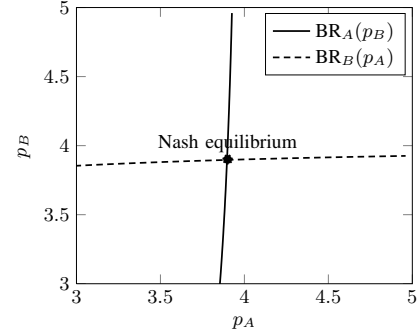


Figure 2. ISP's (symmetric) best response functions, with $t = 1.0$, $\alpha = 1.5$, $d_0 = 1$, and $p_0 = 1.0$. There is a unique Nash equilibrium.

We can make several remarks when we compare the case of positive transit price to the situation where it is null. We first get some positive effects for ISPs, even though they are charged that extra price: with positive prices, the (weakly-dominated) Nash equilibrium $(0,0)$ does not stand anymore, since attracting users by lowering one's price yields subscription revenues but also incurs costs due to transfer charges. In the same vein, with a high sensitivity to prices ($\alpha > 2$), we are able to derive a non-null Nash equilibrium. Figure 3 compares the Nash equilibrium prices in terms of α when $t = 0$ and for various values of t , illustrating the results of Propositions 1 and 2.

We plot in Figure 4 the utility of each stakeholder for $\alpha = 1.5$, when the transit price varies. Both ISP and IBP profits are first increasing, then decreasing, the maximum of each one being reached for a strictly positive transit price. It is interesting to note, and somewhat counter-intuitive, that the sum of ISPs profits is first increasing (even if for a short period) with the transit prices they have to pay to the IBP. They can first compensate the payment by an increase in their subscription fee, without decreasing too much demand. Indeed, the transit price t prevents ISPs from decreasing their




Figure 3. Equilibrium prices of ISPs, as a function of α , for different values of the transit price t .

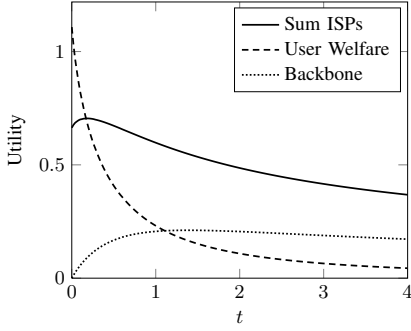


Figure 4. Utilities of the stakeholders at Nash equilibrium prices, as a function of t , with $\alpha = 1.5$, $p_0 = 1$, $d_0 = 1$.

prices too much, which reduces the extent of the “prisoner’s dilemma” situation among ISPs: both ISPs stay with higher price values, and improve their profit. Similarly, the IBP profit first increases with t but then tends to decrease, because this forces ISPs to increase their price, hence the total volume of traffic decreases.

On the other hand, notice that ISPs do not benefit from the transit price (in the situation plotted on the figure) if that price is set by the IBP. Another negative effect of the transit price is that user welfare always decreases with t , because of the increasingness of the ISP equilibrium price proved in Proposition 2.

Assuming that the Nash equilibrium is unique, we now prove that transit pricing has to be prevented to maximize social welfare.

Proposition 3. *Assume that the Nash equilibrium is unique, then the user welfare (4) at this point is decreasing with the transit price t . Hence, for every price sensitivity α , it is maximized for $t = 0$.*

Proof: If $\alpha \leq 1$, the Nash equilibrium does not depend on t . Hence the proposition is valid in this case.

If $\alpha > 1$, from Proposition 2, the unique and symmetric price at Nash equilibrium strictly increases with the transit price t , while the user welfare (4) is a strictly decreasing function of the ISPs price. Hence the result. ■

IV. IBP PRICE DETERMINATION

We now come to the first level of decision, where the IBP price is chosen first but anticipating what would be the Nash equilibrium at the ISP game level. As described in Section II-B, we consider three scenarios: (i) at least two IBPs engaged in a Bertrand competition and leading the transit price to zero, (ii) a monopolistic IBP setting the transit price maximizing his profit, and (iii) a regulator determining the transit price to maximize social welfare. From here, we use a numerical solver, and draw the profits and utility of each stakeholder at the transit price determined for each scenario. All curves shown are for the parameter values $p_0 = 1$ and $d_0 = 1$.

Figures 5, 6 and 7 respectively show the utility of ISPs, users, and IBP, as a function of the user sensitivity α , comparing the above three scenarios for a fixed t (and with the relative weight of UW versus provider revenues taken to $\lambda = 5\%$). Maximizing social welfare seems to be a good trade-




Figure 5. (Sum of) profit of ISPs as a function of the user price sensitivity α .

off here, when $\alpha < 2$, because it offers a balance between user welfare and providers’ revenues (of course, that balance depends on the value of λ). We observe that the monopolistic scenario is better than the competitive one for ISPs as soon as the user sensitivity is larger than a given threshold. From Propositions 1 and 2, that threshold is always below 2, since for $\alpha > 2$, the ISPs profit is null when $t = 0$. We observe that




Figure 6. User welfare as a function of the user price sensitivity α .

user price-sensitivity exacerbates the competition at the ISP level, that becomes extremely fierce when the transit price

is null, as pointed out in Proposition 1. As a consequence, for $\alpha > 2$ subscription prices are null and user welfare (and as a consequence social welfare) is therefore infinite in the multiple-IBP scenario. On the other hand, user welfare seems to grow linearly with $\alpha < 2$ for the two other scenarios. We can also remark on Figure 6 that competition is better for users than maximizing social welfare, since the IBP is not taken into account (see also Figure 7). Similarly, from Figure 7, the




Figure 7. IBP profit as a function of the user price sensitivity α .

IBP profit is larger for large users' sensitivity to prices (in the monopolistic case and up to $\alpha = 2$ for the social welfare maximizing case, since being zero for this last situation for $\alpha \geq 2$), and is of course always zero when $t = 0$.

We also display on the same figure (Figure 8) the transit price t maximizing social welfare, the corresponding sum of ISPs revenues, the IBP revenue and User Welfare, when α is fixed but the factor λ in the social welfare varies (weighing user welfare with respect to provider profits). We can see that the optimal price t decreases: the user welfare is getting more and more importance, and this resumes to decreasing t , up to a moment when $t = 0$. Non-surprisingly, the user welfare increases while the revenue of the IBP decreases with λ . Interestingly, the revenue of ISPs admits a maximum for a strictly positive λ . This is due to the fact that the optimal price of the IBP (paid by ISPs) quickly goes to zero when λ increases, hence an initial positive effect.




Figure 8. Transit price maximizing social welfare (5), with $\alpha = 1.5$ as well as the utilities of ISPs and users. The user welfare is expressed in monetary unit thanks to the conversion factor λ .

V. IF THE IBP IS ALSO AN ISP

In practice, it is often the case that a provider is at the same time an IBP and an ISP (for example Deutsche Telekom in Germany, or AT&T in the USA). We want to investigate that scenario, by assuming that ISP A is itself the monopolistic IBP, and compare it with the case where the IBP is a different provider (and chooses t to maximize his profit) as analyzed in the previous section (that scenario being hereafter referred to as the “differentiated case”). Under that new assumption, the utility function of ISP B does not change, whereas that of A becomes:

$$U_A = \sigma_{APA} + \sigma_B t.$$

We assume that the decision order remains the same: ISPs set their user price after the transit price t has been decided.

The results are displayed for the provider profits and user welfare respectively in Figures 9 and 10, in terms of α . It can be seen that, whatever the value of α , ISP B loses approximately the same amount because of the monopoly of A . On the other hand, A wins more when α is small (small sensitivity to price for end users). Surprisingly in Figure 10,




Figure 9. Profit of ISPs as a function of the user price sensitivity α .

end users benefit from ISP A being also the IBP, this benefit increasing with the sensitivity α . Actually, classical antitrust analysis identifies this fact as one of the main benefits from a vertical merger. Indeed, vertical mergers can avoid the double marginalization that the presence of a upstream monopoly and a downstream oligopoly may cause – see [11] for an explanation.




Figure 10. User welfare (increasing curves) and IBP optimal price (decreasing curves) in terms of the user price sensitivity α .

To better understand this evolution in terms of α , we also display respectively in Figures 10 and 11 the IBP optimal price, ISP equilibrium price and mass of end users accessing the network when the IBP is an ISP or not. Interestingly, the IBP reduces its transit price if it is also an ISP (Figure 10), and the same happens for both ISPs' price at equilibrium (but to a lesser extent for ISP *A*). As a consequence, more users access the network (Figure 11). This might actually explain the transit price decrease by the IBP, that is now more sensitive (through its ISP component) to the proportion of ISP subscribers.




Figure 11. Mass of subscribers (increasing curves) for each ISP and optimal IBP price (increasing curves) as a function of the user price sensitivity α .

VI. CONCLUSIONS

In this paper, we have built and analyzed an economic model encompassing the main actors and aspects of the current Internet market structure, including its hierarchical nature, with backbone and local providers. Regarding the main question addressed, concerning the economic relationships between those types of providers in a competitive context, the following conclusions can be drawn.

First, with users very sensitive to prices, then if ISPs are not charged by backbone providers, the competition among ISPs to attract subscribers is so fierce that it leads to null profits for ISPs, and questions their survivability. Imposing transit prices to ISPs to reach the backbone prevents that non-desirable outcome, however the value of that transit price must be properly determined: when appropriately set, it can ensure the survivability of the network providers and a revenue shared among local and backbone providers, while keeping subscription prices to a reasonable level. In that context, the structure of the backbone market is of prominent importance, since even allowing backbone providers to fix the transit price may lead to null transit prices because of competition. For those reasons, it seems justified that the transit prices be set by a regulator, and its value chosen to reach a globally satisfying outcome (taking into account the user behavior, and the competition among ISPs).

The situation when a backbone provider is also an ISP has been considered, since it occurs in several countries. From a global and user point of view, such a situation is economically desirable since it drives user prices down, while yielding the integrated entity a significantly larger revenue due to the

vertical merger, which can be used to maintain and improve the network. We recognize here an argument of net neutrality opponents, of the backbone provider having to charge the traffic to ISPs (but not their own, here) to guarantee the network sustainability. Nevertheless, the regulator intervention may still be needed here, to ensure that the dominant position of the merged actor does not prevent other ISPs from staying in or entering the local access market.

That last aspect (regulation of a dominant position) is beyond the scope of this paper, but can be considered in a future work. Another interesting direction would be to explicitly include the content provider side of the market in our model. In particular, with regard to the net neutrality debate, we would aim at investigating the justifications of a differentiated pricing that would be applied by ISPs depending on the content.

ACKNOWLEDGMENTS

This work has been supported by EuroNF NoE's project INNIS on network neutrality for all authors, and by the French ANR VERSO project CAPTURES on competition between providers and Inria's ARC project MENEUR for the first three authors.

APPENDIX

We show the result for ISP *A*, since it is symmetric for *B*. Let us first differentiate the profit of ISP *A* with respect to its set price p_A :

$$\begin{aligned} \frac{\partial U_A}{\partial p_A} &= \frac{(p_A^{-\alpha} + p_B^{-\alpha} + p_0^{-\alpha})p_A^{-\alpha-1}(-\alpha(p_A - d_0t) + p_A)}{(p_A^{-\alpha} + p_B^{-\alpha} + p_0^{-\alpha})^2} \\ &\quad + \frac{\alpha p_A^{-2\alpha-1}(p_A - d_0t)}{(p_A^{-\alpha} + p_B^{-\alpha} + p_0^{-\alpha})^2} \\ &= \left(\frac{p_A^{-\alpha}}{p_A^{-\alpha} + p_B^{-\alpha} + p_0^{-\alpha}} \right)^2 (-K(p_B)p_A^\alpha + L(p_B, t)p_A^{\alpha-1} + 1) \end{aligned} \quad (6)$$

with $K(p_B) \stackrel{\text{def}}{=} (\alpha - 1)(p_B^{-\alpha} + p_0^{-\alpha})$ and $L(p_B, t) = \alpha d_0 t (p_B^{-\alpha} + p_0^{-\alpha})$. Hence (7) has the same sign as the function

$$S(p_A, p_B, t) \stackrel{\text{def}}{=} -K(p_B)p_A^\alpha + L(p_B, t)p_A^{\alpha-1} + 1.$$

Notice that, since $\alpha > 1$, K and L are positive functions (strictly for K).

(1) We have $\frac{\partial S}{\partial p_A}(p_A, p_B, t) = p_A^{\alpha-2}(-\alpha K(p_B)p_A + (\alpha - 1)L(p_B, t))$, hence $S(p_A, p_B, t)$ is strictly increasing with p_A until $p_A = \frac{\alpha-1}{\alpha} \frac{L(p_B, t)}{K(p_B)} > 0$, and strictly decreasing after. Moreover S is 1 when $p_A = 0$ (since $\alpha > 1$), and it goes to $-\infty$ when $p_A \rightarrow \infty$. Hence, by continuity of S , there is a unique value p_A^* (that depends on p_B and t) such that $S(p_A^*, p_B, t) = 0$, and $S(p_A, p_B, t) > 0$ (resp. $S(p_A, p_B, t) < 0$) if and only if $p_A < p_A^*$ (resp. $p_A > p_A^*$). Consequently $\text{BR}^t(p_B) = p_A^*$ is single-valued. Another consequence is that we have a criterion to compare a price p_A with the best response $\text{BR}^t(p_B)$:

$$p_A < \text{BR}^t(p_B) \quad \text{if and only if} \quad S(p_A, p_B, t) > 0. \quad (8)$$

(2) The continuity is a consequence of Berge's maximum theorem [12]. The hypotheses of the theorem are valid here, so that the best-response price application is upper hemicontinuous. Since the best response is single-valued, it is a continuous function.

(3) The best response $\text{BR}^t(p)$ is lower bounded by $\frac{\alpha}{\alpha-1}d_0t > d_0t$. Indeed, for every p , $S(\frac{\alpha}{\alpha-1}d_0t, p, t) = 1 > 0$.

It is also upper bounded by $x \stackrel{\text{def}}{=} \max(\frac{1 + \alpha d_0t}{\alpha - 1}, p_0^{\frac{\alpha}{\alpha-1}})$. Indeed, $S(x, p, t) = x^{\alpha-1}(p^{-\alpha} + p_0^{-\alpha})(\alpha d_0t + (1-\alpha)x) + 1 \leq -x^{\alpha-1}(p^{-\alpha} + p_0^{-\alpha}) + 1 \leq -x^{\alpha-1}p_0^{-\alpha} + 1 \leq 0$.

(4) Let u and v be two prices for ISP B with $u < v$. To show that the best response is strictly increasing with the price, it is sufficient to show that $S(\text{BR}^t(v), u, t) < 0$ thanks to Relation (8):

$$S(\text{BR}^t(v), u, t) = (\text{BR}^t(v))^{\alpha-1}(u^{-\alpha} - v^{-\alpha}) \times ((1-\alpha)\text{BR}^t(v) + \alpha d_0t) < 0,$$

where we have used the fact that $S(\text{BR}^t(v), v, t) = 0$, and the (previously proved) fact that $\text{BR}^t(v) > \frac{\alpha}{\alpha-1}d_0t$.

(5) Finally, for every $t^+ > t^-$ we have

$$S(\text{BR}^{t^+}(p), p, t^-) = (t^- - t^+)\alpha d_0(\text{BR}^{t^+}(p))^{\alpha-1}(p^{-\alpha} + p_0^{-\alpha}) < 0,$$

hence the best response strictly increases with the transit price.

When $\alpha \leq 1$, $p = \infty$ is a strictly dominant strategy for each ISP (i.e., the unique best response whatever the strategy of the other ISP). This can be seen with the derivative of its profit (7) being strictly positive.

We now consider the case $\alpha > 1$. From Lemma 1, the best-response function of each ISP i is continuous, and bounded by strictly positive values. Let us denote by m_i (resp. M_i) the lower (resp. upper) bound of BR_i . Then consider the application

$$g : [m_A, M_A] \times [m_B, M_B] \mapsto [m_A, M_A] \times [m_B, M_B] \\ (p_A, p_B) \rightarrow (\text{BR}_A(p_B), \text{BR}_B(p_A)).$$

Since g is continuous and $[m_A, M_A] \times [m_B, M_B]$ is a compact convex subset of \mathbb{R}^2 , from Brouwer's fixed point theorem, it has a fixed point that constitutes a Nash equilibrium with strictly positive prices in $[m_A, M_A] \times [m_B, M_B]$.

From Lemma 1, the best response function, common to both ISPs ($\text{BR}_A = \text{BR}_B = \text{BR}$), is single valued and strictly increasing. Suppose that, at Nash equilibrium, $p_A \neq p_B$, for instance $p_A < p_B$ (should the indexes be permuted). Then $p_B = \text{BR}(p_A) < \text{BR}(p_B) = p_A$, hence a contradiction. This proves that every Nash equilibrium is symmetric.

Still from Lemma 1, prices at Nash equilibrium are strictly positive, hence the first order necessary condition implies that the derivative of an ISP profit is zero at such a point. From (6) it follows that

$$2(p^{\text{NE}})^{-\alpha} + p_0^{-\alpha} = \frac{\alpha(p^{\text{NE}} - d_0t)(p^{\text{NE}})^{-\alpha}}{\alpha(p^{\text{NE}} - d_0t) - p^{\text{NE}}}.$$

Plugging this relation into the utility function $U_i = \frac{(p^{\text{NE}})^{-\alpha}(p^{\text{NE}} - d_0t)}{2(p^{\text{NE}})^{-\alpha} + p_0^{-\alpha}}$ leads to $U_i = \frac{\alpha - 1}{\alpha}p^{\text{NE}} - d_0t$. In

particular, since, by Lemma 1, $p^{\text{NE}} > \frac{\alpha}{\alpha-1}d_0t$, then the profit is strictly positive.

Assume that the Nash equilibrium is unique. We denote by BR^t the best-response of ISPs with transit price t . We also denote by $p^{t,\text{NE}}$ the (symmetric) price set by ISPs at Nash equilibrium. Then, for all $z < p^{t,\text{NE}}$ and every $t > 0$ we have

$$z < \text{BR}^t(\text{BR}^t(z)). \quad (9)$$

Indeed, the best-response being lower bounded by a strictly positive value, (9) is verified when $z \rightarrow 0$. If there exists $z < p^{t,\text{NE}}$ that does not satisfy the inequality, then by continuity of the best-response functions, there exists $\hat{p} < p^{t,\text{NE}}$ for which $\hat{p} = \text{BR}^t(\text{BR}^t(\hat{p}))$. But this means that $(\hat{p}, \text{BR}^t(\hat{p}))$ is a Nash equilibrium, which contradicts the hypothesis of $p^{t,\text{NE}}$ being the unique Nash equilibrium.

Now, recall from Lemma 1 that the best-response is a strictly increasing function of t . Furthermore, it is a strictly increasing function of the other ISP's price. Let $0 < r < t$, and assume that $p^{t,\text{NE}} < p^{r,\text{NE}}$ (i.e., the equilibrium price is not increasing with the transit price). Then we have:

$$p^{t,\text{NE}} < \text{BR}^r(\text{BR}^r(p^{t,\text{NE}})) < \text{BR}^r(\text{BR}^t(p^{t,\text{NE}})) \\ < \text{BR}^t(\text{BR}^t(p^{t,\text{NE}})) = p^{t,\text{NE}}.$$

The first inequality comes from (9) together with the hypothesis $p^{t,\text{NE}} < p^{r,\text{NE}}$. The second is due to BR^r increasing with the transit price, and BR^r increasing with the price set by the opponent ISP. The last one is due to the increasingness of BR^r in the transit price, and the last equality stems from $p^{t,\text{NE}}$ being a symmetric Nash equilibrium price. Finally, this shows a contradiction. Hence $p^{t,\text{NE}}$ increases with t .

REFERENCES

- [1] C. Metz, "Interconnecting ISP networks," *Internet Computing, IEEE*, vol. 5, no. 2, pp. 74–80, 2001.
- [2] J. Musacchio, G. Schwartz, and J. Walrand, "A two-sided market analysis of provider investment incentives with an application to the net-neutrality issue," *Review of Network Economics*, vol. 8, no. 1, p. 3, 2009.
- [3] N. Economides, "The economics of the Internet backbone," *Handbook of Telecommunications. Amsterdam: Elsevier Publishers*, 2005.
- [4] P. Coucheny, P. Maillé, and B. Tuffin, "Network Neutrality Debate and ISP Inter-Relations: Traffic Exchange, Revenue Sharing, and Disconnection Threat," Inria, Rapport de recherche, Jan. 2012. [Online]. Available: <http://hal.inria.fr/hal-00659278>
- [5] N. Badasyan and S. Chakrabarti, "A simple game-theoretic analysis of peering and transit contracting among Internet service providers," *Telecommunications Policy*, vol. 32, no. 1, pp. 4–18, 2008.
- [6] J. Prufer and E. Jahn, "Dark clouds over the Internet?" *Telecommunications Policy*, vol. 31, no. 3–4, pp. 144–154, 2007.
- [7] J. Laffont, S. Marcus, P. Rey, and J. Tirole, "Internet interconnection and the off-net-cost pricing principle," *RAND Journal of Economics*, pp. 370–390, 2003.
- [8] D. Fudenberg and J. Tirole, *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.
- [9] M. J. Osborne and A. Rubinstein, *A Course in Game Theory*. MIT Press, 1994.
- [10] P. Reichl, R. Schatz, and B. Tuffin, "Logarithmic laws in service quality perception : Where microeconomics meets psychophysics and quality of experience," *Telecommunication Systems*, 2012, to appear, DOI: 10.1007/s11235-011-9503-7.
- [11] W. Kip Viscusi, J. Harrington, and J. Vernon, *Economics of Regulation and Antitrust*. MIT Press, 2005.
- [12] C. Berge, *Espaces topologiques. Fonctions multivoques*, ser. Collection Universitaire de Mathématiques. Paris: Dunod, 1959, vol. III.