

Economic Effects of Multi-Homing and Content Delivery Networks on the Internet

Abstract

Peering points between different network providers are among the bottlenecks of the Internet. Two technical solutions to bypass them are content delivery networks (CDN) and multi-homing (MH). While quality of data delivery can thus be improved, there is no research that analyses the economic effects of those technologies on the Internet. This paper argues that CDNs and MH reduce the commodity nature of data being transported on the Internet and put terminating Internet service providers (ISPs) in a position to engage in price discrimination. By formalizing CDNs and MH as facilitators of price discrimination we show that - in a static model with locked in end users and paid content - ISPs will exploit their market power and generate profits. As a consequence ISPs have incentives to degrade ordinary peering quality and shift demand to higher quality methods of access to EUs such as CDN and MH.

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. 1998a; Laffont et al. 1998b), i.e. the price one provider pays to the other for termination of traffic with an end user (EU). This paper departs from this “classical” view on communications pricing by considering content oriented pricing of termination instead of purely data oriented inter carrier settlement. To the authors’ knowledge this issue is not covered in the extant literature. We show how an ISP with access to EUs can discriminate against content providers and charge higher prices for termination. The discussion is related to and uses results from research on two sided markets (Armstrong 2006; Rochet 2006), vertical integration (Rey 2006; Tirole 1988), telecommunications pricing (Laffont 2003; MacKie-Mason 1995; Shakkottai 2006), Net Neutrality (Crowcroft 2007; Sidak 2006; Wu 2003) and quality of service (QoS) (Soldatos 2005; Wang 2001).

The existing literature has ignored the possibility that content providers (CPs) and terminating internet service providers (ISPs) directly interconnect. In contrast consider the following two situations: Firstly, it is commonplace that content providers (CPs) directly buy transit from terminating ISPs, thus effectively paying them for preferential access to end users. This practice is called multi homing (MH) and is responsible for exponentially growing routing tables (Bu et al. Mai 15, 2004). Secondly, content delivery networks (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 EUs without the need to traverse several ISPs’ networks (Vakali et al. 2003).

Both technologies provide viable means to improve the speed and reliability of data transport from a CP’s website to customers (i.e. EUs). 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 Internet bottlenecks (Akella et al. 2003). Both technologies serve as ways to bypass these peerings and gain more direct access to the EUs; thus increasing the probability of timely delivery of data to the end user.

The paper is structured as follows: First we explain the relevant entities of the current Internet that we need for a formal model. Then we present a formalized treatment of four scenarios that show how CDN and MH affect ISPs incentives to price traffic in a static model with EUs that cannot switch their provider. Lastly we discuss consequences of our model and sketch out an agenda for further research.

2. STATUS QUO: THE MARKET FOR INTERNET CONNECTIVITY

In this Section we present and discuss the entities and concepts necessary to model the Internet.

2.1 Overview

Figure 1 and Figure 2 show in an idealized manner the structure of the Internet (Shakkottai 2006; Uludag 2007). Figure 1 focuses on the interconnection aspect. Several ISPs interconnect with each other. Figure 2 focuses on the hierarchical structure of the Internet. Data first flows up the hierarchy from the CP to its ISP and across a peering point back down via an ISP to the EU.

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




Figure 1: Simplified model of the Internet.




Figure 2: A tiered view of the Internet without MH or CDN.

2.2 Internet Service Providers

ISPs provide connectivity end users and content providers. ISPs interconnect at peering points and the originating ISPs pay the access fee p_w to terminating ISPs. 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. If they have capacity problems on the backbone it is within the ISPs' decision to extend capacity and there are no interdependencies with other ISPs. 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.

2.3 Points of Interconnection

In Figure 1 and Figure 2 the circles with arrows represent a point of interconnection or peering point where different ISPs interconnect their networks to form the Internet. There are two dominant modes of interconnection: Peering and transit. Peering 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. Furthermore, peering agreements commonly do not involve any quality agreements while transit is sold with guaranteed service levels.

Peering points with peering agreements are among the major bottlenecks of the Internet (Akella, Seshan, and Shaikh 2003). There are several reasons for this. Firstly it always takes both parties in a peering agreement to agree to an extension of a peering point in order to increase its usable capacity. Since the telecommunications firms who own the Internet infrastructure do not tend to be overly cooperative among themselves these peering points represent excellent opportunities for power play or generally uncooperative behavior. See (Cave 2002, vol. 2 chap. 9) and (Cremer 2000) for a controversial discussion and also (Armstrong 1998; Badasyan 2008; Cremer 2000; Foros 2005) for further details on interconnection practices. Ways for CPs to bypass overloaded (and therefore high latency and packet dropping) peerings are multi-homing 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 (so 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 out-payments 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.

2.4 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 IPS 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 un-prioritized traffic transportation across peering points. Canonical analysis(Laffont 2000; Laffont, Rey, and Tirole 1998a; Laffont, Rey, and Tirole 1998b; Laffont 2003) usually assumes the following model of Internet payments:

$$CP \rightarrow ISP_o \leftrightarrow_a ISP_t \leftarrow EU .$$

(t=terminating, o=originating), ignoring where the CP gets funding from and emphasizing the analysis of the inter ISP settlement a that has an influence on the prices paid to the ISPs.

Focusing on payments for content viewing, we model the payments flows according to the following scheme:

$$ISP_t \leftarrow_a ISP_o \leftarrow_{p_w} CP \leftarrow_p EU .$$

Ignoring payments from the EU to the terminating ISP for access to the Internet, payments flow from the EU along the value chain of content delivery to the terminating ISP. Here p is

the price paid by the EU for viewing some content, p_w is the price paid from the CP to the ISP for reaching the EU. If the ISP receiving p_w cannot terminate the traffic it has to pay an access charge to another ISP able to terminate the traffic.

When a CDN is involved in content delivery, the CP has to pay the cost of the CDN, too. Modifying the way CPs get access to the EUs and modifying payments accordingly will be the key part of this paper. The two variations we will consider are: multi homing and CDN delivery. With MH, the terminating ISP is directly connected with the CP, while with CDNs a neutral third party mediates between CP and ISP_t . Under MH payment flows are

$$ISP_t \leftarrow_{p_w} CP \leftarrow_p EU$$

and the originating ISP is cut out of the equation. With CDN delivery, the payments are:

$$ISP_t \leftarrow_{p_w} CDN \leftarrow_{p_w + c_{cdn}} CP \leftarrow_p EU.$$

We model CDNs as fully competitive entities only passing on costs.

2.5 End Users

Unlike CPs, EUs 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 EUs are bound to their ISP; providing the ISP with a monopoly over terminating traffic to them.

2.6 Multi Homing

Multi homing (shown in Figure 3) is the common practice that CPs interconnect with several different ISPs. There are several instances of big content providers that are present in public peering points in order to gain direct access to ISPs with many EUs. Google for example is not just with one ISP in the USA but has its own transatlantic fiber capacity and is present in (DE-CIX 2008) where it interconnects with several European ISPs.

2.7 Content Delivery Networks

CDNs 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: Firstly content is brought closer to the EU without passing through inter ISP peerings thus making its delivery faster. Secondly the CDN buys transit from the ISP where it needs to terminate traffic and thus gets a guaranteed bandwidth it can extend as needed. The CDN then delivers the cached content from the mirror site to the EU. By using the services of a CDN a CP does not need to multi-home with every possible network. We assume that CDNs do not make any profit and charge efficient prices.




Figure 3: A tiered view of the Internet with MH.




Figure 4: The Internet with CDNs

2.8 Quality of Service

Quality of service (QoS) refers to technologies that enable the Internet to guarantee certain bounds on technical parameters such as packet loss, delay and jitter of packet transmission. By tagging each data packet on the internet with a quality label, routers are able to prioritize packets with higher quality requirements. The quality differentiation capabilities of the Internet protocol are currently not being used in the public Internet and all types of traffic are treated equally. 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, methods to bypass the main bottlenecks are used by commercial CPs. Two relevant methods are multi-homing and CDN. An important aspect of both technologies is that they “decommoditize” Internet traffic. With multi-homing, the source of traffic is known to the terminating ISP and with CDN it is possible to infer the increased quality requirements from the fact that a CP pays a CDN. Thus terminating ISPs are put in a position discriminate against traffic sources.

3. A Static Model of Multi Homing and CDN

In the following we will analyze four combinations of market form faced by the CP and strategic behavior of the ISPs that result from the use of MH or CDN. The columns of Table 1 represent the possible actions a terminating ISP with can take: With MH it can perfectly discriminate against the known source of traffic and extract profits; with CDNs it can discriminate against all traffic that requires higher quality and has demonstrated this through the use of a CDN (self selection). The rows represent the market situation the CP is facing. Under competition many CPs are serving a market; with a monopoly, only one CP serves a specific market with one product. The cells of the table show what the ISP will do to extract the maximum revenue from the CP for terminating traffic with EUs. The analysis as well as the conclusions apply to business models for paid content on the Internet.

ISP CP	Perfect/1st degree price discrimination (MH)	Class of Service/2 nd degree price discrimination (CDN)
Content Competition	1: ISP sets access price until monopoly price for content is reached	2: ISP sets monopoly price for termination, inefficient since ISP sets average price
Content Monopoly	3: ISP extracts all revenue from CP through two part tariff	4: ISP sets monopoly access price but cannot capture all rents from CPs (dbl. marginalization)

Table 1: The four extreme combinations of market power and ability to discriminate

3.1 Scenario 0: No Discrimination by ISP Possible

The situation when the ISP has no way of differentiating traffic and therefore cannot engage in price discrimination is not shown in Table 1. It corresponds to the “normal” way traffic is exchanged on the Internet. Traffic transport is a commodity and is priced at cost. The price for terminating traffic p_w is equal to the ISP’s cost c . The charge p_w which is levied by the terminating ISP feeds through to the originating CP as a rise in connection costs to be included in the charge to the EU. The whole structure is as efficient, as the market between CP and EUs permits. This treatment ignores the complexity of determining access charges and their effect on prices. For a detailed discussion see (Laffont 2003).

3.2 Scenario 1: Perfect Discrimination by ISP and Content Competition

Assume that the terminating ISP can perfectly discriminate against CPs while the CPs face competitive markets. With perfect competition, only efficient firms are active in each independent market and offer the same price. Therefore firms in one market are identical as far as the ISP is concerned. It can perfectly discriminate against each market. By setting its price $p_{w,i}$ charged to the CPs of segment i equal to the monopoly price the ISP can extract monopoly profits while leaving no profits to the CPs.

Each CP sets the competitive price p_i in segment i equal to its marginal costs. For simplicity we assume that marginal costs of production of the CPs are zero. Therefore $p_i = p_{w,i}$. The ISP with marginal costs c_i per market segment determines $p_{w,i}$ by solving the following problem:

$$\max_{p_{w,i}} [(p_{w,i} - c_i)D_i(p_{w,i})]. \quad (1)$$

For example suppose that the demand function in one of the market segments i is given by $D_i(p_{w,i}) = 1 - p_{w,i}$. Then the target function of the ISP is:

$$\max_{p_{w,i}} [(p_{w,i} - c_i)(1 - p_{w,i})]. \quad (2)$$

Solving this gives for price and output quantity

$$p_{w,i} = \frac{1 + c_i}{2}, \quad (3)$$

$$D_i(p_{w,i}) = q_i = \frac{1 - c_i}{2} \text{ and} \quad (4)$$

$$\Pi = p_{w,i}q_i - c_i = \left(\frac{1 - c_i}{2}\right)^2. \quad (5)$$

Interpretation: The ISP converts each competitive market segment ‘ i ’ into a perfect monopoly and extracts the maximum profit. The output quantity is reduced but the CPs make no profit. This result would not hold with users being able to switch between ISPs. Switching users without switching costs would restore competitiveness of the market.

3.3 Scenario 2: Market Segmentation by ISP and Content Competition

This situation is similar to scenario one only with a reduction in the ability to discriminate against CPs. With a CDN as mediator only CDN (i.e. quality sensitive) and non-CDN traffic can be distinguished. CDNs buy transit from the ISP who can thus differentiate between quality sensitive traffic from CDNs and ordinary traffic through normal peerings with other ISPs.

The price taken by the competitive CPs in this scenario depends on the price p_w set by the ISP which acts like a monopolist. The ISP faces an aggregate demand for quality traffic. Assuming that the CDN segment is competitive, the ISP can set monopoly prices for CDN traffic which raises marginal costs of the CDN providers and thus costs for CPs. In the following treatment we assume the CDN costs to be fully absorbed into the marginal cost of the CP. For simplicity both are assumed to be constant and zero. Again the CPs are competitive in each market segment. In each of n market segment the CPs have the target function

$$\max_{p_i} [(p_i - p_w)D_i(p_i)] \quad (6)$$

where p_i is the price charged to the EUs and p_w is the uniform price for CDN traffic paid to the ISP. Different from scenario one, now the ISP cannot set different $p_{w,i}$ per market segment but just one averaged price p_w for CDN traffic. Under perfect competition price equals marginal costs and the CPs set prices

$$p_i = p_w, \forall i. \quad (7)$$

Interpretation: It is obvious from equation 6 that the ISP influences the pricing decisions of the CP's through its price p_w . However, since the ISP can only set a uniform price for all CDN traffic, it has to optimize across an aggregated demand function $D = \sum D_i$. Therefore prices in the EU market need not be those an integrated monopolist would have chosen. In this situation inefficiencies are present due to the fact that the ISP charges an average monopoly price to the CPs. Some segments profit from this by paying a price below the monopoly level, thus improving efficiency, while other segments pay more than the optimal monopoly price which might result in some markets not being served.

3.4 Scenario 3: Perfect Discrimination by ISP and Content Monopoly

In this case, the ISP has great power over the CPs since it can target them individually and set individually profit maximizing prices. At first sight this situation seems prone to double marginalization with inefficiencies induced by two cascaded monopoly decision. However, the ISP can anticipate this sort of inefficiency and avoid it by setting an efficient usage price p_w for access to its network and then extract all profits of the monopoly CP through a fixed fee A . This way of pricing is known in the literature as franchising (Tirole 1988). The total price set by the ISP is

$$T(q) = A + p_w q. \quad (8)$$

Assuming the ISP sets $p_w = c$ efficiently at the marginal cost of providing the service (the variable q is the output quantity of the ISP), the CP chooses the optimal end user price p :

$$\max_p [(p - c)D(p) - A]. \quad (9)$$

This results in the monopoly EU price

$$p = \frac{1+c}{2}. \quad (10)$$

By setting the fixed fee A equal to the profit of the CP, the ISP can now extract all profit without distorting the target function of the CP. The same result could have been achieved with a “tax” $\tau \Pi$ levied by the ISP.

Interpretation: An interesting aspect of the profit based charge A is that it does not alter the price paid by consumers but shifts profits from the CP to the ISP. This can be a problem since monopoly profits could be the reward for innovation and if those profits are taken away, innovation might not be profitable any more. Therefore this situation is good for static efficiency since the output decision of the CP is not changed by the ISPs behavior. However, this becomes a problem when considering the development over time. No matter how much of the CP’s revenue is extracted by the ISP, the EU price for the content stays the same and thus there is no competitive pressure from EUs on the ISP.

3.5 Scenario 4: Market Segmentation by ISP and Content Monopoly

In this case each of n CPs serves a monopoly market as in case three but the ISP cannot differentiate between those markets and can only optimize its revenue based on gross demand (as in case two) by all fully differentiated CPs.

This case is similar to a situation know as double marginalization (Tirole 1988). In this situation the CPs (with constant marginal cost set to zero) set their price for content as a standard monopolist. Their target functions are:

$$\max_{p_i} [(p_i - p_w) D_i(p_i)] \quad (11)$$

where p_i is the price charged to the EU and p_w is the price paid to the ISP.

Since the ISP can only 2nd degree price discriminate we will analyze the ISP’s problem based on average figures for price (12) and demand (13) from all CPs demanding quality (using CDN):

$$p = \frac{1}{n} \sum_{i=1}^n p_i, \text{ and} \quad (12)$$

$$D = \frac{1}{n} \sum_{i=1}^n D_i(p_i). \quad (13)$$

Assuming for simplicity that aggregate demand is linear: $D(p)=1-p$, the solution to this problem is identical to the treatment in section 3.2 and yields optimal prices of

$$p = \frac{1+p_w}{2} \quad (14)$$

which corresponds to an average output quantity per CP of

$$q = \frac{1 - p_w}{2}. \quad (15)$$

Using this quantity, the optimization problem of the ISP with marginal output cost c for CDN traffic then is:

$$\max_{p_w} \left[(p_w - c) \left(\frac{1 - p_w}{2} \right) \right] \quad (16)$$

resulting in

$$p_w = \frac{1 + c}{2}. \quad (17)$$

Now we can go back to formula (14) and paste p_w from equation (17)

$$p = \frac{3 + c}{4} \quad (18)$$

for the final consumer price. Comparing the sum of the profits of ISP Π_{isp} and CP Π_{cp} we see that it is smaller than the profit Π_{int} of an integrated monopoly provider:

$$\Pi_{isp} + \Pi_{cp} < \Pi_{int} \quad (19)$$

$$\frac{(1 - c)^2}{8} + \frac{(1 - c)^2}{16} < \frac{(1 - c)^2}{4}. \quad (20)$$

The end user price $p = (1 + c)/2$ in monopoly is lower than that in equation (18) (as long as $c < 1$ which must be because with $D = 1 - p$, $1 > p \geq p_w \geq c$ must be true if no losses are made and there is a positive demand). Note that the above discussion is only true for the average CP. For each individual CP the actual price will depend on the each demand D_i in formula (13) and thus the average price in formula (18) might never be observed.

Interpretation: On average across EUs' demand for the perfectly differentiated markets this situation is suboptimal since prices could be lower and revenues could be higher. In addition to this double marginalization problem there exists a problem due to the averaging of price and demand of the different CPs. Since the price set by the ISP is targeted at the average CP, it will typically be either too high or too low. Thus there is a second source of inefficiency. ISPs could also opt for a different pricing model and charge (as in scenario 3) a fixed fee plus an efficiently set usage fee p_w . Thus inefficiencies due to the ISPs behavior would be removed. However, since the ISP can only set an average $\tilde{A}$ it is impossible to extract all rents from the CPs. $\tilde{A}$ would even make some market segments with low profits unattractive to serve since profits are too low to cover the fixed fee. Thus one has to choose between double marginalization reducing output in some segments and franchising resulting in some markets not being served at all.

3.6 Synopsis of the Four Scenarios

We presented four situations that CPs and ISP might face when using MH or DCN. Scenario one shows that an access monopoly can be exploited by the ISP. Scenario two modifies this

situation by only allowing 2nd degree price discrimination by the ISP. Scenario three demonstrates how an ISP can – through a franchising fee – extract all monopoly profits from CPs. Scenario four modifies this situation and demonstrates the effects of double marginalization as an alternative to a franchising fee with less information needs.

The first conclusion from the above analysis is that in all four scenarios welfare is below optimal and prices are above the competitive level. It is even possible as shown in cases two and four that high average prices can exclude some markets from being served at all. The first general result therefore is that CDNs and multi-homing reduce welfare due to enabling discriminative behavior by the terminating ISP. This result ignores the welfare gains of being able to deliver higher QoS by the use of those technologies.

The second result which is common to all cases is that ISPs' price setting reduces CPs' profits. In some cases all profits may be extracted. While this in itself does not need have a negative effect on welfare in a static environment it can be detrimental when considering monopoly revenues of CPs as the reward for innovation. If ISPs extract these profits innovation might become unprofitable for CPs.

Furthermore with CDN and MH, ISPs are but in a position to exploit their access monopoly and create monopolies from otherwise competitive markets. The only precondition for this is a quality requirement of the service that does not allow the use of ordinary peerings.

To put the derived results into perspective it is important to be aware of the limitations of their applicability. Firstly, the assumption that EUs are perfectly locked in with their ISP can be challenged. Consumers that are able to switch their ISP will certainly not tolerate monopoly prices for content 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. Secondly a large part of the content business is financed by advertisements. Such a financing model changes the considerations of the CP since it must not optimize across a vertical structure but across two distinct markets (Armstrong 2006; Littlechild 2006; Rochet 2003).

4. CONCLUSIONS AND FURTHER RESEARCH

While this paper implicitly relies on the untested hypothesis that CDN and MH are used to improve quality, this assumption is not needed for the general results to hold. The central insight that price discrimination is possible on the internet stays valid. The four scenarios presented above are extreme cases that are unlikely to be observed in the real world. However, with considerable switching costs for changing ones ISP and contracts that run between one and two years, ISPs have considerable market power. It is derived from a temporary access monopoly they can exploit only when CPs switch from using the known internet business model with access reciprocal charges between ISPs to business models that remove reciprocity.

Trying to predict the future is a risky thing but with the above analysis in mind one change that will happen to the internet is not hard to envision. Since CDNs and MH promise profits, there are incentives for ISPs to shift traffic from ordinary peerings to CDN or better yet MH. Slow extension of peering capacity is the most obvious way to degrade quality and thus increase demand for better access to EUs.

The analysis also directs the attention to a weak spot in network research. While academics were looking for the perfect solution to the QoS problem and (as of today) failed to come up with a solution that would be accepted by the market participants, the problem was solved through “innovation at the edges” and created business models that are able to (at least partially) cater for quality needs.

The presented analysis leaves open and poses many research question that need to be addressed in further work. Firstly the four presented scenarios look at extreme cases that could be considered the corners of the solution space. Sampling from the center would for example involve a more dynamic scenario with EUs that can switch providers. This would constrain the ISP in its monopoly power in the paid content case or confront the CP with a two sided market in an ad-financed business model.

Furthermore, the question about the economic efficiency of CDN 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?

Besides economic questions that need further clarification there is also a more technical side. Can CDN and MH fully replace inter carrier agreements on quality parameters of traffic? Which quality mechanisms are necessary inside one carriers network to complement the peering bypass capability of CDN and MH with the ability to deliver to the EUs workstation?

Assuming the business model of the Internet changes according to the predicted way (more CDN and MH), which effects does this have on the Internet as a whole. First effects can be seen in the ever growing routing tables. Which effects do more and more “shortcuts” through the Internet have? Especially the effect on peer to peer networks and the applicability of the described model to peer to peer structures pose interesting questions.

That the lack of quality guarantees on the Internet is a problem is a widely accepted fact (Soldatos 2005) and many solutions have been proposed. However, the solutions that have been proposed so far have had a rather small impact on the Internet. The final solution to the QoS problem therefore might look much different than what is envisioned today.

REFERENCES

- Akella, Aditya, Srinivasan Seshan, and Anees Shaikh. 2003. An empirical evaluation of wide-area internet bottlenecks. In *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*. Miami Beach, FL, USA: ACM.
- Armstrong, Mark. 1998. Network Interconnection in Telecommunications. *The Economic Journal* 108:545-564.
- Armstrong, Mark. 2006. Competition in two-sided markets. *RAND Journal of Economics* 37:668-691.
- Badasyan, Narine & Chakrabarti, Subhadip. 2008. A simple game-theoretic analysis of peering and transit contracting among Internet service providers. *Telecommunications Policy* 32:4-18.
- Bu, Tian, Lixin Gao, and Don Towsley. May 15, 2004. On characterizing BGP routing table growth. *Computer Networks* 45 (1):45-54.
- Cave, M.; Majumdar, S. K. & Vogelsang, I. 2002. *Handbook of telecommunications economics*: Elsevier Boston, Mass.
- Cremer, Jacques; Rey, Patrick & Tirole, Jean. 2000. Connectivity in the Commercial Internet. *The 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.
- DE-CIX. 2008. <http://www.de-cix.net/content/clients.html>.
- Foros, Oystein; Kind, Hans Jarle & Sand, Jan Yngve. 2005. Do internet incumbents choose low interconnection quality? *Information Economics and Policy* 17:149-164.
- Laffont, J. J. & Tirole, J. 2000. *Competition in Telecommunications*: MIT Press.
- Laffont, Jean-Jacques, Patrick Rey, and Jean Tirole. 1998a. Network Competition: I. Overview and Nondiscriminatory Pricing. *The RAND Journal of Economics* 29:1-37.
- Laffont, Jean-Jacques, Patrick Rey, and Jean Tirole. 1998b. Network Competition: II. Price Discrimination. *The RAND Journal of Economics* 29:38-56.
- Laffont, Jean-Jacques; Marcus, Scott; Rey, Patrick & Tirole, Jean. 2003. Internet Interconnection and the Off-Net-Cost Pricing Principle. *The RAND Journal of Economics* 34:370-390.
- Littlechild, S. C. 2006. Mobile termination charges: Calling Party Pays versus Receiving Party Pays. *Telecommunications Policy* 30 (5-6):242-277.
- MacKie-Mason, J.K. & Varian, H.R. 1995. Pricing the Internet. *Public Access to the Internet*:269-314.

- Rey, P. & Tirole, J. 2006. A Primer on Foreclosure. *Handbook of Industrial Organization* 3.
- Rochet, J. C. & Tirole, J. 2003. Platform Competition in Two-Sided Markets. *Journal of the European Economic Association* 1:990-1029.
- Rochet, Jean-Charles & Tirole, Jean. 2006. Two-sided markets: a progress report. *The RAND Journal of Economics* 37:645-667.
- Shakkottai, Srinivas & Srikant, R. 2006. Economics of network pricing with multiple ISPs. *IEEE/ACM Trans. Netw.* 14:1233-1245.
- 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.; Vayias, E. & Kormentzas, G. 2005. On the building blocks of quality of service in heterogeneous IP networks. *Communications Surveys & Tutorials, IEEE* 7:69-88.
- Tirole, J. 1988. *The Theory of Industrial Organization*: MIT Press.
- Uludag, S.; Lui, K.S.; Nahrstedt, K. & Brewster, G. 2007. Analysis of Topology Aggregation techniques for QoS routing. *ACM Computing Surveys* 39 (3).
- Vakali, A., and G. Pallis. 2003. Content delivery networks: status and trends. *Internet Computing, IEEE* 7 (6):68-74.
- Wang, Z. 2001. *Internet QoS: Architectures and Mechanisms for Quality of Service*: Morgan Kaufmann.
- Wu, Tim. 2003. Network Neutrality, Broadband Discrimination. *Journal of Telecommunications and High Technology Law* 2:141.