

Price Setting in Two-Sided Markets for Internet Connectivity

Thorsten Hau and Walter Brenner

Institute of Information Management, University of St. Gallen
Müller-Friedberg-Strasse 8, CH-9000 St. Gallen, Switzerland
Tel.: +41 (0)71 224 3803; Fax: +41 (0)71 224 3296
thorsten.hau@unisg.ch

Abstract. Due to a lack of incentives, Internet peerings are a notorious bandwidth bottleneck. Through the use of direct interconnection and content delivery networks, content providers are able to provide better services to their customers. These technologies have a profound impact on the business models of internet service providers. Instead of competing for consumers and keeping uplink connection costs low, ISPs face a two-sided market in which they compete for EUs and generate revenues on the CP side of the market. This work presents a formal model for the providers' pricing decision towards content providers and discusses consequences for the Internet.

Keywords: Quality of Service, Network Economics, Peering, Internet.

1 Introduction

The Internet is made up of many independent sub-networks – so called “autonomous systems” (AS). Generally speaking these ASs correspond to different carriers or Internet service providers (ISPs): firms that own and operate the infrastructure (cables and routers) that make up the Internet. These ISPs have customers who are either content providers (CP) with mostly outgoing traffic or end-users (EU) with mostly incoming traffic. To form the Internet, each ISP offers all of its customers' connectivity to all of the other ISPs' customers. In order to uphold this universal connectivity, the ISPs have to exchange user traffic, an activity that is governed by contractual agreements between the ISPs and physically enabled by infrastructure that interconnects their networks. Thus, even though the Internet consists of many independent sub-networks, each user can reach every website on the web. However, it is also well known, that ISPs are not very cooperative in their peering behavior [7]. The decision to interconnect usually is more expensive for one party than for the other and therefore peerings tend to have smaller capacity than what would be optimal.

The rules regulating the exchange of traffic between ISPs have been subject to extensive treatment in the literature. Issues like hot potato routing [22, 24] and determination of access charges have been extensively studied and are quite well understood [8, 10, 17]. However, the available literature studies an idealized Internet in which




Fig. 1. Physical Interconnection 1

there are EUs, CPs and ISPs that have relations as depicted in fig. 1. The Internet is modeled as a strictly hierarchical system in which traffic flows from a CP to its ISP, is exchanged with the requesting user’s ISP and is then sent to the EU.

A key characteristic of this setup is that traffic is exchanged through a peering point. Depending on the contract between two ISPs this traffic exchange may happen in exchange for a payment or “for free”. Due to lacking incentives to extend peering capacity sufficiently, these peerings represent major traffic bottlenecks [1, 7, 15].

In contrast to the available work [17, 18, 23], this paper focuses on two important variations of this idealized model of Internet infrastructure as shown in figures 2 and 3. The existing literature has ignored the possibility that content providers and terminating internet service providers interconnect directly.

There are two modes of “direct interconnection” that we will consider. Firstly, a content provider can directly buy transit from the terminating ISPs, thus effectively paying them for preferential access to end-users. This practice shown in fig. 2 is called multi-homing (MH) and contributes to exponential growth of routing tables [5]. Secondly, content delivery networks (CDNs) shown in fig. 3 are a popular way to enhance the flow of information on the Internet. A CDN uses local caches to keep distributed images of content close to EUs without the need to traverse several ISPs’ networks [26]. Both technologies provide viable means to improve the speed and reliability of data transport from a CP’s website to EUs. They allow bypassing peerings and gaining more direct access to the EUs, thus increasing the probability of timely delivery of data to the end-user. The motivation to use CDN or MH is to provide better quality of service (QoS) with the following chain of causality: Traversing peerings degrades user experience by creating delays \rightarrow QoS access to EUs through CDN and MH creates better user experience \rightarrow more visitors to a website \rightarrow higher revenues from selling ad space on the website.

This paper uses the economic theory of two-sided markets to understand the pricing decision an ISP has to make with respect to the charges levied on the CP side of the market in settings such as those in Fig. 2 and 3. Neither the Internet as a whole, nor individual internet service providers (ISPs) can straightforwardly be considered platforms that optimize their revenue from two sides of a market. With the standard Internet business model, each terminating ISP lacks the power to charge content




Fig. 2. Physical Interconnection 2




Fig. 3. Physical Interconnection 3

providers that are signed up with another ISP. There are technical as well as contractual barriers to charge some remote content provider for single data packets it sends to an ISP's network. The access charge (the interconnection fee) exchanged between two ISPs is only an imperfect tool to exploit an access monopoly on the Internet due to the fact that it is often reciprocal or zero for external reasons [2, 13, 17]. This is a key difference between the Internet and telephone services (PSTN), where for each call, sender and receiver can be identified and billed per unit of time and a per unit settlement between providers is possible [4, 6, 10]. With CDNs or MH, the property of PSTN that the participating parties can be identified (and are billable customers) is recreated.

The CDN is a third party mediating between CP and ISP but the ISP can charge the CDN for delivery of traffic which will pass this cost on to its CPs. This situation is different from the case of access charges between different providers (as analyzed by [17] with a focus on the Internet or [9] with a focus on PSTN) because in a CDN relationship there is no reciprocity or two way access which is an important condition for that model to be applicable. For the rest of this article we simplify the role that CDNs play on the Internet by treating them as pure mediators between atomic CPs and ISPs. They aggregate CP demand but do not engage in strategic actions. This simplification allows us to model the situation of fig.1 and fig. 2 in the same way. In the last section of this work we sketch a path to relaxing this rather strong assumption.

The paper is structured as follows: Firstly we review the relevant literature on two-sided markets and related topics in telecommunications pricing. Then we explain the abstracted situation we wish to understand and motivate our use of a two-sided market model. Thirdly, we present a formalized model for an ISP facing a two-sided market, deriving results from the market setup. We derive optimal prices charged by ISPs to CPs and CDNs that wish to directly interconnect with them. Lastly we summarize our findings and discuss implications and future research topics.

2 Literature Review

Armstrong's discussion of competition in two-sided markets [3] provides much of the foundation for this work. Two-sided markets are markets where a platform optimizes profit across two distinct sets of customers instead of just one. In the credit

card industry, the card issuing company would be the platform and the merchants accepting the card constitute one group of customers while the buyers using the card to pay form the other. Armstrong analyzes three distinct settings with different customer behaviors and levels of platform competition. The situation relevant for this work is termed “competitive bottleneck”: One group of customers can use any number of platform providers simultaneously, while the other group chooses only one of the competing platforms. In our problem, this situation corresponds to EUs being subscribed to only one single ISP while CPs can deal with any number of ISPs at the same time.

Rochet et al. [21] provide a comprehensive overview of the current literature on two-sided markets. They define two-sided markets as markets in which not only the total price but also the price structure influences the number of transactions on the market. For the case at hand, the ISP provides the platform on which transactions between EUs and CPs can take place. They also provide definitions for membership and usage externalities. In the first case one party profits from the sheer presence of the other, while a usage externality is a network effect that arises in an transaction between members of the two sides. They also discuss the effects of fixed and variable prices on the platform. Since variable prices reduce the externality exerted by one group of customers on the other, participation incentives are reduced.

Laffont et al. [17] are not directly concerned with two-sided markets. This work analyzes the access charge paid from one ISP to another for passing traffic on to that ISP’s network. In their model the ISP optimizes the prices it charges to CPs and EUs subject to the access charges it pays (for sending traffic to an EU on another ISP’s network) and receives (for terminating traffic with its own EUs). In their model the access charge turns out to be a pure tool for reallocating termination costs between EUs and CPs. In the common case of zero access charges all termination costs are born by the EU which corresponds to a subsidy to CPs.

Musacchio et al. [19] compare the effects of single and multi-homing of CPs. They provide explicit formulations of welfare under both regimes and offers results for an economy with many ISPs. However, they do not model EU and CP demand separately but base their model on the assumption of click rates of EUs as a measure of demand for both customer groups and only differentiate CPs from EUs via the per-click price.

3 Problem Description

This work uses the theory on two-sided markets to explore two special cases of interconnection that are different from the symmetric and reciprocal case studied by [17]. The standard model of Internet traffic exchange as shown in fig. 1 follows the pattern $CP \rightarrow ISP_o \leftrightarrow_a ISP_i \leftarrow EU$ (t =terminating, o =originating, a =access charge) as shown in fig. 4. CP and EU pay their respective ISP and the ISPs exchange traffic for a fee a . This scheme ignores the source of the CP’s funding and emphasizes the analysis of the inter-ISP settlement a , which has an influence on the prices paid to the ISPs. By contrast, this work focuses on the setup $EU \rightarrow ISP_i \leftarrow CP \leftarrow Adv$ as shown in fig. 5. EUs derive utility from high quality of service (QoS) access to CPs’ websites while




Fig. 4. Business View 1




Fig. 5. Business View 2

CPs generate profits from selling ad space to third parties. There is no monetary flow between CPs and EUs. Both, however, may exchange money with the ISP, which acts as a profit maximizing platform. This situation corresponds to the majority of today's Internet business models. CPs create websites that appeal to many EUs, thus generating page views that translate into value of ad-space on that site. (Fig. 1, 2 and 3 focus on physical interconnection while Fig. 4 and 5 depict the business view on connection relationships.)

This CP business model has received wide attention in the two-sided markets literature as it corresponds to the business model of newspapers [3, 20, 21]. This work, however, does not consider the business model of the CP but that of the ISP. CPs pay a transaction-independent price for direct connection to EUs through buying bandwidth from the terminating ISP. EUs on the other hand pay a flat rate fee to the ISP to be connected to the Internet and no transaction based fee for viewing content. There is no payment between EU and CP. The case with payments between the CPs and the EUs has been analyzed in [12].

In the sense of the two-sided market literature we have the following setup: Platform = ISP, single homing side = EU, multi-homing side = CP. The platform charges both sides a lump sum fee for facilitating transactions. This is more reasonable than a linear fee since for EUs, flat rates are the common pricing model and CPs commonly buy a certain bandwidth or a fixed traffic volume. Furthermore the price for Internet services delivered by an ISP might depend on the data volume but rarely on the value of a transaction. Therefore we assume that there is no linear payment that reduces the size of the externalities exerted on the other side of the market, respectively.

4 ISPs as Platforms in Two-Sided Markets

The analysis in this section is related to the competitive bottleneck case of [3]. Competitive bottlenecks arise, when one firm has a monopoly over access to certain customers. Suppose there are two ISPs in a region denoted $ISP_i, i \in \{1, 2\}$. There are also two groups of agents. Group one agents are called end-users (EUs) while group two members are called content providers (CPs) or websites. There is a fraction n_j^i of agents of Group j participating on platform i . In other words, ISP_i has n_1^i subscribed EUs and n_2^i directly interconnecting customers from the CP side.

The setup is such that two ISPs are present in a market and serve two distinct groups of EUs with Internet connectivity. EUs are single-homing with their ISP. This

means that they are only subscribed with one ISP at a given time. CPs on the other hand multi-home. They may be connected to zero, one or two ISPs in order to reach potential customers (EUs).

To analyze this situation we start by modeling the target function of two ISPs that compete in a market for EUs. The ISPs maximize their respective profits. Symbolically,

$$\pi^i = n_1^i p_1^i + n_2^i p_2^i - C^i(n_1^i, n_2^i), i \in \{1; 2\} \quad (1)$$

which is a function of the number of EUs times the price they have to pay, plus the number of CPs times the price they have to pay minus the cost for connecting the two types of customers. The fraction $n_1^i \in [0, 1]$ of EUs that are customers of ISP_{*i*} is given as a function of the utilities offered by the two ISPs:

$$n_1^i = \phi^i(u_1^i, u_1^j) = \frac{1}{2} + \left(\frac{u_1^i - u_1^j}{2t} \right), \forall i \neq j. \quad (2)$$

The function ϕ^i is thus increasing in the first argument and decreasing in the second. Note that $n_1^i + n_1^j = 1$ holds since EUs do not multi-home. To specification of EU-demand in equation (2), i.e. the fraction of EUs that are signed up with either ISP is described in a Hotelling [14, 25] way. This implies that the two ISPs share the market equally if they are undifferentiated from the consumers' point of view. If, however, one ISP offers superior utility, it can capture more than half of the market.

The utility EUs get from subscribing to ISP_{*i*} is given by

$$u_1^i = U^i(n_2^i) - p_1^i = \alpha_i n_2^i - p_1^i. \quad (3)$$

It equals the gross utility they get from being connected with superior QoS to n_2^i directly interconnected CPs minus the price they have to pay for that connection. The function U^i is increasing in n_2^i since more content in better quality is always better than less. The parameter α_i can be interpreted as the utility an EU derives from being able to reach one high QoS CP. The EUs perceive the ISP with more CPs connected with QoS as providing a better connection to the Internet.

The fraction $n_2^i \in [0, 1]$ of CPs that is connected to ISP_{*i*} is given by

$$n_2^i = \phi_2^i(n_1^i, p_2^i) = 1 - F(\gamma^i) \text{ with } \gamma^i = \frac{p_2^i}{n_1^i}. \quad (4)$$

It is a function of the number of EUs that can be reached through ISP_{*i*} and the price charged. The number of CPs the ISP can persuade to directly interconnect depends on the parameter $\gamma^i = p_2^i/n_1^i$. This parameter is calculated as the fraction of the fixed price for connectivity over number of reachable EUs. Thus it can be interpreted as the perceived price per EU. The distribution $F(\gamma^i)$ then yields the fraction of CPs that are not willing to pay that price and $1 - F$ yields the fraction of CPs that are willing to pay that price because their expected revenue per EU covers the expense.

The CPs do not deal exclusively with a single ISP but may be connected to zero, one or two ISPs, depending on their participation constraint being fulfilled. Therefore in general $n_2^i + n_2^j \neq 1$.

While equation (4) only depends on factors under control of ISP_i, equation (2) also depends on factors controlled by the other ISP. This reflects the fact that there is competition for EUs, but none for CPs.

Costs for interconnection are defined as

$$C^i(n_1^i, n_2^i) = cn_2^i \quad (5)$$

This implicitly includes the assumption that the cost of the access network is not part of the considerations for interconnecting with CPs. This assumption is justified by the fact that access networks largely represent sunk costs.

Now, in order to solve the ISPs' optimization problem $\max \Pi^i$, assume that the platforms have reached an equilibrium and offer utility \hat{u}_1^i to their \hat{n}_1^i EUs, respectively. That is, we keep these values fixed while varying the others. This corresponds to today's situation in many markets for DSL or cable. There is some churn, but by and large networks operate in saturated markets with stable customer numbers. Since (4) defines n_2^i as a function of p_2^i , we can eliminate p_2^i and only have n_2^i left as a dependent variable. Thus, given an equilibrium $(\hat{u}_1^i, \hat{n}_1^i)$, we can solve for the optimal number of CPs n_2^i .

Rewriting equation (3) as $p_1^i = U^i(n_2^i) - u_1^i$ we can insert this expression into (1) to get

$$\begin{aligned} \Pi^i &= \hat{n}_1^i (U^i(n_2^i) - u_1^i) + \overbrace{p_2^i (1 - F(p_2^i / \hat{n}_1^i))}^{n_2^i} - C(\hat{n}_1^i, n_2^i) \\ &= \hat{n}_1^i (\alpha_1 n_2^i - \hat{u}_1^i) + (p_2^i - c) n_2^i \end{aligned} \quad (6)$$

This expression shows that given an arbitrary equilibrium we can explicitly write the profit of the platform as a function of the price charged to its group two customers (i.e. CPs). The platform can thus easily calculate the optimal price and the resulting number of CPs, given its current competitive situation on the EU side of the market.

To give a concrete example, we define the distribution F and explicitly calculate the profit maximizing price p_2^i . Let the distribution function F be given by the probability density function $f(\gamma) = 1/\tau, \forall \gamma \in [0; \tau]$ of the uniform distribution. γ represents the expected revenue from ad-clicks per EU and τ represents the maximum price a CP is willing to pay for access to such an EU. The corresponding cumulated distribution function is

$$F = \gamma/\tau = p_2^i / n_1^i \tau \quad (7)$$

Any other distribution function would work as well. However, the normal distribution for example is not easily manipulated and thus would only allow a numerical solution to the problem at hand.

Now we insert (4) and (7) into (6)

$$\begin{aligned}\Pi^i &= \hat{n}_1^i(\alpha_1 n_2^i - \hat{u}_1^i) + (p_2^i - c)n_2^i \\ &= \hat{n}_1^i(\alpha_1 n_2^i - \hat{u}_1^i) + [(1 - n_2^i)\tau\hat{n}_1^i - c]n_2^i\end{aligned}\quad (8)$$

and find the maximizer of the resulting expression:

$$\begin{aligned}\frac{\partial \Pi^i}{\partial n_2^i} &= \hat{n}_1^i \alpha_1 + (1 - 2n_2^i)\tau\hat{n}_1^i - c \stackrel{!}{=} 0 \\ n_2^i &= \left(1 - \frac{c - \hat{n}_1^i \alpha_1}{\tau\hat{n}_1^i}\right) \frac{1}{2}.\end{aligned}\quad (9)$$

This is the optimal number of CPs the ISP should allow on its platform (since the 2nd order condition for a maximum holds). Together with (4) and (7) this yields the optimal price to CPs

$$p_2^i = \frac{1}{2}(c + \hat{n}_1^i \tau - \hat{n}_1^i \alpha_1). \quad (10)$$

Therefore, CPs pay a price that is calculated on the basis of the cost they cause, increased by a factor relating to their per-EU-valuation and decreased by the externality they exert on the EUs. The factor 1/2 should not be over-interpreted since it is an artifact of the definition of the distribution function in (7).

We thus have calculated the optimal number of CPs and the optimal price that an ISP should charge atomistic and ex-ante identical CPs for quality interconnection.

5 Conclusions and Further Research

To sum up, we have firstly explained two phenomena of the Internet that fundamentally change the way CPs and EUs are interconnected. CDNs and MH foster more direct links between these two user groups with only one mediating ISP instead of many. Employing the theory of two-sided markets we then went on to show how direct interconnection puts the ISP into a position to charge CPs directly. In the main section we showed how the optimal price p_2^i can be calculated for any given equilibrium on the EU side of the market.

While today it is uncommon to explicitly charge content providers for delivering traffic to their customers, there are clearly developments in the marketplace that can be understood in the above context. Google's effort to provide free W-Lan to customers in the US is only one example. Google wants to control the platform over which its content is delivered so that the profits it makes on the advertisement side cannot be extracted by ISPs.

To interpret the results obtained, let's first compare the predicted price to today's bill and keep regime. In today's peering agreements between ISPs, the fee for carrying traffic is very often zero. As [17] point out, this corresponds to a subsidy to CPs, since EUs carry most of the transmission cost. In our two-sided market framework on the other side, the CPs have to bear the cost they cause. They may be furthermore charged by the ISP, depending on their willingness to pay. This charging is balanced

by a “bonus” for the externality they exert on the ISP’s EUs. Since the difference between being subsidized and paying bottleneck prices can be quite large, there will probably be a transitional period before ISPs can leverage their whole power in charging CPs. However, the presence of charges to content providers in itself does not represent a market failure. As long as ISPs are competing in the market for EUs, the profits they make on the CPs are used to compete in the EU market [3]. A waterbed effect might occur, but would merely be a sign of imperfect competition in the EU market [11].

Secondly, the last term in (10) illustrates a very interesting result. Imagine that the ISP could perfectly discriminate between two different groups of CPs. The group that exerts a higher externality on the customers through its presence would pay a lower price than the group with the lower externality effect. Thus, CPs that are very important to EUs will pay a low price to the bottleneck ISP, while those CPs, the presence of which is less valued by EUs, will pay a high price for access to EUs. Thus, a power balance could develop, in which CPs are charged by the network if they have low market power; or charge the network, if their content is highly desired by EUs.

Lastly, look again at the externality exerted by CPs. Here might lie an interesting option for future ISP business models. The ISP could try to capture some of the externality. This could happen for example through transaction dependent charges. Aside from contractual problems this would fulfill many ISPs’ long standing vision to capture some of the profits of the content business. This development can already be witnessed in the mobile sector where Vodafone provides high quality ad-financed content to its customers.

An important aspect of this work that requires further research is the effect of the two-sided markets phenomena on the quality of standard peerings. As it stands today, peerings do not generate revenue for ISPs but only costs. With revenues from direct interconnection there is obviously a strong incentive for ISPs to move as many CPs as possible to a paying interconnection model. The ultimate consequence of this would be that, in order to foster a self selection process, standard peering quality would be considerably degraded to make sure that all customers with a willingness to pay are in the paying group. While such price discrimination is welfare enhancing, it is crucially important the market for EUs is competitive since otherwise, ISPs are in a position to appropriate rents.

This paper demonstrates the use of two-sided market theory to analyze the decision problems faced by Internet service providers in more complex setups than the standard peering scenario examined in earlier works. A first analysis demonstrates that new business models such as content delivery networks and multi-homing can fundamentally change the rules for interconnection pricing. This work thus extends the work on Internet interconnection [17] and the work on voice interconnection such as [16] or [2] (as well as the references cited therein).

As this is only a first step towards a thorough understanding of the new rules of interconnection pricing brought about by new interconnection regimes, there remains considerable work to be done:

Firstly, the presented analysis cuts short some more in depth equilibrium analysis by assuming a market equilibrium as given.

Furthermore, a more thorough analysis of the effects of the ISPs actions on the secondary markets for advertisements would be interesting. How do two vertically dependent two-sided markets interact?

In a similar line of thought, the aspect that CDNs are intermediaries between ISPs and CPs has been used as a starting point of the analysis but is then abstracted from in the further analysis. This can be justified by assuming that CDNs only pass on costs but their role certainly deserves more attention, especially since CDNs are potent players in the Internet market. A further topic to be analyzed is the role of peer to peer traffic.

The paper has shown an aspect of the quality of service debate that has been under-researched. The market for Internet interconnection has a considerable influence on the deliverable quality of Internet services. Understanding these markets (the contribution of this work) and “engineering” them to function better (future research) propose challenging research topics that might shape the next generation of networks.

References

1. Akella, A., Seshan, S., Shaikh, A.: An empirical evaluation of wide-area Internet bottlenecks. In: Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement, Miami Beach, FL, USA (2003)
2. Armstrong, M.: Network Interconnection in Telecommunications. *The Economic Journal* 108, 545–564 (1998)
3. Armstrong, M.: Competition in two-sided markets. *RAND Journal of Economics* 37, 668–691 (2006)
4. Berger, U.: Bill-and-keep vs. cost-based access pricing revisited. *Economics Letters* 86, 107–112 (2005)
5. Bu, T., Gao, L., Towsley, D.: On characterizing BGP routing table growth. *Computer Networks* 45, 45–54 (2004)
6. Cambini, C., Valletti, T.M.: Network competition with price discrimination: ‘bill-and-keep’ is not so bad after all. *Economics Letters* 81, 205–213 (2003)
7. Cremer, J., Patrick, R., Tirole, J.: Connectivity in the Commercial Internet. *The Journal of Industrial Economics* 48, 433–472 (2000)
8. DeGraba, P.: Reconciling the off-net cost pricing principle with efficient network utilization. *Information Economics and Policy* 16, 475–494 (2004)
9. Economides, N., Lopomo, G., Woroch, G.: Strategic Commitments and the Principle of Reciprocity in Interconnection Pricing. Stern School Business, NYU (2008)
10. Gans, J.S., King, S.P.: Using bill and keep interconnect arrangements to soften network competition. *Economics Letters* 71, 413–420 (2001)
11. Genakos, C., Valletti, T.: Testing the ‘Waterbed’ Effect in Mobile Telephony. CEIS Working Paper No. 110 (2008)
12. Hau, T., Wulf, J., Zarnekow, R., Brenner, W.: Economic Effects of Multi Homing and Content Delivery Networks on the Internet. In: Proceedings of the 19th ITS European Regional Conference, Rome (2008)
13. Hermalin, B.E., Katz, M.L.: Sender or Receiver: Who Should Pay to Exchange an Electronic Message? *The RAND Journal of Economics* 35, 423–448 (2004)
14. Hotelling, H.: Stability in Competition. *Economic Journal* 39, 41–57 (1929)
15. Kushman, N., Srikanth, K., Katabi, D.: Can you hear me now?! it must be BGP. *SIGCOMM Comput. Commun. Rev.* 37, 75–84 (2007)
16. Laffont, J.-J., Rey, P., Tirole, J.: Network Competition: I. Overview and Nondiscriminatory Pricing. *The RAND Journal of Economics* 29, 1–37 (1998)

17. Laffont, J.-J., Marcus, S., Rey, P., Tirole, J.: Internet Interconnection and the Off-Net-Cost Pricing Principle. *The RAND Journal of Economics* 34, 370–390 (2003)
18. Laffont, J.-J., Rey, P., Tirole, J.: Competition between telecommunications operators. *European Economic Review* 41, 701–711 (1997)
19. Musacchio, J., Walrand, J., Schwartz, G.: Network Neutrality and Provider Investment Incentives. In: Conference Record of the Forty-First Asilomar Conference on Signals, Systems and Computers, 2007. ACSSC (2007)
20. Rochet, J.C., Tirole, J.: Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1, 990–1029 (2003)
21. Rochet, J.-C., Tirole, J.: Two-sided markets: a progress report. *The RAND Journal of Economics* 37, 645–667 (2006)
22. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
23. Shakkottai, S., Srikant, R.: Economics of network pricing with multiple ISPs. *IEEE/ACM Trans. Netw.* 14, 1233–1245 (2006)
24. Teixeira, R., Shaikh, A., Griffin, T., Rexford, J.: Dynamics of hot-potato routing in IP networks. In: Proceedings of the joint international conference on Measurement and modeling of computer systems (2004)
25. Tirole, J.: *The Theory of Industrial Organization*. MIT Press, Cambridge (1988)
26. Vakali, A., Pallis, G.: Content delivery networks: status and trends. *IEEE Internet Computing* 7, 68–74 (2003)