

2010

# Vertical Platform Interaction on the Internet: How ISPS and CDNS Interact

Thorsten Hau

*University of St. Gallen, thorsten.hau@unisg.ch*

Walter Brenner

*Institute of Information Management University of St. Gallen, walter.brenner@unisg.ch*

Follow this and additional works at: <http://aisel.aisnet.org/ecis2010>

## Recommended Citation

Hau, Thorsten and Brenner, Walter, "Vertical Platform Interaction on the Internet: How ISPS and CDNS Interact" (2010). *ECIS 2010 Proceedings*. 114.

<http://aisel.aisnet.org/ecis2010/114>

This material is brought to you by the European Conference on Information Systems (ECIS) at AIS Electronic Library (AISeL). It has been accepted for inclusion in ECIS 2010 Proceedings by an authorized administrator of AIS Electronic Library (AISeL). For more information, please contact [elibrary@aisnet.org](mailto:elibrary@aisnet.org).



**VERTICAL PLATFORM INTERACTION ON THE INTERNET:  
HOW ISP AND CDN INTERACT**

Journal:	<i>18th European Conference on Information Systems</i>
Manuscript ID:	ECIS2010-0256.R2
Submission Type:	Research Paper
Keyword:	Internet, Internet economics, Economic modeling, Electronic markets



# VERTICAL PLATFORM INTERACTION ON THE INTERNET: HOW ISP AND CDN INTERACT

Hau, Thorsten, University of St. Gallen, Mueller-Friedberg Str. 8, 9000 St. Gallen,  
Switzerland, thorsten.hau@unisg.ch

Brenner, Walter, University of St. Gallen, Mueller-Friedberg Str. 8, 9000 St. Gallen,  
Switzerland, walter.brenner@unisg.ch

## Abstract

*Content delivery networks (CDN) are important players on the Internet. They provide services to websites that improve the quality of service (QoS) and provide end users (EU) with a better experience, i.e. faster loading web pages. CDNs can be seen as platforms that cater for two distinct groups of customers. On the one side, they have content providers (CP), i.e. web sites; on the other side they need to collaborate with Internet service providers (ISP) to reach EUs. We construct a formal model that demonstrates the pricing decisions of ISPs and CDNs and contrast it to the standard types of pricing Internet access and traffic. As a modelling tool we use theory on two-sided markets and bottleneck platforms. We find that ISPs have relatively high market power and extract profits from CDNs to compete for EUs.*

*Keywords: CDN, Network Economics, Internet, Interconnection, Two-Sided Markets*

## 1 INTRODUCTION

Content delivery networks (CDNs) (Buyya, Pathan and Vakali, 2008) are potent players on the market for Internet infrastructure. The market leader Akamai claims to “routinely deliver between ten and twenty percent of all Web traffic” (Akamai, 2009). However, despite the relevance of those firms, there is little research available that analyses the economic impact of CDNs on the Internet. We will argue that new methods to enhance the quality of data delivery to end users (EU) on the Internet (i.e. CDN) are changing its economic structure in the following way: Peering agreements have for a long time shielded content providers from being charged by terminating Internet service providers (ISPs) that control access to EUs. With CDNs, ISPs gain more power to charge CPs and optimize their profits across both customer groups. In other words, Internet peering agreements and the existence of routes with multiple transit ISPs split the market for data termination into two separate parts which were not controlled by a single entity. The advent of CDN has brought these two parts together again onto one platform with contractual relationships with both sides. This work analyses the new value chain of content delivery to EUs and uses a formal model of two vertically interacting platforms to understand the effects of the changing Internet structure.

The Internet is a network of interconnected networks. Local, national or global Internet service providers own these networks. In order to enable data exchange across the Internet these ISPs physically interconnect their networks. This is frequently done in specialized computing centers called Internet Exchange Points (IX) which are commonly operated by third party providers.

A CDN is a network of interconnected servers that are distributed around the Internet (Buyya et al., 2008). Some firms distribute their servers in the major Internet Exchange points like DECIX in Frankfurt (DE-CIX, 2008) while others choose to place several servers in more local facilities owned by individual ISPs. The servers are then physically connected to the network of the ISPs. This interconnection is governed by a contractual agreement between CDN and ISP. To interconnect its servers among each other a CDN uses a mixture of public Internet and private bandwidth. Thus CDNs use a construct which is partly an overlay on the public Internet and partly a parallel infrastructure. On the business side, CDNs follow a relatively simple business model. Their value proposition is faster data delivery to customers through their capacity to locally store data. They sell this capability to any website that wishes to provide a better EU experience through shorter loading times. The primary usage today is the delivery of IP-video streams.

Websites or more generally content providers on the Internet buy capacity on the CDN and upload high bandwidth content like video onto the CDN servers. The CDN automatically determines in which local servers to store the data. The CP only needs to change its site’s html code to link to the CDN’s server instead of its own. When the data is requested from EUs, it is served from the CDN server instead of the origin server. The customers of the CDN benefit from this in two ways. Firstly they do not need to engage in load balancing activity, i.e. managing server capacity which is a first component to provide a website with good loading speed and thus quality of experience for the EU. The second value proposition of CDNs is their closeness to EUs. By being able to provide content from within different ISPs’ networks, they can provide faster data delivery by avoiding long routes across different transit ISPs’ networks and thus improve loading time of web pages.

With the broad adoption of CDN services – today Amazon cloudfront (Amazon, 2009) services offer CDN even to the smallest CP – the Internet undergoes an important change because the business objectives of ISPs are changing from single-sided to two-sided profit maximization. ISPs have a temporal termination monopoly over bringing data to their customers. Therefore they could theoretically charge all companies wishing to transport data to EUs monopolistic fees. However, due to the reality of the Internet terminating ISPs were not able to charge CPs because there was commonly no business relationship between CP and terminating ISP. The two-sided market / platform was split into two separate parts. This situation however is changing with the advent of CDNs because now the CP has a business relationship with the terminating ISP intermediated by the CDN. The two sides of the market are brought together again on one platform.

In this work we will take a rigorous approach to understand the role of CDNs on the Internet and their relationship to ISPs, EUs, and CPs. We will model the vertical structure consisting of CPs, CDNs and ISPs by using two-sided markets (Armstrong, 2006) logic. This approach is new in two aspects. Firstly it shows how the two-sided platform perspective can be applied to better understand the business of providing Internet infrastructure. Secondly constructing a model of a vertical structure that consists of interrelated platforms is new. The authors are not aware of any publication that combines findings from the literature on two-sided markets and on vertical restraints in a formal model.

The remainder of the paper is structured as follows: Firstly we review the relevant economic literature to provide a reference frame for our work. We then construct descriptive and formal models that capture relevant aspects of the interaction between EUs, ISPs, CDNs and CPs. We finish the paper with conclusions and an outlook on future work.

## 2 RELEVANT ECONOMIC CONCEPTS

A comprehensive literature overview on vertical restraints can be found in (O'Brien, 2008). This body of knowledge concerns the competitive situations that arise when an upstream firm provides a critical input for a downstream firm. An example for such a situation is a manufacturer-retailer relationship. The upstream manufacturer might for example provide a monopoly good and one then analyzes the effect different retail structures have on the price level for final customers. One important conclusion of this literature is that small changes in model assumptions may have important effects on the result of the modeling. For example with private negotiations a market with a monopoly supplier and downstream duopoly may lead to efficient marginal cost retail prices whereas the ability to publicly commit to a certain price level may lead to the monopoly outcome (Dobson and Waterson, 2007).

The literature on access pricing in telecommunications (Laffont, 1994; Laffont and Tirole, 2000; Dewenter and Haucap, 2006) is an application of the literature on vertical constraints. This literature considers the economic problems that arise when one carrier provides a vital input for another carrier. Typically, this vital input is forwarding of voice or data communication. Situations that have been abundantly analyzed are interconnections between a local incumbent and a new entrant. Here the incumbent owns the infrastructure that allows it to terminate phone calls with an EU. The new entrant into the market needs to buy this service from the incumbent which is typically reluctant to provide this service since this means supporting a competitor.

A related problem is considered by the literature on interconnection / access between providers where the roles are reciprocal. Whereas above we were faced with a one-way problem (i.e. one company needs access to the customers of another), in this two-way problem each of two firms seek access to the customers of the other firm. Thus the problem is symmetric. The analysis usually considers high termination fees but not with the goal to foreclose a competitor but with the aim of earning supernormal profits. Situations like the high termination charges for calls to mobile handsets are a typical application of this literature (Laffont, Rey and Tirole, 1998a; Laffont, Marcus, Rey and Tirole, 2003; Laffont, Rey and Tirole, 1998b; Laffont, 1996).

The literature on interconnection on the Internet is much less elaborate than that on voice telephony. Important works are Laffont et al. (Laffont et al., 2003; Laffont, 2001) and Cremer (Cremer, 2000). These works analyze the economics of Internet interconnection. Shakkottai and Srikant (Shakkottai and Srikant, 2006) provide an analysis of interconnection taking account of the hierarchical nature of the Internet.

Laffont et al. (Laffont et al., 2003) analyze 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.

One important driver of the two previously presented literatures and models is the principle of reciprocity. This means that what one firm does – for example raising the price for another firm – always has repercussions not only through a change in demand but also through the other firms ability to change the price it charges itself, i.e. to retaliate. This two-way relationship between the two firms which are competing and cooperating at the same time is especially visible in (Laffont, 2001). Here reciprocity is shown to lead to a zero on-net vs. off-net price difference. Thus, besides the technical reasons there are also economic reasons for ISPs to not charge Internet traffic that goes off their own network differently from that traffic that stays on their network. This result, however, is highly sensitive to the assumptions made in the paper.

A two sided market is a market in which a platform optimizes profit across two distinct sets of customers instead of just one; in other words not only the total price level but also the price structure is relevant (Rochet, 2006). 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. Here not only the total cost for making a payment via the card is relevant to determine total demand but the distribution of charges between merchants and customers is important, too. Would the bulk of all charges be borne by the holders of the credit card as opposed to merchants, the profit of credit card companies would be considerably smaller. Other typical examples include night clubs with different prices for men and women and the mostly advertisement-financed media such as TV or newspapers. Relevant papers in this line of research are (Rochet, 2006; Armstrong, 2006; Rochet, 2003).

Rochet et al. (Rochet, 2006) provide a comprehensive overview of the current literature on two-sided markets. They 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 a 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 usage externality exerted by one group of customers on the other, participation incentives are reduced.

Armstrong's discussion of competition in two-sided markets (Armstrong, 2006) provides much of the foundation for this work. He 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.

Musacchio et al. (Musacchio, Walrand and Schwartz, 2007) compare the effects of single and multi-homing of CPs. They provide explicit formulations of welfare under both regimes and offer results for an economy with many ISPs.

Noam (Noam, 2001, p. 111 ff.) takes another perspective and analyzes informally the question whether and how ISPs might exploit their access monopoly. He introduces the notion of "third-party neutrality". The presence of intermediaries prevents the monopoly from being able to price discriminate against CPs since arbitrage would be possible in the presence of different prices and the inability of the ISP to recognize the true origins of traffic. In this sense CDNs are possibly partly relying on such arbitrage as a revenue source. CDNs get cheaper access to ISPs networks than CPs would if they were to interconnect directly with them.

### **3 A DESCRIPTIVE MODEL OF THE CONTENT DELIVERY VALUE CHAIN ON THE INTERNET**

This section provides an intuitive approach to understanding the difference between the standard method of data delivery on the Internet and the use of CDN services.

Figure 1 depicts the standard Internet way of delivering data and the corresponding ad-financed business model. A CP collects advertising money which the advertiser is willing to pay for contact to EUs. Thus the CP acts as a platform facilitating interaction between advertiser and EU. To realize the access to EUs, the CP needs to deliver data across several ISPs' networks to the EU. With millions of EUs, millions of CPs and thousands of ISPs, one ISP will usually not have CP and EU on its network.

At least two ISPs will be involved in data delivery. The source-ISP has the CP as a customer and the terminating ISP has the EU as a customer. Possibly there are also several transit-ISP's involved that have no business relationship to the endpoints of the communication at all. This situation has made it virtually impossible for ISPs to exploit their termination monopoly.




Figure 1: Classical Internet Connectivity

Compare this situation to that in figure 2: Now there are no source and transit ISPs any more but the CDN takes that role. The CDN has a business relationship with both CP and terminating ISP.

The important aspect captured by the different presentations in the two figures is that in figure 1 ISPs compete in the market for EU-access (selling DSL or cable contracts) and on the other side ISPs compete for connecting CPs to the Internet. These two ISPs are not identical and there is no relationship between these two businesses. In other words, ISPs make money from the endpoints of any traffic flow but are always only connected to one of those, either source or sink. Figure 2 in contrast shows the market for data delivery and the market for content as well as their interaction in the “new” world with CDNs. The terminating ISP has – mediated by the CDN - contact to CPs and can engage in two-sided optimization of its profit function. One central reason for this to be possible is that CDNs are like CPs only senders of traffic and not receivers. This situation – from a modeling perspective – gets rid of all reciprocity concerns.




Figure 2: The content delivery value chain.

Hau et al. (Hau and Brenner, 2009) have shown how - when one ISP has a business relationship with both source and sink of a traffic flow on its network and furthermore has a contractual agreement with both sides – which is typically the case when CPs decide to multi-home (i.e. buy direct interconnection to an ISP's network to improve performance of their website) - the ISP will use two sided markets logic to maximize its profit.

In the following we will only consider the market for data delivery as shown in isolation in figure 3. The ISP with its control over the access bottleneck acts as a platform. It collects fees for Internet access from EUs and for termination from CDNs. The vehicle to charge the CDNs are transit payments, i.e. the CDN buys bandwidth from the ISP. The CPs are customers of the CDN and pay the CDN for delivery of their data. The relationship between CDNs and ISPs is complex. On the one hand CDNs buy access bandwidth as a vital input from ISPs. On the other hand, ISPs also have an interest in having CDNs present in their networks since this enhances their quality proposition to their EUs. Furthermore the more CPs a CDN has, the more attractive it is to ISPs and vice versa, the more points of interconnection a CDN has with different ISPs, the more attractive it is to CPs. Thus the CDN functions as a platform and optimizes its profit across ISPs and CPs simultaneously.




Figure 3: The utility and monetary flows in the data delivery market.

In the next paragraph we will model this situation as a sequential game. First the ISP sets its profit maximizing price as a platform. Then the CDN follows and determines the price it charges to CPs under consideration of the value of one side of customers to the other side. Thus we show how the two platforms interact.

#### 4 A FORMAL MODEL OF VERTICALLY INTERACTING PLATFORMS

To formally model the described situation we follow a two step approach. Firstly we follow the bottleneck case described by Armstrong (Armstrong, 2006) to model the ISP. Then in a second step we model the CDN as a price taker (with respect to ISPs) and let it decide on its CP-price based on the price it has to pay to ISPs. We assume that the ISP charges only fixed fees to its customers on both sides. For EUs this is obviously viable considering common flat-rates for DSL or cable. For CDNs, presence or usage fees are both viable but (variable) usage fees introduce a problem due to incompatible units of measure on the ISP and CDN side. We will therefore use fixed fees for the CDN. Furthermore we assume duopolistic competition between ISPs. Then we can derive the following results (Hau and Brenner, 2009; Armstrong, 2006):

Suppose there are two ISPs in a region denoted  $ISP_i, i \in \{1; 2\}$ . There is a fraction  $n_j^i \in [0, 1]$  of agents of Group  $j$  participating on platform  $i$ . In other words,  $ISP_i$  has  $n_1^i$  subscribed EUs and  $n_2^i$  CDNs. EUs are single-homing with their ISP. This means that they are only subscribed with one ISP at a given time. CDNs on the other hand multi-home. They may be connected to zero, one or two ISPs in order to reach potential customers (EUs). The ISPs maximize their respective profits

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

which is a function of the number of EUs times the price they have to pay, plus the number of CDNs times the price they have to pay minus the cost for connecting the two types of customers.

The fraction  $n_1^i$  of EUs that are customers of  $ISP_i$  is given as a function of the utilities offered by the two ISPs as introduced by (Hotelling, 1929):

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

It is thus increasing in the first argument and decreasing in the second. Note that  $n_1^1 + n_1^2 = 1$  holds since EUs do not multi-home and two duopolistic ISPs share a market of total size ‘one’. 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

$$u_1^i = U^i(n_2^i) - p_1^i = \alpha_1 n_2^i - p_1^i \quad (3)$$



and equals the gross utility they get from being connected to  $n_2^i$  CDNs minus the price they have to pay for their Internet connection. The function  $U^i$  is increasing in  $n_2^i$  since more content in better quality is always better than less. The parameter  $\alpha_1$  can be interpreted as the utility an EU derives from being able to reach one CDN. The EUs perceive the ISP with more CDNs connected with QoS as providing a better connection to the Internet.

The fraction  $n_2^i \in [0,1]$  of CDNs that is connected to  $ISP_i$  is given by

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

This is a function of the number of EUs that can be reached through  $ISP_i$  and the price charged. The number of CDNs the ISP can persuade to directly interconnect depends on the parameter  $\gamma^i = p_2^i/n_1^i$  which is the price per EU reached.  $1 - F(\gamma^i)$  then yields the fraction of CDNs that are willing to pay that price.

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.  $\gamma$  represents the expected revenue from ad-clicks per EU and  $\tau$  represents the maximum EU value. The cumulated distribution function is given by

$$F = \gamma/\tau = p_2^i/n_1^i\tau. \quad (5)$$

This yields

$$n_2^i = 1 - \frac{p_2^i}{\tau n_1^i}. \quad (6)$$

We define costs for interconnection as

$$C^i(n_1^i, n_2^i) = cn_2^i \quad (7)$$

which implicitly includes the assumption that the cost of the access network is not part of the considerations for interconnecting with CDNs. 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 ISPs have reached a competitive equilibrium and offer utility  $\hat{u}_1^i$  to their  $\hat{n}_1^i$  EUs, respectively. 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 shares. Since (6) 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 CDNs  $n_2^i$  and the price  $p_2^i$ .

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

$$\begin{aligned}
\Pi^i &= \hat{n}_1^i (U^i(n_2^i) - u_1^i) + p_2^i \overbrace{(1 - F(p_2^i / \hat{n}_1^i))}^{n_2^i} - C(\hat{n}_1^i, n_2^i) \\
&= \hat{n}_1^i (\alpha_1 n_2^i - \hat{u}_1^i) + (p_2^i - c) n_2^i \\
&= \hat{n}_1^i (\alpha_1 n_2^i - \hat{u}_1^i) + [(1 - n_2^i) \tau \hat{n}_1^i - c] n_2^i .
\end{aligned} \tag{8}$$

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

We now find the maximizer of the resulting expression:

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

This is the optimal number of CDNs the ISP should allow on its platform (since the 2<sup>nd</sup> order condition for a maximum holds). Together with (6) this yields the optimal price to CDNs.

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

This shows that the CDNs 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; it is an artifact of the definition of the distribution function (5).

The next step is to model the reaction of the CDN<sup>k</sup> to this price. Let the CDN's profit function be

$$\Pi^k = n_3^k p_3^k - \sum_{i=1}^m p_2^i, \forall i, i \in \{\text{ISP}_i \text{ is customer of CDN}^k\} . \tag{11}$$

This is the number of CPs times the price charged to CPs minus the price paid to ISPs. We neglect further costs and substitute the expression

$$\alpha_i = a_i n_3^k \tag{12}$$

in equation (10) to represent the assumption that the number of CPs on the CDN is a proxy for content quality and thus the externality exerted by CPs on EUs. We assume the demand by CPs for CDN services to be

$$D: n_3^k = \sum_{i=1}^m \tau_i n_1^i - p_3^k . \tag{13}$$

The first order condition for the profit maximization problem of the CDN is given by

$$\frac{\partial \Pi}{\partial p_3^k} = 0. \quad (14)$$

Solving this yields the price charged by CDNs to CPs

$$p_3^k = \frac{1}{4} \left( -a \sum_{i=1}^m n_1^i + 2 \sum_{i=1}^m (n_1^i \tau_i) \right), \quad (15)$$

the resulting number of CPs served

$$n_3^k = \frac{1}{4} \left( a \sum_{i=1}^m n_1^i + 2 \sum_{i=1}^m (n_1^i \tau_i) \right) \quad (16)$$

and for the present case with  $m = 2$  the maximum profit of the CDN

$$\Pi^k = \frac{1}{16} \left( -16c + a^2 \sum_{i=1}^m (n_1^i)^2 + 4a \left( \sum_{i=1}^m n_1^i \right) \left( \sum_{i=1}^m (n_1^i \tau_i) \right) + 4 \left( -2 + \sum_{i=1}^m (n_1^i \tau_i) \right) \left( \sum_{i=1}^m (n_1^i \tau_i) \right) \right). \quad (17)$$

The quadratic nature shows that there is a strong incentive for CDNs to acquire as many EUs as possible. Plotting the function for two ISPs and a parameter choice of  $a = 20$ ,  $\tau_1 = 500$ ,  $\tau_2 = 400$  and  $c = 100$  we can graph the Profit as a function of the number of EUs.




Figure 4: Maximum profit as a function of the number of reachable EUs.

This graph shows how the CDN will set its price dependent on the number of EUs. Since EUs can only be acquired in discrete units (i.e. per ISP) only the four corners of the graph are plausible solutions. The function shows quadratic growth in the number of EUs which shows firstly that EUs are a valuable resource to CDNs and that there are considerable advantages for large CDNs.

## 5 CONCLUSIONS AND FURTHER RESEARCH

This work presents a formal model of a typical value chain for the delivery of high bandwidth content on the Internet. We first motivated the work by explaining the changing structure of the Internet away from a 'flat' network towards a more 'differentiated' network in which different types of data are delivered by different means to the EUs. In fact, the common method to deliver high bandwidth content to EUs today is the usage of CDNs. Local caches store data and distribute local copies to EUs, thus reducing load in the core network.

By means of a formal model we argued that this technological change also entails an economic change. The ‘normal’ Internet provided a mechanism to separate the two end of the content delivery value chain. The complex system of peering and transit agreements ensured that ISPs were not able to act as platforms. This resulted in competitive pricing of Internet access for CPs and - depending on the regulation of the last mile – in either relatively competitive markets for EUs’ Internet access (as in Europe) or mostly duopolistic structures as in the United States. The technological change towards CDN has been changing this situation towards a market in which ISPs can act as platforms. They have at least temporary control over the bottleneck facility ‘EU-access’ and can contract directly with CDNs. The CDNs also have a platform business model since their revenue with CPs and their cost for access to ISPs’ networks are reciprocally dependent. Many CPs imply a high externality on EUs which lowers the price they need to pay to ISPs while many ISPs means many EUs which allows them to charge higher prices to CPs.

We constructed a simple model of a sequential game in which first the ISP chooses its optimal pricing structure and then the CDN chooses its price conditional on the price it has to pay to the ISPs. Instead of solving the discrete optimization problem the CDN is facing (which ISPs to connect to?) we based the analytical treatment on the number of EUs reachable and then argued that only the corners of the solution space are plausible solutions. With this simple model we are able to illustrate that pricing decisions in a platform focused Internet are different from pricing on the “old” Internet with disconnected sides of the platform. It is interesting to note that the profit function of the CDN is quadratic in the number of reached EUs, indicating that CDNs have a strong incentive to connect to as many ISPs as possible.

The presented work is only the starting point for a deeper understanding of the economics of an Internet with differentiated delivery technologies. The first way to extend the presented analysis is the refinement of the presented model. Instead of the simple two stage game a more sophisticated setup with players that react to each others’ decisions would provide further insights into endogenous equilibria as a result of players’ interaction. Furthermore, assuming that CDNs compete in the market for CPs could lead to marginal cost pricing on their side.

We furthermore ignored important costs such as the setup cost for connection to an ISP. Depending on the modeling assumptions for the cost function it would be easy to implement cases where it is optimal to only connect to a subset of ISPs. However, assuming that the interconnection between CDNs and ISPs takes place in IXs’ facilities would make high fixed costs for interconnection implausible.

On a theoretical level, the authors are not aware of any publication that attempts to analyze the vertical interaction of several platforms. In that field the present article can only be a starting point for further development of a theory.

In conclusion, the presented work analyzes two under-researched problem domains: ISP-CDN interconnection and vertical interaction of platforms. This interdisciplinary problem domain appears to be offering more interesting problems and much potential for further research.

## References

- Akamai (2009) <http://www.akamai.com/html/customers/index.html>.
- Amazon (2009) Amazon Cloudfront, 03.07.2009.
- Armstrong, M. (2006) Competition in two-sided markets., *RAND Journal of Economics*, 37, 668-691.
- Buyya, R., Pathan, M. and Vakali, A. (2008) Content Delivery Networks, *Lecture Notes In Electrical Engineering*, 418.
- Cremer, J. R., Patrick & Tirole, Jean (2000) Connectivity in the Commercial Internet, *The Journal of Industrial Economics*, 48, 4, 433-472.
- DE-CIX (2008) <http://www.de-cix.net/content/clients.html>.
- Dewenter, R. and Haucaup, J. (2006) *Access Pricing: Theory and Practice*, Emerald Group Publishing.
- Dobson, P. and Waterson, M. (2007) The competition effects of industry-wide vertical price fixing in bilateral oligopoly, *International Journal of Industrial Organization*, 25, 5, 935-962.
- Hau, T. and Brenner, W. (2009) Price Setting in Two-sided Markets for Internet Connectivity, *LNCS*, Aachen.
- Hotelling, H. (1929) Stability in Competition, *Economic Journal*, 39, 153, 41-57.

- Laffont, J.-J., Marcus, S., Rey, P. and Tirole, J. (2003) Internet Interconnection and the Off-Net-Cost Pricing Principle, *The RAND Journal of Economics*, 34, 370-390.
- Laffont, J.-J., Rey, P. and Tirole, J. (1998a) Network Competition: I. Overview and Nondiscriminatory Pricing, *The RAND Journal of Economics*, 29, 1-37.
- Laffont, J.-J., Rey, P. and Tirole, J. (1998b) Network Competition: II. Price Discrimination, *The RAND Journal of Economics*, 29, 38-56.
- Laffont, J. J. and Tirole, J. (2000) *Competition in Telecommunications*, MIT Press.
- Laffont, J. J. G., I.; Tirole, J. & Geras, I. (1996) Creating competition through interconnection: Theory and practice, *Journal of Regulatory Economics*, 10, 227-256.
- Laffont, J. J. M., S.; Rey, P. & Tirole, J. (2001) Internet Peering, *The American Economic Review*, 91, 287-291.
- Laffont, J.-J. T., Jean (1994) Access pricing and competition, *European Economic Review*, 38, 1673-1710.
- Musacchio, J., Walrand, J. and Schwartz, G. (2007) Network Neutrality and Provider Investment Incentives, 1437-1444.
- Noam, E. (2001) *Interconnecting the Network of Networks*, MIT Press.
- O'Brien, D. P. (2008) The Antitrust Treatment of Vertical Restraints: Beyond the Possibility Theorems, In: *The Pros and Cons of Vertical Restraints* (Ed., Swedish Competition Authority), Tryck AB, Stockholm.
- Rochet, J. C. T., J. (2003) Platform Competition in Two-Sided Markets, *Journal of the European Economic Association*, 1, 990-1029.
- Rochet, J.-C. T., Jean (2006) Two-sided markets: a progress report, *The RAND Journal of Economics*, 37, 645-667.
- Shakkottai, S. and Srikant, R. (2006) Economics of network pricing with multiple ISPs, *IEEE/ACM Trans. Netw.*, 14, 1233-1245.