

On the Economics of Internet Peering

Pio Baake¹, Thorsten Wichmann²

January 1998

Abstract

We discuss economic rationales behind peering decisions in the Internet. In the first part of the paper we analyze the decision about a bilateral peering agreement between two commercial Internet service providers (ISPs) who are in Cournot competition. In the second part we discuss multilateral peering between commercial ISPs and an academic research network (ARN). The latter is organized as club of academics who share the cost of their network. It is discussed whether peering threatens the existence of the ARN and under what circumstances a commercial ISP would want to use strategic pricing to win all ARN-members as customers.

JEL classifications: L11, L31, L86

Keywords: Internet Economics, Interconnection Agreements

1) Humboldt-Universität zu Berlin, Institut für Öffentliche Wirtschaft und Wirtschaftspolitik, D-10178 Berlin, Spandauer Str. 1, baake@wiwi.hu-berlin.de.

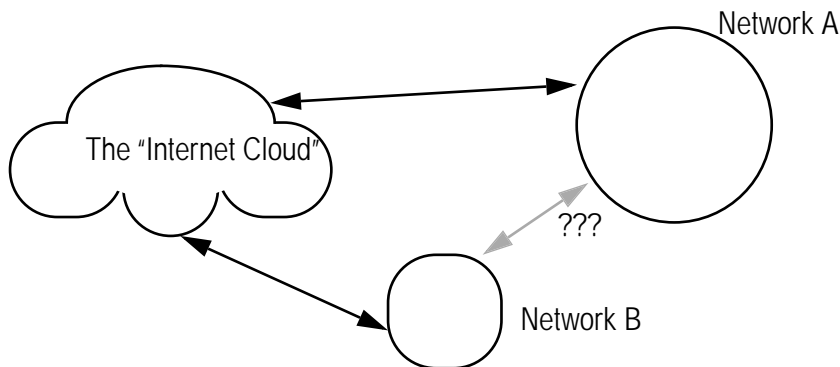
2) Corresponding author. Berlecon Research, Käthe-Niederkirchner-Str. 30, D-10407 Berlin, tw@berlecon.de.

1 Introduction

Interconnection among all parties is one of the main features of the Internet. Every computer connected to the Internet is able to communicate with any other computer on this network. This communication takes place by sending data “packets” back and forward between both machines. Along its way, a packet typically passes several “routers” that determine which way it has to take. This can be a quite winding path; e.g., it is not untypical in Europe that Internet traffic between two computers in the same city but on different networks is routed via the United States.

Typically, such a detour implies a lower quality of service: transferring the same amount of data takes longer and response times (the so-called “latency”) increases. While this does not matter very much for e-mail, it is an important factor for interactive activities like working on a remote computer via telnet sessions or surfing the Internet. Such delays can often be avoided when Internet service providers (ISPs)¹ agree upon an additional, usually local, interconnection between their networks so that traffic between them does not have to be routed through the Internet cloud (figure 1).

Figure 1: The Interconnection Decision



Such interconnection agreements can be bilateral as well as multilateral. In addition, multilateral exchanges can either consist of a central router to which all networks are connected, as is the case in most Commercial Internet Exchanges (CIX) or they can constitute special high-speed networks, as the Metropolitan Area Exchanges (MAE) in the United States.²

1. ISPs in this paper are defined as organizations providing a basic connection to the Internet and managing their own network, however large that may be.

Most CIX and bilateral arrangements are based on cost-sharing agreements: every ISP pays for his own cost of accessing the interconnection point while the costs for the routing equipment are shared. Both direct costs are usually negligible compared to the total cost of running a network. All participants agree to accept all Internet traffic addressed to their network without any further charge. This mutual agreement is called “peering”. Such agreements are thus based on a “sender keeps all” settlement, which distinguishes them from most telephony interconnection agreements (*Farnon/Hurdle, 1997*).

In addition to the performance improvement, such a peering agreement also eases the traffic on the ISPs main lines into the “Internet cloud”. For European ISPs this line is usually a rather expensive intercontinental connection to some ISP in the USA, to whom the European ISP is only customer, not peer and has to pay for the interconnection service.

According to this argument, a peering agreement seems to be of mutual advantage at first sight. However, there do exist some arguments against peering agreements. Two reasons are frequently mentioned by larger ISPs which at least are points against peering with smaller providers. The first can be called “backbone free-riding”. A national ISP *A*, e.g., has to build up and maintain a backbone network connecting the different regions, while an ISP *B* concentrating on a single region does not. If both ISPs agree to peering, *B* uses *A*’s backbone capacity for free to exchange traffic with *A*’s customers in distant regions.

The second reason could be called the “business-stealing effect”. Consider a customer of a large ISP *A*, whose main sites of interest are also customers of the same ISP. For most of his Internet usage, the customer will enjoy a good quality of service in the form of low latency. By switching to a smaller ISP *B* who, for whatever reason, is cheaper and which is not directly connected to network *A*, the customer would usually have to accept a higher latency when accessing his favorite sites on *A*’s network. With a peering agreement between *A* and *B* this barrier to switching falls and *A* might lose customers to *B*. Thus, the decision whether to peer has to take into account all of these possible consequences. While advantageous from the point of view of the customer, ISPs might be reluctant to peer.

2. *Bailey* (1997) distinguishes between peer-to-peer bilateral, hierarchical bilateral, third-party administrated and cooperative Internet interconnection agreements.

In the history of the Internet, peering has traditionally been one of the main modes of interconnection.³ Recently, however, larger ISPs in the USA have terminated several peering agreements. In April 1997 UUNET, one of the major ISPs in the US, announced that it will only continue peering with ISPs that meet certain minimum requirements with respect to their backbone capacity and their quality of service. Already in January MCI, another major ISP, announced a similar policy: it only continues to peer with “national level networks” that also meet certain technical criteria. Most major ISPs in the United States are following this route.

The mentioned criteria are usually only met by other large ISPs, so that effectively the peering agreements between these major and small, often only local, ISPs have been cancelled. Such small ISPs have been offered the status of a customer, meaning that they are now being charged for the interconnection service. Within the Internet community, this change of policy has caused an outcry (*Bielski, 1997*). Also the FCC started investigating, which role its telecommunications regulation should play for Internet services (*Werbach, 1997*).

Within Germany, a somewhat related case has been fiercely debated during 1997. The major commercial ISPs have set up a Commercial Internet Exchange, which is organized as a multilateral peering agreement. The commercial ISPs would like to see the academic research network provided by the DFN to participate in the mutual peering arrangement. Currently this network hosts a considerable number of Germany’s computers connected to the Internet and a large amount of all German Internet traffic starts or ends there.

However, the DFN declined to participate in any peering arrangement. Instead, it suggests that ISPs who want to be well connected to the DFN-network become (commercial) customers of the DFN. For the ISPs that means that they not only have to bear the full technical interconnection costs (leased lines, computing equipment) but in addition are charged for this access considerably more than research institutions (*Heese 1997*). This “peering battle” was fought in the newspapers for almost a year, until the parties settled for a compromise in late 1997. The DFN will participate in the CIX, although it will be compensated for doing so by the commercial ISPs for whom this agreement is still cheaper than the previous rules.

3. For a data-rich overview of Internet interconnection agreements in the USA cf. *Srinagesh (1997)*.

This behavior of large Internet service providers motivates our paper. While at first sight refusing peering agreements seems to stand in stark contrast to possible gains in form of larger customer satisfaction (and thus price) and in form of less international interconnection capacity needed (and thus cost savings), we have also shown some reasons why certain ISPs might want to refrain from peering. In this paper, we want to model the peering decision and to isolate the different factors of influence.

This is done in two steps: We first investigate the decision about a bilateral peering agreement between two commercial ISPs who are in Cournot competition. In a second step we discuss multilateral peering between the commercial ISPs and an academic research network (ARN). The latter is organized as a club of academics who share the cost of their network. Special attention is given to the question, whether peering may threaten the existence of this club.

The remainder of the paper is organized as follows: In section 2 the basic model is set up and in section 3 the peering decision is discussed. Section 4 investigates peering between the academic network and commercial ISPs and section 5 concludes.

2 The Basic Model

Our model is based on two rather simple Internet service providers. Both providers are in Cournot competition with each other. Each provider connects his customers with a high-quality backbone. Since this paper discusses the Internet, both networks are interconnected in some way. We assume that interconnection initially is achieved by a link of much lower quality than the network's backbone. One might want to think of this link as the usual USA connection the typical European Internet service provider has.

We consider a scenario where a) the interconnection quality can be improved by direct peering agreements between the two ISPs, and b) a direct peering agreement c.p. decreases network costs for the ISP. As justification for the latter think of two German ISPs exchanging a substantial amount of traffic via their USA connection. The question we want to answer is, under which circumstances both ISPs would be willing to agree to bilateral peering.

2.1 Assumptions

We assume the existence of two ISPs, i and j . The inverse demand functions for access to ISP i 's network i , $j = 1, 2$ and $i \neq j$ are assumed as:

$$(1) \quad p_i(z, n_i): (p_{iz} < 0, p_{in_i} > 0, p_{izz} < 0, p_{in_in_i} < 0) \\ \text{where: } z = z_i + z_j, n_i = z_i + az_j, 0 < a \leq 1.$$

In this inverse demand function z_i and z_j denote the number of users connected to the respective networks. The first term in the function, z , captures the pure quantity effect. The larger the total number of Internet users, the lower the price as users with a lower willingness to pay join the Internet: $p_{iz} < 0$. In addition a network effect, modelled by n_i , exists. The larger the number of users, the larger each user's utility from being on the Internet and the higher is her willingness to pay: $p_{in_i} > 0$.

We distinguish the size of the network effect n_i and n_j for both networks where each n is a measure for the perceived network size from the point of view of a customer on the respective network. The number of customers on the second network – who can only be reached with lower quality – is discounted by a factor a . The better the quality of interconnection between both ISPs, the larger a .

If there is no direct interconnection between both networks, the interconnection quality is at its lower level, a_0 . A peering agreement increases a up to its upper bound $a = 1$ where users do not experience any difference between accessing resources on their own and on the other network.

For provider i 's cost function we assume

$$(2) \quad c_i(z_i, z_j, a): c_{iz_i} > 0, c_{iz_j} \geq 0, c_{ia} < 0.$$

Costs for an ISP increase with a rise in the number of Internet users on their network ($c_{iz_i} > 0$). Also an additional user on network j will raise costs on the network i ($c_{iz_j} \geq 0$). This is due to the fact that the Internet networks are always interconnected, even if the interconnection is via the Internet cloud. This additional user will at least sometimes access information on the network i , not only producing traffic on i 's backbone but also on his connection between the two networks.

On the other hand, ISPs can decrease their costs by deciding to peer ($c_{ia} < 0$). Generally, the costs for setting up an interconnection point for two networks are rather low, especially if both providers serve the same region or adjacent areas. In comparison, the cost reductions following from a

reduction in the necessary capacity for the connection into the Internet cloud are rather high. This assumption requires that existing customers of network j do not access considerably more information on network i when the interconnection quality improves. Otherwise the additional costs due to j 's customers accessing resources on i 's network might overcompensate the cost savings from fewer intercontinental lines. While $c_a > 0$ does not change the model's outcome considerably, we regard $c_a < 0$ as the more interesting case. The model implications also do not change if we account for a fixed interconnection cost $k(a)$.

2.2 Profit Maximization

Each ISP maximizes its profits given by

$$(3) \quad \pi_i = p_i(z, n_i)z_i - c_i(z_i, z_j, a)$$

taken as given the network size of the second provider. Thus the first-order condition for provider i is given as

$$(4) \quad \pi_{iz_i} = (p_{iz} + p_{in_i})z_i + p_i - c_{iz_i} = 0.$$

Condition (4) yields the optimal network size $z_i^*(z_j, a)$ as a function of the size of the second network and the interconnection quality. The second-order condition to be met is $\pi_{iz_i z_i} < 0$ as we assume interior solutions. The reaction functions for provider i can be obtained in the usual way from the first-order condition (4) as:

$$(5) \quad z_i^* z_j = -\frac{\pi_{iz_i z_j}}{\pi_{iz_i z_i}} \text{ and } z_i^* a = -\frac{\pi_{iz_i a}}{\pi_{iz_i z_i}}.$$

Stability of the Cournot equilibrium requires that

$$z_i^* z_j > \frac{1}{z_j^* z_i}.$$

3 The Peering Decision of Commercial ISPs

The peering decision in our model is a decision about interconnection quality a . Both commercial ISPs will agree to peer if they can increase their profits by this agreement. If both providers are in equilibrium initially, the condition to be met for making a peering agreement profitable can be derived from equation (3) by using the envelope theorem:

$$(6) \quad \begin{aligned} \frac{d}{da} \pi_i^* &= \pi_{i^* a} + \pi_{i^* z_j} \frac{d}{da} z_j^* \\ &= p_{in_i} z_j z_i^* - c_{ia} + [(p_{iz} + p_{in_i} a) z_i^* - c_{iz_j}] \frac{d}{da} z_j^* > 0 \end{aligned}$$

However, it is not possible to derive general conditions for preferences and technology from (6) under which peering is profitable or unprofitable. Therefore we proceed by first analyzing the different factors determining the outcome of equation (6) and subsequently discussing some examples by simulation.

First of all, the right-hand side of equation (6) can be divided into three parts: The first term, $\pi_{i^* a}$, captures the direct effect of a change in interconnection quality a on i 's profits. The second part, $\pi_{i^* z_j}$, is the indirect effect of a change in the competitor's network size on i 's own profit. And, finally, the last part $\frac{d}{da} z_j^*$ contains the reaction of the other provider's optimal network size when a changes.

The direct effect of an increase in interconnection quality is given by $p_{in_i} z_j z_i^* - c_{ia}$ and is positive by our assumptions about costs and preferences. Apart from the network sizes, its value is determined by cost as well as by price derivatives. It is the former, which are mainly pointed out by proponents of peering agreements. Since $c_{ia} < 0$, an increase in a raises π_i^* through this channel and thus increases the chances that peering is profitable. The more elastic costs are with respect to an increase in a (the smaller c_{ia}), the larger this effect and the larger $\pi_{i^* a}$: The chances that peering is profitable rise.

The network effect captured by $p_{in_i} > 0$ determines the second reaction. A rise in a increases the perceived network size n_i for i 's customers and thus c.p. the price they are willing to pay for their Internet usage. This effect becomes less important as the total size of both networks and the quality of their interconnection increase since both raise n_i and we have assumed $p_{in_i} < 0$.

The second term $\pi_i^* \frac{d}{dz_j} z_j^*$ in condition (6) consists of two parts. Consider $\pi_i^* = (p_{iz} + p_{in_i} a) z_i^* - c_{iz_j}$ first. From the first order condition (4) for i follows that $p_{iz} + p_{in_i} = -(p_i - c_{iz_i})/z_i^*$ which leads to $p_{iz} + p_{in_i} < 0$ since prices are set above marginal cost. A rise in the competitor's network size decreases i 's profits.

Just like the direct effect also the indirect one can be split into two parts, a cost effect as well as a price effect. As for the cost effect, π_i^* falls in c_{iz_j} , i.e., in the cost increase for provider i induced by an additional user on network j who also produces traffic on i 's network. This cost effect might be especially important if "backbone free-riding" is a serious issue. If, for example, ISP i is regionally dispersed, she incurs large costs from extending the capacity of her large and expensive backbone to cope with the additional traffic from an increase in z_j . For a regional provider, these costs are usually much lower.

The price effect in π_i^* is driven by the Cournot competition. If j increases her network size, provider i has to adjust its price to keep her network size z_i^* . The strength of this effect depends on the price derivatives. The larger p_{in_i} and the smaller $|p_{iz}|$, the lower is also $|\pi_i^*|$. In the crucial case of $\frac{d}{da} z_j^* < 0$ both effects raise the chances of peering being profitable. Note that the effect of p_{in_i} works in the same direction as above with $\pi_i^* a$.

Since $\pi_i^* < 0$, the sign of the second term in (6) is determined by the sign of $\frac{d}{da} z_j^*$. Hence, the smaller $\frac{d}{da} z_j^*$, the higher the chances that peering is profitable. If this term is negative, peering will always be profitable.

Under reasonable assumptions such a decrease in the competitor's network size will be accompanied by an extension of one's own network: Taking the total derivative of provider i 's first-order condition, which has to equal zero in equilibrium, leads to

$$\pi_i^* \frac{d}{dz_i} z_i^* + \pi_i^* \frac{d}{dz_j} z_j^* + \pi_{iz_i a} = 0.$$

Combining this with equation (5) for provider i implies

$$\frac{d}{da} z_i^* - z_i^* \frac{d}{dz_j} z_j^* - z_i^* a = 0.$$

With the additional assumptions of reaction functions falling in the opponent's network size ($z_i^* < 0$) and rising in interconnection quality ($z_i^* a > 0$) it follows that $\frac{d}{da} z_i^* > 0$ if $\frac{d}{da} z_j^* < 0$. (The opposite, however, does not hold.)

Hence, while in such a situation peering decreases j 's network size, it increases i 's network – certainly a potential source of conflict that has been discussed as “business stealing” in the introduction. However, it is easy to see from condition (6) that peering might still be profitable despite losing market share. Even if provider j loses customers to i , peering might still be advantageous for her if only her profit decrease $\pi_j^* z_i$ from the rise in the other provider's network size is not too large. In this case the positive direct effect $\pi_j^* a$ dominates.

Summarizing the different elements determining the peering decision, we can conclude that this decision is not only influenced by cost but also by price effects. While large cost savings from fewer intercontinental lines raise the chances of peering being profitable, a large cost sensitiveness to increases in the competitor's network size decreases the chances. The latter is usually more important for large, regionally dispersed providers than for small ones. However, even if costs would increase from interconnection, profits need not fall since customers are willing to pay for the increased network size.

The combined influence of the different effects mentioned can best be shown by calculating some numerical examples. For doing this, we have chosen the following inverse demand function:

$$p_i = (1 + \alpha n_i - z)^{0.5}$$


as well as the cost functions

$$c_i = \beta_i((2 - a)z_i + z_j),$$

which both satisfy the above restrictions on costs and preferences.


Consider first the baseline setting of identical providers. In this setting the equilibrium is symmetric and both providers either unanimously favor peering or reject peering. Given the above demand and cost functions, both ISPs will agree to peer for most parameter values. Figure 1 shows how network size and profit change for both providers as the level of interconnection quality rises from $a = 0.2$ to $a = 1$. We have chosen an initial value greater zero for a since there is always some low-quality interconnection in the Internet. The other parameters for this scenario are $\alpha = 0.1$ and $\beta_1 = \beta_2 = 0.1$.

Figure 1: Symmetric Solution



Next consider providers with different cost structures. We assume for this example that the first provider has lower marginal costs for each network size ($\beta_2 = 0.046$). Figure 2 shows first of all that the low-cost provider will be larger and more profitable.

Figure 2: Differences in Cost Structure




Starting with $a = 0.2$, the second provider would prefer any increase in interconnection quality to her current state. For the first provider the outcome of her peering decision depends very much on the actual interconnection level chosen. Up to rather high levels of a profits are lower than in the start position. Only with a very high interconnection quality are profits higher than initially.

Note that the first ISP's network size z_1^* decreases in a . However, even with a decreasing network size and market share her profits rise if interconnection quality is above a certain level, which is due to cost savings and price increases. The latter have been made possible by the higher quality of the interconnected network.

Note also that in this example network size changes as well as profit changes are much smaller for the first than for the second provider. Even if perfect interconnection is not possible for some reason, the second provider could therefore use part of his increase in profits from better interconnection to compensate the first for his fall in profits.

Figure 3: Low Network Importance



We have already pointed out above that the network effect can possibly become very important as it may lead to higher profits despite a falling market share. Figure 3 illustrates this. Parameter values are the same as for figure 2, except that now $\alpha = 0.01$ which corresponds to a lower willingness to pay for an increase in perceived network size, p_{in_i} . In this scenario, the first provider would prefer to keep his low-quality link. Of course the sec-

ond provider could again compensate her competitor for the profit fall just like in the previous scenario.

The importance of the network effect is probably one of the main differences between consumer and business networks. It is reasonable to assume that business users generally have a higher willingness to pay for a high quality interconnection than users who only send e-mail and do some occasional surfing. Thus for the question whether peering will take place between “business networks”, this effect is probably more important than for the same question concerning “consumer networks”.

4 Peering between Academic and Commercial ISPs

In the previous section we have discussed the peering decision for commercial ISPs which are in Cournot competition with each other. However, commercial ISPs are a rather new group on the Internet, which has started out as a network of academic research networks (ARNs) based on cost sharing. Therefore an ARN might consider the above discussion as not appropriate since it does not consider itself being a commercial network. In fact, one of the “peering battles” set out in the introduction, that between the German academic research network DFN and a group of commercial ISPs is exactly such a case where the group of commercial ISPs was already peering and wanted the DFN to join this multilateral peering agreement.

We will therefore discuss in this section the peering decision in such a scenario, where one side is commercial and the other is based on cost-sharing. Our basic setup for this question will constitute of two identical commercial providers, who already are perfectly interconnected on one side and of a “club”, the academic research network (ARN) on the other side. The latter is non-profit and based on cost sharing. While members of the ARN are able to leave the club and get connected by one of the commercial providers, customers of the latter can not become members of the ARN. Since members can always leave, it is also interesting to know, whether peering might increase this incentive and thus lead to dissolution of the ARN.

We start out with a situation where the ARN already exists due to some historic event not further discussed. It consists of M homogeneous members. The two commercial providers are in a symmetric equilibrium with perfect interconnection as discussed in section 3. The ARN is initially connected to the commercial ISPs with an interconnection quality a_0 where $0 < a_0 < 1$.

Each individual member can always leave the ARN. It will decide to do so if its utility from being a customer of an ISP is larger than from remaining in the academic network, given that the other members remain with the ARN.

We assume that an academic's total utility – when being a member of the ARN – is given by a function depending on the perceived network size of his academic community M^{ARN} as well as of the commercial Internet community N^{ARN} and diminished by his cost share.

$$U^{ARN} = \bar{u}(M^{ARN}, N^{ARN}) - \frac{c(M, z, a)}{M}$$

$$\text{where } M^{ARN} = M, N^{ARN} = az = a(z_1 + z_2)$$

$$\text{and } \bar{u}_{M^{ARN}} > 0, \bar{u}_{N^{ARN}} > 0.$$

If the member switches to one of the ISPs, it has to accept this ISP's price p_i instead of a cost share. In addition it gains a perfect Internet connection to customers of the commercial ISPs while the connection to his research fellows is of a lesser quality. Thus, his utility becomes

$$U_i^{ISP} = \bar{u}(M^{ISP}, N^{ISP}) - p_i$$

$$\text{where } M^{ISP} = aM, N^{ISP} = z = (z_1 + z_2)$$

$$\text{and } \bar{u}_{M^{ISP}} > 0, \bar{u}_{N^{ISP}} > 0.$$

We assume that M is large so that the member's switching decision only has a negligible influence on M .

Given these utility functions, the member will remain with the ARN as long as $U^{ARN} \geq U_i^{ISP}$. Thus, existence of the ARN requires that

$$(7) \quad \bar{u}(M^{ARN}, N^{ARN}) - \bar{u}(M^{ISP}, N^{ISP}) \geq \frac{c(M, z, a)}{M} - p_i.$$

Given that condition (7) holds, we assume that the ARN maximizes its member's utility. We will not discuss possible principle-agent problems between members and management nor any other possible objectives of the ARN's management besides the one assumed.

Both ISPs maximize profits, taking into account the existence of the ARN:

$$\pi_i = p_i(z, n_i)z_i - c_i(z_i, z_j, M, a),$$

$$n_i = z_1 + z_2 + aM.$$

4.1 The Peering Decision

As before, the commercial ISPs will favor peering if it increases their profits. Since interconnection between both ISPs is perfect and the peering decision concerns interconnection with the ARN, their change in profits is now given by

$$(8) \quad \begin{aligned} \frac{d}{da} \pi_i^* &= \pi_{i^* a} + \pi_{i^* z_j} \frac{d}{da} z_j^* \\ &= p_{in_i} M z_i^* - c_{ia} + [(p_{iz} + p_{in_i}) z_i^* - c_{iz_j}] \frac{d}{da} z_j^* > 0 \end{aligned}$$

As both ISPs are symmetric, (8) can be transformed to yield the following condition for profitable peering:

$$\frac{p_{in_i} M z_i^* - c_{ia}}{(p_{iz} + p_{in_i} a) z_i^* - c_{iz_j}} \geq \frac{z_i^* a}{1 - z_i^* z_j^*}.$$

In comparison, the ARN will agree to peer if it can increase its members' utility. This is the case if

$$(9) \quad (z^* + a z_a^*) \bar{u}_{N^{ARN}} - \frac{c_a + c_z z_a^*}{M} \geq 0 \quad \text{where} \quad z_a^* = 2 \left[\frac{z_i^* a}{1 - z_i^* z_j^*} \right].$$

For the commercial ISPs the peering decision is very similar to that discussed above. The term $p_{in_i} M z_i^*$ contains the influence of the ARN's size. The larger this network, the more do ISPs favor peering. The second term c_{ia} is now the cost decrease from better interconnection between academic and commercial network. Finally, the third large term contains an effect one would tend to neglect at first sight: the interconnection decision and the subsequent increase in the total size of the perceived network change the optimal network size for both ISPs. Basically, the increase in interconnection quality a shifts the inverse demand curve p_i upwards. It thus works like an exogenous demand increase and – at least in general – raises z^* . Just like in the previous section, ISP i has to take into account the influence of the changing z_j^* on his own profitability.

The ARN has to take into account a utility effect as well as a cost effect. Since we have assumed $\bar{u}_{N^{ARN}} > 0$, the combined size of the commercial networks, z^* , as well as its change, z_a^* , positively influence the chances of peering. The size of this effect depends on the influence of the commercial ISPs network size on ARN members' utility which is given by $\bar{u}_{N^{ARN}}$. The cost effect

works as one would expect: $c_a < 0$ works in favor of peering, $c_n > 0$ against it. The influence of the latter depends on the peering-induced change in z^* .

Note that the ARN should agree to peering, even if such an agreement implies larger costs for its members, as long as the increase in utility is large enough to meet condition (9).

In the mentioned peering battle between DFN and a group of commercial ISPs in Germany much of the debate concerned the quantity of the derived effects. The commercial ISPs asserted that $\bar{u}_{N^{ARN}}$ is of considerable size for the ARN since its members access considerable amounts of information on the commercial networks. Furthermore they suggested that cost savings from peering would be large for both peering parties. However, the research network questioned the cost savings, neglected the possible utility gains and instead pointed to $p_{in_i} M z_i^*$. As the research network is large and contains a lot of useful information, so they argued, peering would ease access to this material, thus making the commercial networks more valuable. In our model this corresponds to a relatively large p_{in_i} .

4.2 Does Peering Threaten the Existence of the ARN?

In our model ARN members might leave their club at any time and switch to a commercial Internet provider. An interesting question is therefore, whether peering increases their incentive to do so by decreasing the advantages from being a member of the ARN. The question is, how condition (7) changes, as ISPs and ARN agree to peering. It is obvious that this existence threat is larger, the smaller the left hand-side of condition (7) and the larger the right hand-side.

Consider the right hand-side first, where a change in a has the following effect:

$$\frac{d}{da} \left[\frac{c(M, z^*, a)}{M} - p_i \right] = \frac{c_a + c_z z_a^*}{M} - (p_{iz} + p_{in_i}) z_a^* - p_{in_i} M.$$

Assume that $z_a^* > 0$, which will generally be the case as we have pointed out above. Then the second term on the right is negative by the first-order conditions for z_i^* . The third term is positive. The sign of the first term, and thus the overall effect of peering on average cost in the ARN depends on c_a . If the cost savings from peering are rather small, as the ARN suspects, average costs will rise and so will the whole expression.

A relatively large marginal effect of peering on the commercial ISPs' combined network size, z_a^* , will also raise the right hand-side of condition (7).

First of all it will lead to larger average costs for the ARN members. And second it will lead to lower prices in the commercial part of the Internet, making switching more advantageous.

Last, a large club size M implies a rather small right hand-side of (7) if the network effect is important, as the commercial price will be higher the larger M . Thus it will stabilize the ARN.

For the left hand-side we get

$$\begin{aligned} & \frac{d}{da} [\bar{u}(M^{ARN}, N^{ARN}) - \bar{u}(M^{ISP}, N^{ISP})] \\ & = (z^* + az_a^*)\bar{u}_{N^{ARN}} - (M\bar{u}_{M^{ISP}} + z_a^*\bar{u}_{N^{ISP}}) \end{aligned}$$

With the above assumptions on the derivatives of \bar{u} it is clear that a small z^* as well as a large M make this expression small and increase the existence threat on an ARN agreeing to peer. If M is large, an academic loses a large amount of utility when switching to a commercial provider. Thus, also his gains from an increase in a are much higher than with a small M . Likewise, if z^* is small, an increase in interconnection quality raises utility for a member of the ARN by a smaller amount than if the size of the commercial ISPs were larger. Note that the size of M works in the opposite direction as above.

Finally, the effect of z_a^* on the left hand-side is not clear. It depends on the size of a , $\bar{u}_{N^{ARN}}$ and $\bar{u}_{N^{ISP}}$. With a increasing to one and the derivatives not too different, one would suppose that its influence is small, if present at all.

Summarizing the discussion, one can say that peering might indeed threaten the existence of the academic research network. The chances for such a threat will be larger if the cost savings from peering in the ARN are small, if the commercial networks gain considerably in size from peering and if their combined size is nevertheless not too large.

4.3 Strategic Pricing by Commercial ISPs

Even if the ARN remains in existence in the scenario just discussed, the commercial ISPs might have an incentive to use pricing as a strategic instrument to dissolve the ARN by winning all ARN members as customers of their commercial network. In this section we will discuss whether a change in the quality of interconnection will influence the decision of the commercial ISPs whether to use strategic pricing.

To keep things as simple as possible, we introduce two further assumptions: First, we assume that only one provider, i , will use strategic pricing to win the ARN members as customers and the second provider behaves just like before. Second, we assume that price discrimination is possible and i can offer a different price to ARN members than to non-ARN members. The price for the latter is, just like before, the outcome of the Cournot competition for non-academics. We assume that these prices are derived sequentially. First, i will make its strategic pricing decision and afterwards it will engage in Cournot competition with j for the non-ARN members.

ISP i only has an incentive to use strategic pricing for winning the ARN's members as customers if she can increase her profits by doing so. The ISP can only win the members if these are indifferent between either remaining in the ARN or switching altogether to the commercial ISP, thereby dissolving the ARN. This is the case if

$$(10) \quad \bar{u}(M, 2az_i^*) - \frac{c(M, 2az_i^*, a)}{M} = \bar{u}(M, \hat{z}_i + z_j) - \tilde{p}_i,$$

where \hat{z}_i is the number of non-academic customers at provider i and \tilde{p}_i is the discriminatory price for academics on the commercial network.

Since the provider is able to discriminate between academic and non-academic customers, she will choose the connection price for academics such that equation (10) is met. The price for non-academic customers will be given by (1) where z and n_i are such that

$$(11) \quad p_i = p_i(\hat{z}_i + z_j, \hat{z}_i + M + z_j).$$

Provider i 's profit will be given as

$$\tilde{\pi}_i = \tilde{p}_i M + p_i \hat{z}_i - c_i(\hat{z}_i + M, z_j).$$

Note that the cost function no longer depends on a since the ARN has ceased to exist and both commercial providers are perfectly interconnected.

For the second ISP profits and optimal network size are given by

$$\tilde{\pi}_j = z_j p - c_j(z_j, \hat{z}_i + M)$$

with p_j given equivalent to (11). Provider j maximizes her profits taken M and \hat{z}_i as given, which leads to her optimal network size $\tilde{z}_j^*(\hat{z}_i, M)$ from her first-order condition

$$(p_z + p_n)z_j + p - c_{jz_j} = 0.$$

The first-order condition for network i leads to its equilibrium quantity of non-academic customers, \hat{z}_i

$$[M\bar{u}_N(M, \hat{z}_i + \tilde{z}_j^*) + \hat{z}_i(p_z + p_n) + p - c_{i\hat{z}_i}]\hat{z}_i = 0$$

In the remainder we will only consider the interior solution where $\hat{z}_i > 0$. Whether ISP i will pursue this pricing strategy will depend on \tilde{p}_i , the discriminatory price it has to offer the ARN-members. This price depends again on the current state of interconnection, a , as condition (10) shows. We can therefore ask, whether an increase in a raises potential profits from strategic pricing and thus the chances that this strategy is chosen. This derivation can be obtained from i 's profit equation, taking into account (10) as:

$$\frac{d}{da}\tilde{\pi}_i^* = M\left[-\bar{u}_N(M, a(z_i^* + z_j^*))\left(2a\frac{d}{da}z_j^* + 2z_i^*\right) + \frac{c_a}{M} + 2\frac{c_z}{M}\frac{d}{da}z_j^*\right].$$

Compare this, however, to the ARN's condition for agreeing to peering. This condition (9) can be transformed into

$$\bar{u}_N(M, a(z_i^* + z_j^*))\left(2a\frac{d}{da}z_j^* + 2z_i^*\right) + \left(-\frac{c_a}{M}\right) - 2\frac{c_n}{M}\left(\frac{d}{da}z_j^*\right) > 0.$$

Thus, if condition (9) is met and the ARN agrees to peering, than any increase in interconnection quality a decreases the profitability for ISP i of taking over the ARN's customers. Therefore an agreement to peering protects the ARN from dissolution. This means of course also that the chances of such an unfriendly take-over are larger, when the ARN could increase its member's utility by agreeing to peer but does not do so.

Compared to the previous subsection, the consequences of peering for the existence of the ARN are just contrary. It can protect itself from strategic pricing by a commercial ISPs with peering, but under not unrealistic assumptions increases the chance that its members switch to the commercial provider without any further incentive.

However, the Cournot price p_i for voluntary switching to the commercial provider has to fall below the strategic price \tilde{p}_i , as a comparison of conditions (7) and (10) shows. The member's decision discussed in 4.2 is an individual decision and the member has to take into account the disutility from reaching his research fellows with a poorer quality. To compensate for this disutility, the price has to be lower than in the case of strategic pricing, where all members switch at the same time. Thus the strategic pricing threat is stronger and it is therefore still the best strategy for the ARN to agree to peering – although not to too much.

5 Conclusions

In this paper we have discussed economic rationales behind peering decisions in the Internet. The first part of the paper was concerned with bilateral peering between commercial ISPs which are in Cournot competition, while the second part concerned multilateral peering among commercial ISPs and an academic research network based on cost sharing.

While it is not possible to derive unambiguous conditions for peering to be favorable, we have shown that certain conditions and constellations are more favorable to peering than others.

In a peering decision between commercial ISPs, large cost savings, e.g. from fewer intercontinental lines, raise the chances of peering being profitable. In the opposite direction works a large cost sensitiveness to increases in the competitor's network size. The latter effect is usually more important for large, regionally dispersed providers than for small ones.

While it might be the case that customers leave for another, cheaper provider, when their ISP starts peering with this cheaper provider, the “business-stealing effect” mentioned in the introduction, the resulting feasible price raises justified by customers' willingness to pay for the increased network size might even be strong enough to compensate a provider for the resulting revenue loss.

For the peering decision between an academic research network the cost effects work in the same direction. While an increased interconnection quality does not yield higher profits in this environment, it increases the ARN members utility from being a member of the network and should therefore be taken into account.

For the ARN is especially important, whether peering might threaten its existence. We have shown that this might indeed be the case. The chances for such a threat will be larger if the cost savings from peering in the ARN are small, if the commercial networks gain considerably in size from peering and if their combined size is nevertheless not too large.

However, a peering agreement decreases the profitability for a commercial ISP of using strategic pricing to win all members of the ARN as customers. We have shown that this strategic pricing constitutes a stronger threat to the existence of the ARN and that it therefore should agree to a peering arrangement if that increases its members' utility. However, the interconnection quality should not be chosen too high.

6 References

- Bailey, Joseph P. (1997): The Economics of Internet Interconnection Agreements. In: L.W. McKnight; J.P. Bailey (Eds.): Internet Economics, Cambridge: MIT Press, 155-168.
- Bielski, Vince (1997): David Vs. Goliaths, in: The Village Voice Worldwide. <http://www.villagevoice.com:80/ink/cyber.html>
- Farnon, Maria; Scott Hurdle (1997): Settlement Systems for the Internet. In: B. Kahin; J.H. Keller (Eds.): Settlement systems for the Internet, Cambridge, MA: MIT Press, 377-403.
- Heese, Sven (1997): Tarifierung und Preisliste zur Nutzung der Breitband-Infrastruktur des DFN-Vereins, March. <http://www.dfn.de/win/bwin/tarif-konzept.html>.
- Srinagesh, Padmanabhan (1997): Internet Cost Structures and Interconnection Agreements, in: L.W. McKnight; J.P. Bailey: Internet Economics, Cambridge: MIT Press, 121-154
- Werbach, Kevin (1997): Digital Tornado: The Internet and Telecommunications Policy, FCC OPP Working Paper Series 29, March.