Multihoming, Content Delivery Networks, and the Market for Internet Connectivity

Abstract

Peering points between different Internet service providers (ISPs) are among the bottlenecks of the Internet. Multihoming (MH) and content delivery networks (CDNs) are two technical solutions to bypass peering points and to improve the quality of data delivery. So far, however, there is no research that analyzes the economic effects of MH and CDNs on the market for Internet connectivity. This paper develops a static market model with locked-in end users and paid content. It shows that MH and CDNs create the possibility for terminating ISPs to engage in monopolistic pricing towards content providers, leading to a shift of rents from end users and content providers to ISPs. Implications for future innovations are discussed.

Keywords: Internet economics, Quality-of-service, Price discrimination, Paid content, Innovation, Net neutrality

1. Introduction

Research on pricing of data transport has its roots in the literature on telecommunications. Early work on pricing of voice communications established the corner-stones of our thinking about communications pricing. A prominent example for this is the focus on access charges (Laffont et al., 1998), i.e. the price one provider pays to the other for the termination of traffic with an end user. The present paper departs from this "classical" view on communications pricing by also considering content providers (CPs), end users, and content delivery networks (CDNs) instead of only the inter carrier settlement. This issue is not covered in the existing literature. We show how an Internet service provider (ISP) with access to end users can discriminate against CPs and charge monopoly prices for termination. The discussion is related to and uses results from research on one- and two-way access (Buehler & Schmutzler, 2006; Gans, 2006), strategic network pricing (Shrimali, 2008), two sided markets (Armstrong, 2006; Rochet & Tirole, 2006), vertical integration (Rey & Tirole, 2007; Tirole, 1988), telecommunications pricing (Laffont et al., 2003; MacKie-Mason & Varian, 1995; Shakkottai & Srikant, 2006), net neutrality (Crowcroft, 2007; Sidak, 2006; Wu, 2003) and quality of service (QoS) (Soldatos et al., 2005; Wang, 2001).

The existing literature on telecommunications pricing has ignored the possibility that CPs and terminating ISPs directly interconnect. In contrast, consider the following two situations: First, it is commonplace that CPs directly

Preprint submitted to Telecommunications Policy *March 17, 2011*

buy transit from terminating ISPs, thus effectively paying them for preferential access to end users. This practice is called multihoming (MH) and plays a role in the exponential growth of routing tables (Bu et al., 2004). Second, CDNs 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 end users without the need to traverse several ISPs' networks (Pathan & Buyya, 2007; Vakali & Pallis, 2003). Both technologies provide viable means to improve the speed and reliability of data transport from a CP's website to end users. This is due to the fact that peering points, i.e. the points of interconnection between the networks of two ISPs are among the notorious bottlenecks of the Internet (Akella et al., 2003). Both technologies serve as ways to bypass these peerings and to gain more direct access to end users, thus increasing the probability of timely delivery of data to the end user.

The remainder of this paper is structured as follows: Section 2 explains the relevant entities of the Internet that we need for a formal model. Section 3 presents a formalized treatment of six scenarios that shows how MH and CDNs affect ISPs' incentives to price traffic in comparison to the standard situation with peering. The model is static with locked-in end users who cannot switch their provider. In section 4 we discuss the consequences of our model and sketch out an agenda for further research.

2. The Market for Internet Connectivity

Figure 1 and Figure 2 show in an idealized manner the structure of the Internet (Shakkottai & Srikant, 2006; Uludag et al., 2007). Figure 1 focuses on the interconnection aspect. Several ISPs interconnect with each other through points of interconnection (denoted by "I"). Figure 2 focuses on the hierarchical structure of the Internet. Data first flows up the hierarchy from a CP to its ISP and across a peering point back down via an ISP to the end user (EU).

Figure 1: Interconnection structure of the Internet

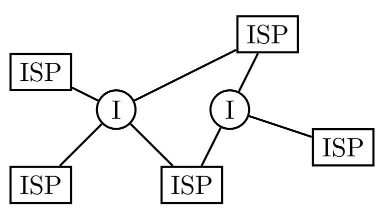

A common approximation (Laffont et al., 2003) we will use is that CPs (web sites) only send traffic and end users only receive traffic. This approximation is justified by the real traffic patterns on the Internet which show that downstream data transmission volume to end users is much bigger than that upstream. This assumption excludes peer to peer relationships from the analysis.

Figure 2: Hierarchical structure of the Internet

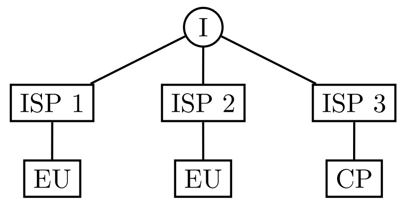

2.1. Internet Service Providers

ISPs provide connectivity to end users and CPs. They interconnect at peering points and the originating ISP pays an access fee a to the terminating ISP. In Figure 2, ISP 3 would pay ISP 2 for delivering data from the content provider to the end user it is connected to. We assume that ISPs have no lack of bandwidth on their backbones and could provide quality assurance to traffic either through excess capacity or network management techniques. Managing capacity on the backbone is within the ISPs' power and there are no interdependencies with other ISPs. Further bandwidth bottlenecks may be present in the peering points and in the access network. We ignore possible problems due to constrained access bandwidth and concentrate on the peering points.

2.2. Points of Interconnection

In Figure 1 and Figure 2 the circles with an "I" represent points of interconnection or peering points where different ISPs interconnect their networks to form the Internet. There are two dominant modes of interconnection: Peering and transit. Peering (Shrimali & Kumar, 2008) is a settlement free agreement to exchange traffic while transit involves payments for exchanged data. Typically peering agreements are used between ISPs of similar size while transit is paid from small ISPs to larger ISPs.

Peering points with peering agreements are among the major bottlenecks of the Internet (Akella et al., 2003). There are several reasons for this. Firstly it always takes both parties in a peering agreement to agree on an extension of a peering point in order to increase its usable capacity (Cremer et al., 2000). Since a capacity extension is costly for both parties, in general the lower of both capacity requirements is realized. See Economides (2002) and Cremer et al. (2000) for a controversial discussion and also Armstrong (1998), Badasyan & Chakrabarti (2008), Cremer et al. (2000) and Foros et al. (2005) for further details on interconnection practices. Ways for CPs to bypass overloaded peerings are multihoming and the use of CDN services.

Transit on the other hand involves a payment from one ISP to the other for the delivery of traffic. With such an agreement a guaranteed bandwidth is bought. The biggest networks (called tier 1 networks) only peer among themselves and charge smaller networks for sending traffic to them. Since small ISPs have to pay for sending traffic to larger networks which is necessary to reach the whole Internet, they optimize their outpayments for transit fees by buying the least amount of bandwidth their users will tolerate. It follows that peerings with peering as well as transit agreements are bandwidth bottlenecks. With transit this is a conscious choice of the buyer, with peering it is a result of non-cooperative behavior.

2.3. Content Providers

Content providers are websites or other service providers that buy connectivity to the Internet from an ISP. CPs are able to multi-home which means they can buy connectivity for one fraction of their traffic from ISP 1 and the rest from ISP 2. Furthermore, they can buy connectivity to the Internet from any ISP anywhere in the world. Therefore, CPs face a market price for ordinary Internet connectivity which is based on perfect competition. This price only includes unprioritized traffic across peering points. Canonical analysis as in Laffont & Tirole (2000) and Laffont et al. (2001, 2003) assumes the following model of payments between network providers:

$$\mathrm{CP} \longrightarrow \mathrm{ISP}_o \xrightarrow{a} \mathrm{ISP}_t \longleftarrow \mathrm{EU} \tag{1}$$

(t=terminating, o=originating, a=access charge). This scheme ignores where the CP gets funding from and emphasizes the analysis of the inter ISP settlement a which has an influence on the prices paid to the ISPs. In contrast, this work focuses on content related charges. Therefore we model the payment flows according to the content delivery value chain:

$$\mathrm{ISP}_t \xleftarrow{a} \mathrm{ISP}_o \xleftarrow{p_w} \mathrm{CP} \xleftarrow{p} \mathrm{EU} \tag{2}$$

Ignoring payments from the end user to the terminating ISP for access to the Internet, payments flow from the end user along the value chain of content delivery to the terminating ISP. Here p is the price paid by the end user for viewing content, p_w is the price paid by the CP to the ISP for reaching the end user. If the ISP that receives p_w cannot terminate the traffic it has to pay an access charge to another ISP that is able to terminate the traffic.

This paper is about two alternatives to this “ordinary” way to deliver data over the Internet. The two variations we will consider are MH and CDNs. With MH, the terminating ISP is directly connected with the CP (Figure 3) while with CDNs a third party mediates between CP and ISP_t (Figure 4). Under MH, payment-flows are:

$$\mathrm{ISP}_t \xleftarrow{p_w} \mathrm{CP} \xleftarrow{p} \mathrm{EU} \tag{3}$$

and the originating ISP is eliminated from the delivery chain. With CDN, the payments are:

$$\mathrm{ISP}_t \xleftarrow{p_w} \mathrm{CDN} \xleftarrow{p_w+m} \mathrm{CP} \xleftarrow{p} \mathrm{EU} \tag{4}$$

The Charge $p_w + m$ levied by the CDN implies that we do not consider the CDN’s pricing decision explicitly but let it add an externally given markup m to its cost for interconnection with the terminating ISP and charge the sum of cost and markup to the CPs.

Figure 3: Hierarchical structure of the Internet with MH

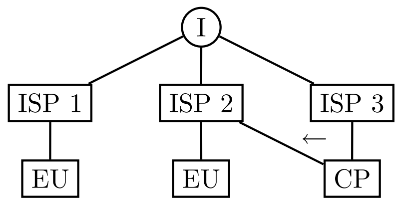

Figure 4: Hierarchical structure of the Internet with CDN

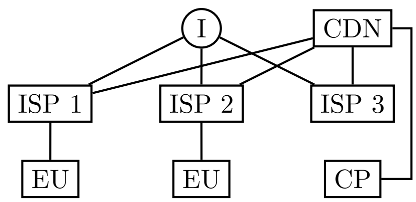

2.4. End Users

Unlike CPs, end users cannot divide their traffic amongst several ISPs and are immobile. They cannot choose their provider globally but need to choose among a small number of local ISPs. In the static model end users are bound to their ISP, providing the ISP with a monopoly over terminating traffic to them.

2.5. Content Delivery Networks

CDNs (Pathan & Buyya, 2007; Vakali & Pallis, 2003) consist of a network of servers that are distributed around the Internet within many ISPs' infrastructures (Figure 4). A CDN takes content from a CP and caches it on those distributed servers which has two effects: First, content is brought closer to the end user without passing through inter ISP peerings thus making its delivery faster. Second, the CDN has a contractual relationship with the ISP where it needs to terminate traffic. The CDN delivers the cached content from the mirror site to the end user. By using the services of a CDN, a CP does not need to multihome with every possible network.

The pricing decisions of a CDN most probably deserve an article of their own since there are several reasonable approaches to model the CDN decision problem. One perspective is that a CDN simply takes its cost for presence at ISPs' sites as given and then optimizes the prices it charges to its customers. From this point of view the CDN's pricing would depend on the competitiveness of the market for its services. On the other hand, a CDN could also be considered a platform that needs to get ISPs (through access to their network, specific contracts and hardware at their computing sites) as well as content providers (as customers hosting content on the CDN servers) on board. Now a two-sided market approach might be feasible since the CDN needs to optimize across two distinct but interdependent customer groups. The present analysis ignores the complexity of that decision through the assumption that the CDN charges w_p plus an additive constant m per unit of content or bandwidth to its customers. For simplicity, we actually ignore m in the formal model below as it has no qualitative effects on our conclusions.

2.6. Quality of Service

Quality of service (QoS) refers to technologies that enable the Internet to guarantee certain bounds on technical parameters of packet transmission such as packet loss, delay and jitter. By tagging each data packet on the internet with a quality label, routers are able to prioritize packets with higher quality requirements (Wang, 2001). The quality differentiation capabilities of the Internet protocol are currently not being used in the public Internet. Speaking in economic terms, traffic differentiation and price discrimination based on the type of data being transported is not practiced on the Internet. Since the Internet cannot assure constant quality levels but there is a demand for improved quality, MH and CDN as means to bypass the main bottlenecks are used by commercial CPs.

3. A Model

In the following we develop a simple model of the market for Internet connectivity as described intuitively above. It allows us to analyze in a rigorous manner how MH and CDN affect ISPs' incentives to price traffic in comparison to the standard situation with peering. We consider two degrees of competition that content providers might face: content competition and content monopoly. These polar cases can be seen as a benchmark for further analysis with intermediate degrees of competition. In total, we thus compare six situations:

1. Content competition without MH or CDN
2. Content competition with MH
3. Content competition with CDN
4. Content monopoly without MH or CDN
5. Content monopoly with MH
6. Content monopoly with CDN

For the analysis we assume that the degree of competition and the type of interconnection are exogenously given. In each situation, we then look at the price building mechanism and see whether the different firms are able to generate positive profits through price discrimination or not.

Suppose there are n markets. In the three situations of content competition all n markets are served by many CPs. In the three situations of content monopoly each of the n markets is served by one CP. For simplicity, assume one terminating ISP (ISP_t) and many originating ISPs (generically denoted by ISP_o). While this is a bit unorthodox, it captures the fact that the end user is locked with the terminating ISP while content providers may freely switch among originating ISPs. (An alternative way is to assume that each market has a unique terminating ISP which would lead to the same results). Moreover, suppose that inverse demand in market i is given by

$$p_i = \alpha_i - \beta_i q_i \tag{5}$$

where q_i is the quantity of bandwidth or content consumed by the end users in that market in a particular period of time and p_i is the price per unit consumed. α_i and β_i are parameters valid for that particular market i. Intuitively, higher prices would discourage Internet services consumption and result in lower bandwidth consumption and vice versa.

Content providers have marginal costs c_i in market i and zero fixed costs. Let the total marginal cost of the traffic in market i by the content provider be $\hat{c}_i$. This includes both its own marginal cost c_i as well as the price per unit of traffic levied by the ISP it is connecting to. Let the marginal cost of ISP$_t$ be c_t and the marginal cost of ISP$_o$ be c_o per unit of traffic.

We are considering a two-stage game. In the first stage, content providers decide on the type of connection. Here there are three choices: without MH/CDN, MH, CDN. While it is likely that multihoming and content delivery networks provide a better quality of service, we abstract from such quality issues and focus purely on the possibility of discrimination that multihoming or content delivery networks make possible. In the standard case without multihoming or content delivery networks, all the content providers connect to ISP$_o$. So the ISP$_t$ who has monopoly power over the end user cannot identify the source of traffic. In case of multihoming, the content providers connect directly to the ISP$_t$. In case of content delivery networks, all content providers connect to the content delivery network, so the ISP$_t$ again cannot identify the specific market the traffic is coming from.

3.1. Content Competition

In a competitive market, CPs set their price for end users equal to their total marginal costs $\hat{c}_i$. Quantities arise according to our demand function given by equation (5). In the second stage of the game we thus have:

$$p_i^* = \hat{c}_i \tag{6}$$

$$q_i^* = \frac{\alpha_i - \hat{c}_i}{\beta_i} \tag{7}$$

In the following we look at the first stage of the game to determine the price setting behavior of the ISPs and thereby also the cost $\hat{c}_i$ the content provider faces in the second stage.

3.1.1. Content Competition without MH or CDN

Consider a standard Internet interconnection situation of content competition (CC) where CPs do not use multihoming or content delivery networks. ISP$_o$ has no market power and thus charges his total marginal cost $p_w = a + c_o$

to CPs, that is, the access charge a he has to pay to ISP_t for terminating the traffic plus his own marginal cost c_o. His profits are zero. ISP_t cannot identify the traffic and thus levies a fee at marginal cost as well:

$$a^{CC} = c_t \tag{8}$$

As a result, the ISP_t's profits are zero. The total marginal cost for a CP can now be written as $\hat{c}_i = p_w + c_i = a^{CC} + c_o + c_i = c_t + c_o + c_i$. CPs charge a competitive price from end users that is equal to this total marginal cost. Thus, the CPs' profits are zero as well. By using the expression we found for $\hat{c}_i$ in equations (6) and (7), the market outcome in a situation of content competition without MH or CDN can finally be summarized as follows:

$$\begin{aligned}
p_i^{CC} &= c_i + c_o + c_t && (9)\\
q_i^{CC} &= \frac{\alpha_i - (c_t + c_o + c_i)}{\beta_i} && (10)\\
\pi_{CP,i}^{CC} &= 0 && (11)\\
\pi_{ISP_o}^{CC} &= 0 && (12)\\
\pi_{ISP_t}^{CC} &= 0 && (13)
\end{aligned}$$

Intuition: This is our benchmark situation. All prices and charges are at set to a competitive level at total marginal costs and all firms earn zero profits in equilibrium. In the context of this paper the central point is that ISP_t has no contractual counterpart from which rents could be extracted.

3.1.2. Content Competition with Multihoming

Now, consider a situation of content competition (CC) where CPs use multihoming (MH). The ISP_t can now identify the traffic and would levy access charges to maximize its profits. Suppose it levies a fee a_i in market i. Then, the CP will charge a competitive price given by $p_i = \hat{c}_i = a_i + c_i$. His profits are zero. Taking demand from equation (5) into account, the total profit of the ISP_t (including all n markets) can be written as:

$$\pi_{ISP_t} = \sum_{i=1}^{n} \left[(a_i - c_t) \underbrace{\left(\frac{\alpha_i - (a_i + c_i)}{\beta_i} \right)}_{q_i} \right] \tag{14}$$

As first order condition for a maximum of this expression we have:

$$\frac{\partial \pi_{ISP_t}}{\partial a_i} = \frac{\alpha_i - 2a_i - c_i + c_t}{\beta_i} \stackrel{!}{=} 0 \tag{15}$$

Solving for a_i gives us the optimal access charge per unit of traffic in market i:

$$a_i^{CC,MH} = \frac{\alpha_i - c_i + c_t}{2} \tag{16}$$

By using $\hat{c}_i = a_i + c_i$ combined with (16) in equations (6) and (7) we find the equilibrium prices and quantities; using (16) in (14) gives us the ISP$_t$'s profit in equilibrium. The market outcome in a situation of content competition with MH can thus be summarized as follows:

$$p_i^{CC,MH} = \frac{\alpha_i + (c_t + c_i)}{2} \tag{17}$$

$$q_i^{CC,MH} = \frac{\alpha_i - (c_t + c_i)}{2\beta_i} \tag{18}$$

$$\pi_{CP,i}^{CC,MH} = 0 \tag{19}$$

$$\pi_{ISP_t}^{CC,MH} = \sum_{i=1}^{n} \left[\frac{(\alpha_i - (c_t + c_i))^2}{4\beta_i} \right] \tag{20}$$

Intuition: By comparing the profit of the ISP$_t$ with the situation from above we see that multihoming and the possibility to identify traffic allow the ISP$_t$ to convert each competitive market into a perfect monopoly and to extract the monopoly profit while the CPs still make zero profits.

3.1.3. Content Competition with Content Delivery Networks

Now, consider a situation of content competition (CC) where CPs use content delivery networks (CDN) for interconnection. In this case, the traffic comes under the filter of the CDN so the ISP$_t$ is forced to levy a uniform access charge a across all markets (in technical terms, a has no index i anymore). A CP in market i passes this access charge on to the end user so the price charged is $p_i = \hat{c}_i = a + c_i$. His profits are zero. Taking demand from equation (5) into account, the total profit of the ISP$_t$ (including all n markets) can be written as:

$$\pi_{ISP_t} = (a - c_t) \sum_{i=1}^{n} \left[\underbrace{\left(\frac{\alpha_i - (a + c_i)}{\beta_i} \right)}_{q_i} \right] \tag{21}$$

As first order condition for a maximum of this expression we have:

$$\frac{\partial \pi_{ISP_t}}{\partial a} = \sum_{i=1}^{n} \left(\frac{\alpha_i - 2a - c_i + c_t}{\beta_i} \right) \stackrel{!}{=} 0 \tag{22}$$

Solving for a yields the optimal access charge per unit of traffic:

$$a^{CC,CDN} = \frac{\sum_{i=1}^{n} \left(\frac{\alpha_i - c_i + c_t}{\beta_i} \right)}{2 \sum_{i=1}^{n} \left(\frac{1}{\beta_i} \right)} \tag{23}$$

Note again that this access fee, in contrast to the multihoming case, is uniform across all markets. By using $\hat{c}_i = a + c_i$ combined with (23) in equations (6) and (7) we find the equilibrium prices and quantities; using (23) in (21) gives us the ISP$_t$'s profit in equilibrium. The market outcome in a situation of content competition with CDN can thus be summarized as follows:

$$\begin{aligned}
p_i^{CC,CDN} &= a^{CC,CDN} + c_i && (24)\\
q_i^{CC,CDN} &= \frac{\alpha_i - p_i^{CC,CDN}}{\beta_i} && (25)\\
\pi_{CP,i}^{CC,CDN} &= 0 && (26)\\
\pi_{ISP_t}^{CC,CDN} &= \left(a^{CC,CDN} - c_t\right) \sum_{i=1}^{n} \left[q_i^{CC,CDN}\right] && (27)
\end{aligned}$$

Intuition: In this situation the ISP$_t$ has an intermediate degree of discriminatory power compared to the two previous situations. The ISP$_t$ can differentiate prices among different CDNs but it cannot discriminate against every single CP.

3.2. Content Monopoly

Having analyzed the case of content competition above, we now look at three situations where CPs are monopolists in each market. Consider a content provider who is a monopolist in a local market i. Irrespective of the type of connection he chooses, his profit in the second stage can be written as

$$\pi_{CP,i} = (p_i - \hat{c}_i) \underbrace{\left(\frac{\alpha_i - p_i}{\beta_i}\right)}_{q_i} \tag{28}$$

by taking market demand from equation (5) into account. The first order condition for a maximum of (28) is given by:

$$\frac{\partial \pi_{CP,i}}{\partial p_i} = \frac{\alpha_i - 2p_i + \hat{c}_i}{\beta_i} \overset{!}{=} 0 \tag{29}$$

Rewriting this expression leads us to the price p_i^{**} the content provider charges from end users in equilibrium, the quantity q_i^{**} that is supplied, as well as the profit of the content provider $\pi_{CP,i}$ in equilibrium:

$$\begin{aligned}
p_i^{**} &= \frac{\alpha_i + \hat{c}_i}{2} && (30)\\
q_i^{**} &= \frac{\alpha_i - \hat{c}_i}{2\beta_i} && (31)\\
\pi_{CP,i} &= \frac{(\alpha_i - \hat{c}_i)^2}{4\beta_i} && (32)
\end{aligned}$$

In the following we look at the first stage of the game to determine the price setting behavior of the ISPs and thereby also the cost $\hat{c}_i$ the content provider faces in the second stage.

3.2.1. Content Monopoly without MH or CDN

Consider a standard Internet interconnection situation of content monopoly (CM) where CPs do not use multihoming or content delivery networks. ISP$_o$ has no market power and thus charges his total marginal cost $p_w = a + c_o$ from CPs, that is, the access charge a he has to pay to ISP$_t$ for terminating the traffic plus his own marginal cost c_o. His profits are zero. ISP$_t$ cannot identify the traffic and thus levies a fee at marginal cost as well:

$$a^{CM} = c_t \tag{33}$$

As a result, the ISP$_t$'s profits are zero. The total marginal cost for a CP can now be written as $\hat{c}_i = p_w + c_i = a^{CC} + c_o + c_i = c_t + c_o + c_i$. By using this expression in equations (30), (31) and (32), the market outcome in a situation of content monopoly without MH or CDN can finally be summarized as follows:

$$p_i^{CM} = \frac{\alpha_i + (c_t + c_o + c_i)}{2} \tag{34}$$

$$q_i^{CM} = \frac{\alpha_i - (c_t + c_o + c_i)}{2\beta_i} \tag{35}$$

$$\pi_{CP,i}^{CM} = \frac{(\alpha_i - (c_t + c_o + c_i))^2}{4\beta_i} \tag{36}$$

$$\pi_{ISP_o}^{CM} = 0 \tag{37}$$

$$\pi_{ISP_t}^{CM} = 0 \tag{38}$$

Intuition: In contrast to the three situations of content competition, in a standard Internet interconnection situation of content monopoly CPs are able to earn positive profits. These profits cannot be extracted by any of the downstream parties: ISP$_o$ faces a competitive environment and ISP$_t$ has no means of discrimination against the source of traffic.

3.2.2. Content Monopoly with Multihoming

Now, consider a situation of content monopoly (CM) where CPs use multihoming (MH). The CP in market i connects directly with ISP$_t$ while there is no ISP$_o$ in the market. The ISP$_t$ can identify the CP where the traffic is coming from. Assuming that the ISP$_t$ levies a two-part tariff, it will set the per unit traffic rate at:

$$a_i^{CM,MH} = c_t \tag{39}$$

Hence, CPs face total marginal costs of $\hat{c}_i = c_t + c_i$ and without any other fees, following (32), their profits would be $(\alpha_i - (c_t + c_i))^2/(4\beta_i)$. ISP$_t$, however, will set a fixed lump-sum fee at exactly this value to maximize his profits:

$$A_i^{CM,MH} = \frac{(\alpha_i - (c_t + c_i))^2}{4\beta_i} \tag{40}$$

Thus, the CPs' actual profits turn out to be zero. The ISP$_t$'s profit, by contrast, is the sum of all n lump-sum fees. By using $\hat{c}_i = c_t + c_i$ in equations (30) and

(31) we find the equilibrium prices and quantities. In sum, the market outcome in a situation of content monopoly with multihoming is as follows:

$$p_i^{CM,MH} = \frac{\alpha_i + (c_t + c_i)}{2} \tag{41}$$

$$q_i^{CM,MH} = \frac{\alpha_i - (c_t + c_i)}{2\beta_i} \tag{42}$$

$$\pi_{CP,i}^{CM,MH} = 0 \tag{43}$$

$$\pi_{ISP_t}^{CM,MH} = \sum_{i=1}^{n} \left[\frac{(\alpha_i - (c_t + c_i))^2}{4\beta_i} \right] \tag{44}$$

Intuition: End users still pay the monopoly price to the CP (now excluding the ISP_os cost). Multihoming, however, allows the ISP_t to extract all the profits from the CP by levying a two-part tariff. This can be a problem because monopoly profits could be the reward for innovation and if these profits are taken away from the CP, his ambitions to innovate might be suppressed.

3.2.3. Content Monopoly with Content Delivery Networks

Now, consider a situation of content monopoly (CM) where CPs use content delivery networks (CDN) for interconnection. The traffic is coming through the filter of the CDN so the ISP_t cannot identify individual traffic. It has to levy uniform fees for all CDN traffic. We cannot assume that levying a two-part tariff is possible because that itself requires identification of the traffic source. Let a be the fee per unit of traffic. Now, assuming that $a^{CM,CDN} < \alpha_i - c_i$ for all n markets, we get the following prices, quantities and profits:

$$p_i^{CM,CDN} = \frac{\alpha_i + (a + c_i)}{2} \tag{45}$$

$$q_i^{CM,CDN} = \frac{\alpha_i - (a + c_i)}{2\beta_i} \tag{46}$$

$$\pi_{CP,i}^{CM,CDN} = \frac{(\alpha_i - (a + c_i))^2}{4\beta_i} \tag{47}$$

$$\pi_{ISP_t}^{CM,CDN} = (a - c_t) \sum_{i=1}^{n} \left[\frac{\alpha_i - (a + c_i)}{2\beta_i} \right] \tag{48}$$

Note that these expressions still depend on a and only hold if all CPs stay in their respective market. For an equilibrium access charge of $a^{CM,CDN} > \alpha_i - c_i$, the quantity demanded in market i as given by equation (46) would become negative. As a result, CP_i would exit the market. At the same time, the ISP's total profit as given by equation (48) would decrease: profits could only be generated in a subset of markets that is smaller than n. Thus, to decide on the optimal value of the access charge, the ISP does not only have to take the direct positive effect of increasing a into account but also an indirect negative effect that comes from a smaller number of markets to serve.

Table 1: Profit generation along the content delivery value chain

	Without MH/CDN	MH	CDN
Content Competition (CC)	ISP_t charges competitive price (her total marginal costs) from ISP_o as does ISP_o from CP and CP from end users $\pi^{CC}_{ISP_t} = 0$ $\pi^{CC}_{ISP_o} = 0$ $\pi^{CC}_{CP,i} = 0$	ISP_t charges monopolistic access fee from CP in each individual market, CP charges competitive price (her total marginal costs) from end users $\pi^{CC,MH}_{ISP_t} > 0$ $\pi^{CC,MH}_{CP,i} = 0$	ISP_t charges uniform monopolistic access fee through the CDN from CP, CP charges competitive price (her total marginal costs) from end users $\pi^{CC,CDN}_{ISP_t} > 0$ $\pi^{CC,CDN}_{CP,i} = 0$
Content Monopoly (CM)	ISP_t charges competitive price from ISP_o and ISP_o from CP, CP charges monopoly price from end users $\pi^{CM}_{ISP_t} = 0$ $\pi^{CM}_{ISP_o} = 0$ $\pi^{CM}_{CP,i} > 0$	ISP_t charges two-part tariff from CP to extract all profits the CP gets by charging monopoly price from end users $\pi^{CM,MH}_{ISP_t} > 0$ $\pi^{CM,MH}_{CP,i} = 0$	ISP_t charges uniform monopolistic access fee through the CDN from CP, CP charges monopoly price from end users $\pi^{CM,CDN}_{ISP_t} > 0$ $\sum_{i=1}^{n} \pi^{CM,CDN}_{CP,i} > 0$

We could solve the ISP's optimization problem and present a full equilibrium outcome by imposing further assumptions on the parameters in the model. For the following discussion this is of little value, however. It is rather important to see the general characteristics of this situation: On the one hand, for a given a, there may be CPs that do not face any demand. These CPs exit the market and thus earn zero profits. On the other hand, there are CPs that stay in the market and earn positive profits thanks to their positions as monopolists.

Intuition: In sum and on average CPs still earn positive profits in this situation. Also the ISP earns positive profits. For those markets that do not shut down the literature refers to such a situation as "double marginalization": there are two independent firms, upstream and downstream, CP_i and ISP, that both have market power and price at a markup over their cost.

3.3. Summary and Implications for Future Innovations

Table 1 summarizes the profit generation along the content delivery value chain. We presented six situations. In the first three of them we assumed that there is competition among content providers. In a standard Internet interconnection situation with content competition, ISPs as well as CPs make zero profits. In a situation of content competition where CPs use MH, the ISP is able to charge monopolistic access fees from CPs in each individual market, leading to positive profits for the ISP. Facing competition, CPs cannot charge more than their total marginal cost from end users and thus make zero profits as without MH. In a situation where CPs use CDNs, the ISP can differentiate prices among different CDNs but it cannot discriminate against every single CP.

As a result, the ISP only has an intermediate degree of discriminatory power but still makes positive profits. CPs make zero profits.

In three further situations we assumed that content providers are monopolists in their market. In a standard Internet interconnection situation with content monopoly, ISPs make zero profits while CPs use their monopoly power to earn positive profits. In a situation where CPs use MH, the ISP is able to charge a two-part tariff from CPs to extract all monopoly profits that the CPs get from end users. In contrast to the ISP, CPs thus make zero profits in the end . In a situation of content monopoly where CPs use CDNs, the ISP charges a uniform monopolistic access fee from CPs trough the CDN. CPs, in turn, charge monopoly prices from end users. We thus face a "double marginalization" situation where both, ISP and CPs (in sum), make positive profits.

In a nutshell, multihoming gives the ISP monopoly power to exploit her access monopoly to end users. This holds in the case of content competition as well as in the case of content monopoly. Content delivery networks give the ISP some monopoly power as well. However, since ISPs cannot discriminate against every single CP it is less pronounced than with multihoming.

Note that these results may have strong implications for future innovations. Under both assumptions for the degree of competition among content providers, full competition and monopoly, MH and CDNs allow the ISP to earn positive profits while without these technologies her profits are zero. These positive profits may well be used to finance future innovations which may have not been possible without MH or CDNs. For CPs, however, we get a different picture. Under the assumption of content competition nothing changes for CPs with the introduction of MH or CDNs. Their profits are zero in all three situations. Under the assumption of content monopoly, by contrast, all profits of the CPs are shifted away to the ISP with the introduction of MH. With the introduction of CDNs, individual CPs' profits may decrease. In sum and on average, however, CPs' profits remain positive.

In a nutshell, our model suggests that the introduction of MH or CDNs increases the potential to finance future innovations for the ISP. The CPs' potential to innovate, by contrast, remains unchanged or decreases.

3.4. Limitations of the Model

To put the results of our model into perspective it is important to be aware of the limitations of their applicability. First, the assumption that end users are perfectly locked in with their ISP can be challenged. Consumers that are able to switch their ISP will probably not tolerate monopoly prices for content in the long run if the total price for content differs between ISPs. Furthermore, they might not tolerate low quality access to certain content and thus force the ISP to invest in its standard peerings.

Second, a large part of the content business is financed by advertisements. The presented analysis relies on the exchange of money between end user and CP. Since this is not the case in ad-financed business models, the presented analysis cannot be applied to websites that base their business model on selling banner space.

Third, we only analyze two polar cases concerning the degree of competition that CPs might face. Usually, markets are neither pure monopolies nor perfectly competitive. Thus, realistic outcomes will fall somewhere in between the analyzed situations. The assumption of content monopoly may be in part justified by considering temporal monopolies gained through innovation and patents. However, more realistic modeling assumptions should take into account the role of substitutes.

Finally, it is only through the use of MH and CDNs that ISPs are put in a position to exploit their access monopoly and to create monopolies from otherwise competitive markets. The self selective nature of this phenomenon –CPs have to actively choose to give the ISP that power– makes it likely that CPs still get positive payoffs from doing so. Therefore, there are probably other effects at play that create more balanced outcomes than those of the pure MH and CDN situations above and that have not been captured by the model.

4. Conclusion and Further Research

The central insight of this work is that price discrimination is possible in today's Internet. The lever for ISPs to practice this price discrimination is not differentiation of data packets in the style of DiffServ (Wang, 2001) or a congestion based pricing mechanism like Odlyzko's Paris Metro Pricing (Odlyzko, 1999). Much rather differentiation is achieved indirectly through offering enhanced modes of interconnection. The reason for ISPs collaborating with CDNs can thus in part be attributed to the revenue potential ISPs see in it. Furthermore the claim that all data that is transported on the Internet is equal (Wu, 2003) must be rejected after considering the above analysis. The possibility to offer differentiated quality levels of data transport does exist on the Internet and it has not led to a breakdown of connectivity or other adverse developments. Much rather different quality and price levels are evolving on the Internet. Differentiated product offerings are generally thought to be welfare enhancing. However, with monopolistic firms one also has to watch their incentives to migrate customers to the more profitable service classes. The possibility that ISPs degrade standard peering quality to move customers to MH or CDN clearly exists. As mentioned above, the assumptions about the competitiveness of markets made in this work are quite strong. To mend these limitations further research could firstly introduce more sophisticated modes of competition. Horizontal product differentiation in a Hotelling framework (Hotelling, 1929) could offer a starting point for such an analysis. Limiting the termination monopoly power of ISPs could involve the introduction of switching costs (Klemperer, 1995). A further limitation of this work is that it is only applicable to paid content. With today's prevalence of advertisement financed business models, the exploration of two sided market (Armstrong, 2006; Rochet & Tirole, 2006) models should yield further insights into the matter of pricing Internet traffic. In a two sided market it would even be possible that end users get subsidized access to the Internet from their ISP since they are valuable assets. The terminating ISP needs to get end users on its network in order to be attractive to content providers.

Future research should also focus on the problem of the still existing lack of guaranteeable QoS. What can be learned from the success of CDNs and MH and how could these technologies be combined with other technology to further improve QoS on the Internet? Furthermore, the question about the economic efficiency of CDNs and MH must be answered. Under which circumstances are these technologies efficient? Is a global QoS regime - based for example on DiffServ - desirable in the light of the availability of these methods? Can CDNs and MH fully replace inter carrier agreements on quality parameters of traffic? Assuming the business model of providing data transport on the Internet changes towards more CDNs and MH, which effects will this have on the Internet as a whole beyond growing routing tables?

We have formally shown that incentives to degrade standard peering and transit do exist. Further research should refine this result to go beyond the politicized net neutrality debate.

Acknowledgments

The authors would like to thank two anonymous referees for their invaluable comments. The usual disclaimer applies.

References

Akella, A., Srinivasan, S., & Shaikh, A. (2003). An empirical evaluation of wide-area internet bottlenecks. In *Proceedings of the 3rd ACM SIGCOMM Conference on Internet Measurement* (pp. 101–114). New York: ACM.

Armstrong, M. (1998). Network interconnection in telecommunications. *Economic Journal*, *108*, 545–564.

Armstrong, M. (2006). Competition in two-sided markets. *RAND Journal of Economics*, *37*, 668–691.

Badasyan, N., & Chakrabarti, S. (2008). A simple game-theoretic analysis of peering and transit contracting among internet service providers. *Telecommunications Policy*, *32*, 4–54.

Bu, T., Gao, L., & Towsley, D. (2004). On characterizing bgp routing table growth. *Computer Networks*, *45*, 45–54.

Buehler, S., & Schmutzler, A. (2006). On the role of access charges under network competition. In R. Dewenter, & J. Haucap (Eds.), *Access Pricing: Theory and Practice* (pp. 121–147). Amsterdam: Elsevier.

Cremer, J., Rey, P., & Tirole, J. (2000). Connectivity in the commercial internet. *Journal of Industrial Economics*, *48*, 433–472.

Crowcroft, J. (2007). Net neutrality: the technical side of the debate: a white paper. *ACM SIGCOMM Computer Communication Review*, *37*, 49–56.

Economides, N. (2002). The economics of the internet backbone. In M. E. Cave, S. K. Majumdar, & I. Vogelsang (Eds.), *Handbook of Telecommunication Economics* Vol. 2. Amsterdam: North Holland.

Foros, O., Kind, H. J., & Sand, J. Y. (2005). Do internet incumbents choose low interconnection quality? *Information Economics and Policy*, *17*, 149–164.

Gans, J. S. (2006). Access pricing and infrastructure investment. In R. Dewenter, & J. Haucap (Eds.), *Access pricing: Theory and practice* (pp. 41–63). Amsterdam: Elsevier.

Hotelling, H. (1929). Stability in competition. *Economic Journal*, *39*, 41–57.

Klemperer, P. (1995). Competition when consumers have switching costs: An overview with applications to industrial organization, macroeconomics, and international trade. *Review of Economics and Statistics*, *62*, 515–539.

Laffont, J.-J., Marcus, S., Rey, P., & Tirole, J. (2001). Internet peering. *American Economic Review, Papers and Proceedings*, *91*, 287–291.

Laffont, J.-J., Marcus, S., Rey, P., & Tirole, J. (2003). Internet interconnection and the off-net-cost pricing principle. *RAND Journal of Economics*, *34*, 370–390.

Laffont, J.-J., Rey, P., & Tirole, J. (1998). Network competition: I. overview and nondiscriminatory pricing. *RAND Journal of Economics*, *29*, 1–37.

Laffont, J.-J., & Tirole, J. (2000). *Competition in Telecommunications*. Cambridge, MA: MIT Press.

MacKie-Mason, J. K., & Varian, H. R. (1995). Pricing the internet. In B. Kahin, & J. Keller (Eds.), *Public access to the internet*. London: Prentice Hall.

Odlyzko, A. (1999). Paris metro pricing for the internet. In *Proceedings of the 1st ACM Conference on Electronic Commerce* (pp. 140–147). New York: ACM.

Pathan, M., & Buyya, R. (2007). A taxonomy and survey of content delivery networks. Technical Report GRIDS-TR-2007-4, University of Melbourne.

Rey, P., & Tirole, J. (2007). A primer on foreclosure. In M. Armstrong, & R. H. Porter (Eds.), *Handbook of Industrial Organization* Vol. 3. Amsterdam: North-Holland.

Rochet, J.-C., & Tirole, J. (2006). Two-sided markets: a progress report. *RAND Journal of Economics*, *37*, 645–667.

Shakkottai, S., & Srikant, R. (2006). Economics of network pricing with multiple isps. *IEEE/ACM Transactions on Networking*, *14*, 1233–1245.

Shrimali, G. (2008). Surplus extraction by network providers: Implications for net neutrality and innovation. *Telecommunications Policy*, *32*, 545–558.

Shrimali, G., & Kumar, S. (2008). Bill-and-keep peering. *Telecommunications Policy*, *32*, 19–32.

Sidak, J. G. (2006). A consumer-welfare approach to network neutrality regulation of the internet. *Journal of Competition Law and Economics*, *2*, 349.

Soldatos, J., Vayas, E., & Kormentzas, G. (2005). On the building blocks of quality of service in heterogeneous ip networks. *IEEE Communications Surveys and Tutorials*, *7*, 69–88.

Tirole, J. (1988). *The Theory of Industrial Organization*. Cambridge, MA: MIT Press.

Uludag, S., Lui, K.-S., Nahrstedt, K., & Brewster, G. (2007). Analysis of technology aggregation techniques for qos routing. *ACM Computing Surveys*, *39*, 7.

Vakali, A., & Pallis, G. (2003). Content delivery networks: status and trends. *IEEE Internet Computing*, *7*, 68–74.

Wang, Z. (2001). *Internet QoS: Architectures and mechanisms for quality of service*. San Francisco: Morgan Kaufmann Publishers.

Wu, T. (2003). Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, *2*, 141.