Analysis of Price Competition under Peering and Transit Agreements in Internet Service Provision to Peer-to-Peer Users

Luis Guijarro, Vicent Pla, Jose R. Vidal, and Jorge Martinez-Bauset
Universidad Politécnica de Valencia (UPV), Camino de Vera s/n, 46022 Valencia, Spain

***Abstract*—Many studies on Internet Service Provider (ISP) interconnection make simplifying assumptions on the implementation of the service provision. Our work explicitly models the ISP service that is provided to users that run peer-to-peer applications and it analyses the behavior of competing ISPs. The ISPs have agreed to peer each other and each ISP has purchased transit service from one Internet Backbone Provider. The quality of service, the equilibrium prices and the market shares that the competition game yields are computed by means of our model. Our work assesses the strategy of an ISP which provisions its transit link against a competing ISP in terms of competitive advantage and social welfare. And it assesses the effect of the entrance of more competing ISPs.**

***Index Terms*—Competition, game theory, peering agreement, peer-to-peer, price of anarchy, quality of service.**

I. Introduction

Many studies of ISP interconnection use the game theory as the basis of the analysis, and make simplifying assumptions on the implementation of the service provision. There are a small number of studies, however, which explicitly model the service provision, such as [1]. Our contribution belongs to the latter. In addition, most studies assume that the web is the only application over the Internet; it has shown to be false. Peer-to-Peer (P2P) file-sharing applications have reached a level of popularity among Internet users such that the volume generated by such applications presently represents a significant portion of the overall Internet traffic [2]. Therefore, there is a need to model ISP behavior under the presence of P2P applications. The paper explicitly models the ISP service that is provided to users that run P2P applications and it analyzes the behavior of various ISPs which compete for P2P users in a local market. Reference [3] presents a model of one ISP which provides service to users and which uses the transit service provided by an Internet Backbone Provider (IBP). It measures the QoS that an ISP offers to P2P users as the probability that a query generated by a user results in a successful download. It analyzes the effect of the different model parameters on the transit capacity that is required to guarantee a minimum QoS. In our work, we use the same indicator for the QoS, but we extend the analysis by introducing competing ISPs in the local market. Our analysis uses a game theory model, which explains in a more realistic way the ISP behavior and the utility obtained by the P2P users.

This work was supported by the UPV through project PAID-06-09-2751 and by projects TIN2008-06739-C04-02/TSI and TIN2010-21378-C02-02

Other works, such as [1], model the interaction between two competing ISPs, in the presence of a peering agreement and of a local market of users. The objective of [1] is to model and to explain the tussle between policy-based routing —where ISPs enforce their transit agreements— and application-based routing —which is applied by P2P applications creating an overlay network. The QoS received by the users (web users and P2P users) is computed as a weighted average of the available capacities at the transit and peering links. In our work, the QoS received by the P2P users is computed in a different manner from [1] and it is based on [3]. We argue that our QoS measure outperforms the one used in [1] in two ways. Firstly, it gives more insight on the specifics of the P2P operation, so that it models in a more realistic way the utility received by the users. Note that the QoS measure in [1] does not capture these specifics and it is actually used for computing the QoS of both P2P users and web users. And secondly, as it will be shown in Section II, it is independent on the arrival process, whereas in [1] a Poisson arrival process is implicitly assumed since the rationale behind the proposed QoS measure relies on some well known results for the M/G/1 - PS queue and its application to model bandwidth sharing [4].

The main research question addressed in our work is whether a transit link provisioning strategy in the presence of P2P users is effective for those ISPs that agree to peer —as it is known for the case of web users [1]. The case of P2P users raise doubts since a stronger free riding effect is expected than it would be the case when only web users are present. In our work the strategy of an ISP which provisions its transit link is assessed in terms of competitive advantage and social welfare. Our work then proceeds to assess the effect of having more than two competing ISPs. To the best of our knowledge this issue has not yet been addressed, and the analysis of the $M > 2$ ISPs has not yet been even modeled. The paper is structured as follows. In Section II, the model is presented, which includes the service model and the game theory model. In Section III, the model equations for different scenarios are solved, and the results are discussed. And in Section IV, some final conclusions are drawn.

II. Model

We study the scenario shown in Fig. 1, which comprises two ISPs (ISP_1 and ISP_2) that compete for providing service to n users in a local market ($n = n_1 + n_2$). Each ISP

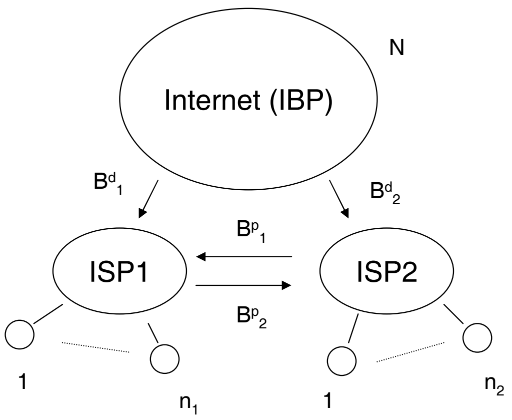

Fig. 1. Model scenario.

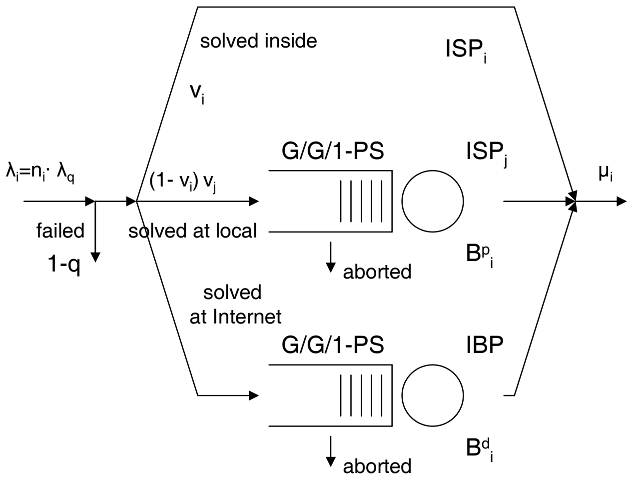

Fig. 2. Basic service model.

obtains transit service from one IBP, which offers a capacity of B_1^d to ISP$_1$ and B_2^d to ISP$_2$. And the ISPs agree to peer through unidirectional capacities denoted by B_1^p and B_2^p. Both local (n) and Internet (N) users are assumed to run only P2P applications. This assumption may correspond either to a common scenario where P2P traffic is responsible for a large fraction of the total traffic, as pointed above, or to a scenario where ISPs reserve a fixed amount of capacity for P2P applications.

A. Basic service model

We model the service that the P2P users receive as the queue system that is shown in Fig. 2, which extends the model developed in [3]. Fig. 2 illustrates the possible outcomes of queries generated by users subscribing to ISP$_i$. Part of the queries generated by the users might fail to locate any copy of the desired content within the P2P community as a whole. This can happen because of several reasons, which are usually not under the control of the ISP. We thus treat the probability that a query is solved (i.e. at least one copy of the desired object is found) as a given parameter, and denote it by q. In general, a fraction of the queries that are solved can be served from within ISP$_i$, while another fraction has to be served from peers residing outside ISP$_i$. Let ν_i be the probability that a solved query is served from within ISP$_i$. To simplify the analysis further, we assume that internal downloads (i.e. data transfers between peers belonging to the same ISP) always complete successfully. This corresponds to the best-case scenario of unlimited bandwidth within the ISP network. If an object cannot be served from within ISP$_i$, either the transit link of limited capacity B_i^d or the peering link of capacity B_i^p will be used. The probability for the former is $(1-\nu_i)\nu_j$ and for the latter is $1-\nu_i-(1-\nu_i)\nu_j$. Since each link is shared by all flows transferring objects from external peers (either local peers or Internet peers), using elastic rate adaptation mechanisms (e.g. TCP), we adopt a Processor Sharing model to describe their dynamics [5]. We consider that users are impatient and tend to abort downloads that do not receive a minimum throughput. This is motivated by the fact that many P2P applications allow the user to inspect the current download rate of an object and abort the transfer. It is reasonable to assume that the abort decision is taken in the early phase of the transfer, relative to the total time required to download the object. Under these assumptions, the shared link is always stable and no capacity is wasted by partial downloads. Furthermore, for the performance metric under consideration, the assumed impatient behavior makes the analysis independent of the arrival process type and the object size distribution. Thus, no specific assumptions on the arrival process type (e.g. Poisson) or object size distribution are needed and a general G/G/1 - PS queue model is considered. Moreover, we will assume that all four links are always fully utilized (i.e. there is always at least one transfer in progress), yielding a throughput equal to the link capacity, either B_i^d or B_i^p. And note that the assumption is consistent with our initial assumption of capacity reservation for P2P users. This assumption will be relaxed in the next subsection, where peering links have an unlimited capacity. Finally, let λ_q be the average content demand of a peer, which can be interpreted as the mean rate at which a user generates queries using one or more P2P applications. Let $\lambda_i = n_i\lambda_q$ be the aggregate generation rate of queries by the population of users belonging to the ISP$_i$. The aggregate system throughput μ_i, which is the rate at which objects are successfully retrieved by the users, can be expressed as

$$\mu_i = \lambda_i q_i \nu_i + B_i^p + B_i^d \tag{1}$$

It follows that the probability that a query generated by a user results in a successful download is $\sigma_i = \mu_i/\lambda_i$, which is given by

$$\sigma_i = q_i\nu_i + \frac{B_i^p}{\lambda_i} + \frac{B_i^d}{\lambda_i} \tag{2}$$

This probability will be used below in order to compute the quality of the service. We follow [3] and determine the probability ν_i that a query generated by a user is served by some other peer belonging to the same ISP. Due to the page limit, readers are referred to [6] for the detailed expressions. We derive an expression of σ_i which depends on α_i, where $\alpha_i \triangleq n_i/n$.

B. Unlimited peering service model

The above assumption that the four links are fully utilized is relaxed for the peering links in this subsection. The model is denoted unlimited peering service model. We argue that this model is more realistic than the basic service model. Note that while the transit cost depends on the traffic, the peering cost does not depend on the traffic and it is independent, to a certain degree, of the link capacity. Therefore, the ISP has no compelling need to limit the peering link capacity as it would be the case with the transit link. Nevertheless, the analysis in Section III will show that the results for the unlimited peering service model are a limiting case when compared with the results for the basic service model. Thus, the basic service model is not invalidated by the unlimited peering service model and it remains a simpler scenario which permits a more tractable formal analysis. The unlimited peering service model assumes that the peering capacity is large enough so that downloads from the competing ISP are always completed successfully. Therefore, data transfers across peering links are transmitted without losses. The model diagram differs from the basic service model diagram depicted in Fig. 2 in that the middle G/G/1 - PS queue is removed. The probability that a query generated by a user results in a successful download now becomes:

$$\sigma_i = \frac{\nu_i}{\lambda_i} = q_i \nu_i + q(1-\nu_i)\nu_j + \frac{B_i^d}{\lambda_i} \tag{3}$$

Readers are again referred to [6] for the detailed expressions. We just note that the expressions yield a cross dependence between σ_1 and σ_2.

C. Demand and offer models

From microeconomic theory, the user preferences are modeled by means of the utility function. We assume that users prefer the service which offers a better quality and which requires a lower payment, so that the proposed expression for the utility is $U_i \triangleq P_i - p_i$, where $P_i \triangleq \log(\beta\sigma_i + 1)$. σ_i is the quality of service computed in the previous section, β is a shape parameter, p_i is the price charged to each user, and the log function has been chosen to model diminishing returns as σ_i increases. Therefore, U_i is a function of p_i and α_i.

Each ISP receives a service income and incurs costs. Both incomes and costs are referred to a given period of time. Each ISP uses flat rate charging and charges a fixed price p_i to all its users. We follow [7], and assume that ISP interconnection costs are of two kinds: (1) transit costs, C_i^d, which comprise an initial startup cost —which we neglect— and a traffic related cost; and (2) peering costs, C_i^p, which are fixed costs related to the deployed equipment at the interconnection site. Given the fully utilized assumption for the transit links, which holds for both service models, the ISP$_i$ obtains a profit $\Pi_i = p_i n_i - C_i^d B_i^d - C_i^p$, which depends on p_i and α_i.

D. Game model and equilibrium

Based on the above models, we model the interaction between the two ISPs and the n users by means of an extensive game with perfect information of the type multi-leader-follower game. The leaders are the ISPs, which, in the first stage of the game, fix their prices p_1 and p_2. In the second stage, the followers, which are the n users, choose which ISP to subscribe to. The preferences of each player are the profit Π_i, for the ISPs, and the utility U_i for the users.

For the second stage, we assume that the number of users, n, is high enough that each individual subscription decision does not modify the utility obtained by the rest of the users. The outcome resulting from such user interactions is described by Wardrop's principle: demand is distributed in such a way that all users choose the ISP which provides them with the highest utility. As a result, the utility U_i (1) is the same for each ISP with positive demand; (2) is higher than the ISP with no demand; and (3) all ISPs which received some demand provide the same utility U_i. In our case, if we focus on the cases when $\alpha_1 \neq 0$ and $\alpha_1 \neq 1$, the Wardrop's equilibrium is given by

$$\begin{aligned} U_1(p_1, \alpha_1) &= U_2(p_2, 1-\alpha_1) \\ P_1(\alpha_1) - p_1 &= P_2(1-\alpha_1) - p_2 \end{aligned} \tag{4}$$

That means that no user will improve its utility by changing its subscription decision, remaining every other user by its ISP. This is indeed the Nash equilibrium condition.

In the first stage, the two ISPs compete for attracting users by setting their respective prices in a simultaneous and independent way. Therefore, we model the first-stage game as a standard one-shot game. The Nash equilibrium for this game is achieved when none of the ISPs can improve its profit by unilaterally changing its price. We can express profit as a function of the prices p_1 and p_2 so that the prices at the Nash equilibrium p_1^* and p_2^* will be such that

$$\begin{aligned} p_1^* &= \arg\max_{p_1} \Pi_1(p_1, p_2^*) \\ p_2^* &= \arg\max_{p_2} \Pi_2(p_1^*, p_2) \end{aligned} \tag{5}$$

Bringing together the Nash equilibrium conditions of the first and the second stages of our game, we state that the ISPs will set prices p_1^* and p_2^* and will obtain market shares α_1^* and $1-\alpha_1^*$ as determined by (5) and (4). The equilibrium conditions can be generalized for M ISPs ($M > 2$).

E. Social welfare

We now introduce the concept of social welfare [8] or total surplus [9], which is defined as the sum of the utilities of all agents in the systems (i.e. users and providers), and we study its maximum value. It is well known in game theory that agent selfishness, such as in a Nash equilibrium, does not lead in general to a socially efficient situation. As a measure of the loss of efficiency due to the divergence of user interests, we use the Price of Anarchy [10], which we denote as PoA and which we defined as the quotient between the maximum value of the social welfare and the social welfare obtained at the Nash equilibrium.

III. Analysis

The objective of solving the equilibrium conditions is to assess the provisioning strategy for the transit link and the increase of the number of competing ISPs. More specifically, in the following cases, we have fixed the set of parameters from [3] and [11]; the number of users, $n = 10^4$ and $N = 5 \cdot 10^7$; and the costs C^d and C^p from [7] and [11]. Then, the equilibrium equations (4) and (5) have been solved for different values of B_i^d or B_i^p chosen by each ISP, and the resulting prices p_1^* and p_2^* and the market shares α_1^* and $1 - \alpha_1^*$ have been computed —asterisks are removed hereafter in order to keep typography lighter. The resulting equilibrium has been evaluated by computing the quality of service σ_i and the utility U_i obtained by the users of each ISP and the profits Π_i gained by each ISP. And finally, the Price of Anarchy has been obtained.

A. Transit link provisioning

In this case, we assess if provisioning the transit link by one ISP is an effective strategy for gaining competitive advantage. We fix the peering capacities $B_1^p = B_2^p = 10$ objects/day. The transit capacity B_2^d varies from 100 to 2000 objects/day and B_1^d is kept fixed at 1000 objects/day. Fig. 3 depicts market shares α_1 and α_2, qualities of service σ_1 and σ_2, and user utility U_i as functions of the transit capacity B_2^d. Fig. 4 depicts prices p_1 and p_2, and ISP profits Π_1 and Π_2 as functions of the transit capacity B_2^d. We see that, as ISP_2 increases its transit capacity B_2^d:

1) ISP_2 increases its market share, improves its quality of service, and increases the equilibrium price and its profit.
2) ISP_1 looses market share, and, although it improves its quality of service, its equilibrium price lowers and its profit also decreases.
3) Users improve the utility they get from both ISPs.

Note that the improvement in ISP_1's quality of service is an indirect one, caused by the decrease of its market share. Note that, fewer users at ISP_1 means lower probability of a query being solved from within the ISP, but it also means that transit (B_1^d) and peering (B_2^p) capacities are shared among fewer users. The balance is positive and σ_1 improves. We conclude that transit link provisioning is an effective strategy for gaining competitive advantage, and this conclusion answers the research question posed in Section I. Indeed, even in the case where only P2P users are present and the competing ISPs have set peering agreements, the effect of free riding on the provisioning ISP is not enough to invalidate the transit link provisioning strategy.

B. Transit link provisioning: effect on social welfare

In this case, we analyze the effect of the transit link provisioning decisions on the social welfare. Both service models

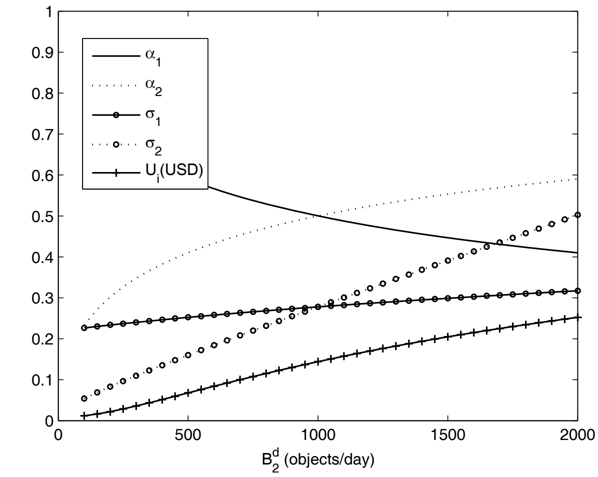

Fig. 3. ISP_1 and ISP_2 market share and quality of service and user utility as functions of ISP_2 transit link capacity.

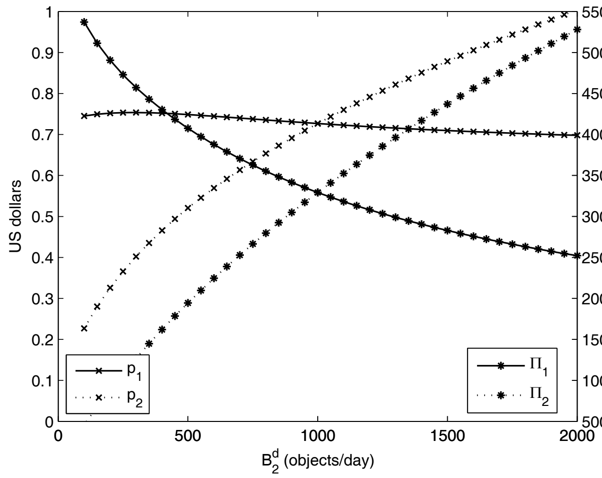

Fig. 4. ISP_1 and ISP_2 price and profit as functions of ISP_2 transit link capacity.

are incorporated. For the basic service model, four different scenarios are defined by the values 10 and 50 objects/day assigned to B_1^p and B_2^p. For the unlimited peering service model, the scenario is $B_1^p = B_2^p = \infty$ objects/day. In each scenario, the transit capacity B_2^d varies from 750 to 1250 objects/day, and B_1^d is kept fixed at 1000 objects/day. The Price of Anarchy (PoA), defined in Subsection II-E, is computed. Fig. 5 depicts the PoA as a function of the transit capacity B_2^d. We see that PoA exhibits a quasi-linear dependence on the absolute value of the difference $(B_2^d + B_2^p) - (B_1^d + B_1^p)$, so that: (1) The socially optimal point, that is, the point of minimum PoA, occurs when $(B_2^d + B_2^p) = (B_1^d + B_1^p)$; and (2) at the optimal point, the value of PoA decreases as the peering capacities increase, reaching the unity value for unlimited peering.

We conclude, firstly, that the competitive advantage that

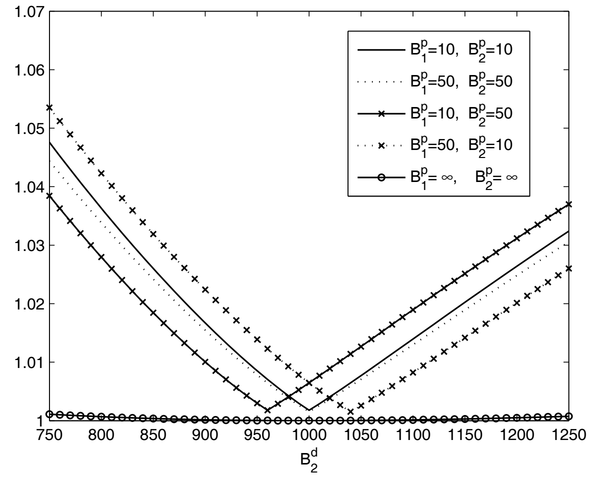

Fig. 5. Price of Anarchy as a function of ISP_2 transit link capacity.

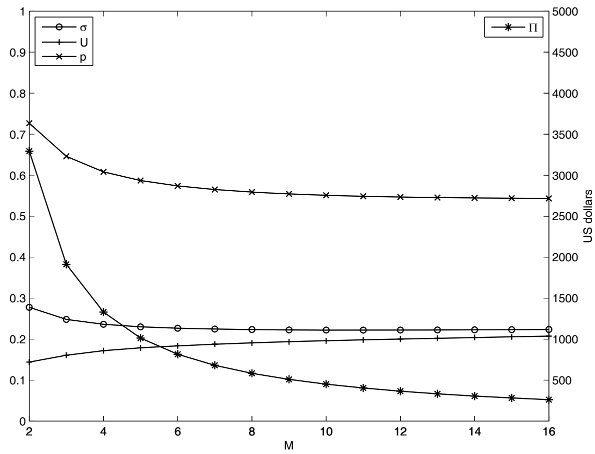

Fig. 6. ISP quality of service, price and profit, and user utility as functions of the number of ISPs.

ISP_2 obtains by provisioning its transit link is balanced by an increase in PoA; and secondly, that high peering capacities —i.e. high enough to be considered as unlimited— lead to best scenarios from the point of view of the social welfare attained.

C. Increasing competition: effect on user utility

We then proceed to increase the number M of ISPs competing in the local market and to assess the effect on the utility of the users. An analysis which extends that presented in Section II has been performed and the equations have been solved. The analysis assumes identical peering and transit capacities: note that, according to the results of the subsection III-B, this case is optimal in terms of social welfare. The model also assumes that the total transit capacity is kept constant, i.e. $B_i^d = B^d/M$. There is therefore an increase on the competition, while the number of users (n) and the resources available (B^d) are kept constant. We fix both the peering capacities $B_1^p = B_2^p = 10$ objects/day and the aggregated transit capacity $B^d = 2 \cdot 1000$ objects/day. The number of ISPs is varied from 2 to 16. Fig. 6 depicts the quality of service σ_i, users utility U_i, price p_i, and ISPs profit Π_i, as functions of M. We see that, as M increases, (1) the quality of service varies slightly, with a minimum around $M = 11$; (2) the price decreases; (3) the user utility increases; and (4) the profits decrease. We conclude that increasing M is beneficial for the users, but ISPs suffers a profit reduction. This result supports the desired outcome of the measures of most telecommunications regulatory authorities aimed at removing entry barriers and increasing competition.

IV. Conclusions

In this paper, we have presented an analysis of the economic interaction of ISPs which compete for providing service to users in a local market and which agree on a peering interconnection. The analysis is based on an explicit model of the ISP service that is provided to users that run P2P applications. From the analysis, the following conclusions have been drawn:

1) Transit link provisioning leads to gaining competitive advantage.
2) The transit link provisioning strategy exhibits a trade-off in terms of social welfare, since the competitive advantage is gained at the expense of an increase in the Price of Anarchy.
3) The entrance of competitors in the local market is beneficial for the users in terms of the utility they receive.

References

[1] J. Wang, D. Chiu, and J. Lui, "A game–theoretic analysis of the implications of overlay network traffic on ISP peering," *Computer Networks*, vol. 52, no. 15, pp. 2961–2974, 2008.
[2] G. Haßlinger, "ISP platforms under a heavy peer-to-peer workload," *Peer-to-Peer Systems and Applications*, pp. 369–381, 2005.
[3] M. Garetto, D. Figueiredo, R. Gaeta, and M. Sereno, "A modeling framework to understand the tussle between ISPs and peer-to-peer file-sharing users," *Performance Evaluation*, vol. 64, no. 9-12, pp. 819–837, 2007.
[4] N. Benameur, S. Ben Fredj, S. Oueslati-Boulahia, and J. Roberts, "Quality of service and flow level admission control in the internet," *Computer Networks*, vol. 40, no. 1, pp. 57–71, 2002.
[5] J. Roberts *et al.*, "Traffic theory and the Internet," *IEEE Communications Magazine*, vol. 39, no. 1, pp. 94–99, 2001.
[6] L. Guijarro, V. Pla, J. Martinez-Bauset, and J. Vidal, "Analysis of Price Competition under Peering and Transit Agreements in Internet Service Provision to Peer-to-Peer Users," *Technical Report TELPOL12d, April*, 2009.
[7] W. Norton, "A business case for ISP peering," *White Paper (v1. 3), February*, 2002.
[8] P. Maillé and B. Tuffin, "Analysis of price competition in a slotted resource allocation game," in *Proc. of IEEE INFOCOM*, 2008.
[9] H. Varian, *Microeconomic analysis*. Norton New York, 1992.
[10] K. Papadimitriou and E. Koutsoupias, "Worst-case equilibria," in *STACS*, 1999, pp. 404–413.
[11] K. Gummadi, R. Dunn, S. Saroiu, S. Gribble, H. Levy, and J. Zahorjan, "Measurement, modeling, and analysis of a peer-to-peer file-sharing workload," in *Proceedings of the nineteenth ACM symposium on Operating systems principles*. ACM, 2003, p. 329.