Competition and Cooperation in Internet Backbone Services*

Margit A. Vanberg

Abstract This paper analyzes the strong network externalities associated with Internet services from a competition policy perspective. In the market for Internet services, network effects are so important that an ISP needs to be able to offer universal connectivity in order to survive in this market. To reach universal connectivity, new entrants to the Internet interconnectivity market need to establish a direct or indirect transit agreement with at least one Tier-1 ISP. The fear that a single Tier-1 ISP could abuse a dominant market position in a transit agreement with lower level ISPs is not substantiated by the analysis. Competitive forces in the market for top-tier Internet interconnectivity are strong. A collusion between Tier-1 ISPs to collectively raise prices in the transit market is also not likely to be stable because the prerequisites for a stable collusion are not fulfilled in the market for top-tier Internet interconnectivity services. The analysis supports the view that competitive forces in the transit market are working and can effectively hinder Tier-1 ISPs from discriminating ISPs that are on lower levels of the Internet hierarchy.

Introduction

This paper discusses the effect of the strong network externalities that are associated with Internet service provision on competition in the market for Internet backbone services. In Internet services network effects are so important that an Internet service

M.A. Vanberg
Centre for European Economic Research (ZEW)
Research Group Information and Communication Technologies
e-mail: vanberg@zew.de

*This paper is based on Chapter 7 of Vanberg (2009).

B. Preissl et al. (eds.), *Telecommunication Markets*, Contributions to Economics,
DOI: 10.1007/978-3-7908-2082-9_3,

provider (ISP) needs to be able to offer universal connectivity in order to survive in this market. To reach universal connectivity, new ISPs need to establish a direct or indirect transit agreement with at least one Tier-1 ISP. The focus in this paper is on understanding the consequences of network externalities on market structure in the Internet backbone services market from a competition policy perspective. U.S. and European competition authorities have studied the effects of network externalities on competition in Internet backbone services extensively.[1]At the focus of their analysis were the proposed mergers of large telecommunications companies (MCI and Worldcom and later of MCIWorldcom and Sprint) with notable market shares in the Internet backbone services market. The question concerning the competition authorities was whether a larger provider of Internet backbone services would have an incentive and the means to discriminate against smaller rivals because of network externalities in the market?

Based on the *disaggregated regulatory approach* (Knieps 1997 and 2006), the logical layer of Internet service provision is analyzed in isolation from the vertically related upstream market for physical network infrastructure (the physical layer) and the downstream market for Internet applications services (the applications layer). The main services provided on the logical layer of Internet service provision are Internet traffic services: *Internet access services*, which are provided on top of local communications infrastructure and serve to transmit Internet traffic between the end-users premises and a point of presence of an ISP's network and *Internet backbone services*, which are provided over long-distance communications infrastructure and serve to transmit data within an ISP's networks and between ISPs' networks. The main network elements of the logical layer are routers and switches which are combined with software and Internet-addressing standards. Furthermore, network management functions and the negotiation of interconnection agreements belong to the logical layer. The communication lines over which Internet traffic is transmitted are part of the physical layer of Internet service provision.

The paper is structured as follows: Section "Network Effects in Internet Service Provision" introduces the specifics of network externalities in the applications layer of Internet service provision and how they relate to the logical layer of Internet service provision. Section "Terms of Interconnection Among ISPs in a Competitive Environment" reviews the terms of interconnection between ISPs which are observable in today's unregulated Internet interconnection markets. Section "Dominance at the Tier-1 Level" reviews the literature on interconnection incentives of ISPs with a focus on the single-dominance case. Section "Collusion on the Tier-1 Level" analyzes whether the Tier-1 ISPs as a group could successfully form a stable collusion on the market for transit services and thereby collectively discriminate ISPs on lower hierarchy levels (collective dominance). Section "Conclusions" concludes the paper.

[1] See European Commission (1998, 2000).

Network Effects in Internet Service Provision

The Internet, as a classical communications network, belongs to the class of goods which exhibit positive external benefits in consumption. Direct external effects are due to the fact that the utility of belonging to the Internet community is directly related to the number of other people and services that can be reached via the Internet. Indirect network effects result from the fact that the more people use Internet services, the more applications and complementary products are offered to Internet users.

The utility derived from the consumption of any network good can be decomposed into a so-called network effect, resulting from the number of people reachable via the network, and a so-called technology effect, resulting from the technological characteristics of the network the user is connected to (Blankart and Knieps 1992: 80). In the context of Internet service provision the network effect can be expected to dominate the technology effect because users are more likely to give up benefits from a preferred technology for a wider reach in the Internet.

One way of maximizing the benefits from the network effect is to have only one network supply its services to all users. This would, however, imply that consumers can derive no benefits from competition over price, product or service quality. As an alternative to a single large network, network interconnection among otherwise independent network operators can allow that users enjoy the positive network externalities associated with a single network while benefiting from product diversity in product dimensions other than network size.

Indeed, the principal attraction of the Internet is that because of interconnection among ISPs anyone connected to the Internet is reachable by all other users of the public Internet, irrespective of the home ISPs these users subscribe to. Internet users expect this universal connectivity from their ISP, that is, the ability to reach all destinations reachable on the public Internet. For universal connectivity all networks need to be either directly or indirectly connected to one another. The strong network effects experienced on the retail level of Internet services provision therefore translate into a demand for Internet interconnection by ISPs on the logical layer of Internet service provision. Still, an ISP's incentives for interconnection may be contradictory, when on the one hand an ISP wants to offer universal connectivity to its customers and therefore will seek to interconnect with rival networks, but on the other hand it could try to gain a competitive advantage by refusing to interconnect with some ISPs, thereby keeping them out of the market and trying to lure the customers of these ISPs to its own network instead.

Terms of Interconnection Among ISPs in a Competitive Environment

The interconnection of networks has three aspects. Firstly, a logical interconnection of the networks needs to define which services are to function across the network boundaries and at which quality. Secondly, a physical interconnection between the network infrastructures needs to be established. Lastly, the ISPs need to negotiate how the costs of the physical interconnection and the costs for the traffic transmission via this interconnection ought to be split.

The advantage of the Transmission Control Protocol/Internet Protocol (TCP/IP) standard is that two IP-based networks can agree to use the TCP/IP protocol and thereby define much of what the logical interconnection parameters will be. ISPs can negotiate further quality of service parameters which they want to guarantee across network boundaries. Advanced services, such as real-time Voice over Internet Protocol (VoIP) capabilities or Television over Internet Protocol (IP-TV) services can, for instance, be offered only to users within one and the same network by running additional protocols on top of the standard TCP/IP protocols.[2] They can however, also be offered across network boundaries, if the ISPs agree to guarantee the required quality parameters.

Negotiations over physical interconnection as well as the financial terms of network interconnection need to address the following questions: (1) where to establish the location of the interconnection, (2) how to cover the costs of the network infrastructure which physically connects the two networks and (3) how the two networks ought to split the costs for traffic transmission to and from the other's network. The following subsections present the typical financial agreements for Internet interconnection services today.

Costing and Pricing of Internet Traffic Services

Early interconnection of IP-based networks in the NSFNET era[3] functioned basically without monetary compensation between the connecting parties. The rationale may have been that traffic flows could be expected to be roughly symmetrical. More importantly, however, the funding for the network infrastructure at this time was in most cases provided by the government. Network administrators therefore considered the effort to install complex traffic metering dispensable. This situation changed fundamentally, when the National Science Foundation (NSF) reduced funding and networks had to become self-supporting, this being the case even more so when commercial ISPs took up business. The need arose to recover network costs according

[2] See, for instance, Buccirossi et al. (2005). According to Marcus (2006: 34) these technologies are already widely deployed for controlling the quality of service within networks.

[3] When computer-networking was increasingly used in the 1970s the U.S. National Science Foundation (NSF) played an important role in the development of network interconnection. The NSF initially funded regional networks in the academic community. In 1986, the NSF build the NSFNET, a long-distance network connecting five sites at which NSF funded supercomputers could be accessed. The NSFNET was a network of high-capacity links spanning the entire United Sates and connecting the supercomputer sites (Rogers 1998). This network was open to interconnection by previously existing regional networks in support of research and communication (Jennings et al. 1986). The NSFNET was therefore the first common backbone, or "network of networks".

to some cost-causation principle. It is no coincidence that interconnection agreements changed dramatically at the time of the privatization of the Internet, and that at the same time concerns regarding the possibility of anti-competitive interconnection agreements started to be intensely analyzed by competition authorities and competition economists.

The costs of providing Internet traffic services include the access costs to network resources of the physical layer as well as the costs of switches and routers, the costs for transmission software and the costs for employed staff. These costs are driven by the geographic extent of the network as well as by the bandwidth of the links making up the network.[4] Most of these costs are long-run variable costs. The short-run marginal costs for any particular product or service provided over a given infrastructure are close to zero. As is typical for network services, most of the costs involved in Internet traffic services are also overhead costs, meaning that they cannot be allocated to the incremental costs of particular products and services. The pricing for Internet backbone services therefore necessarily does not reflect short-run marginal costs or even long-run incremental costs of the service.

In general, the price of a particular product must cover at least the long-run incremental costs of this product. If these are not covered then, from an economic point of view, the product should not be produced. In addition, the entire set of products and services offered must cover all overhead costs of production, that is, all costs which cannot be allotted to the incremental costs of a particular product or service. To cover their considerable overhead costs, network operators use pricing strategies that calculate mark-ups on the incremental costs, which allocate the overhead costs to particular products and services according to the price elasticity of demand for these products and services.

The elasticity of demand for Internet backbone services depends on the possibilities for substitution. To offer universal connectivity, a network provider can combine the components (1) own network services, (2) network services from peering partners, and (3) network services from transit partners. These components are interchangeable to a degree and the amount used will depend on the costs of each of these services. With network interconnection, an ISP can avoid building out its own network to particular regions and customer groups, instead profiting from the network investments made by the interconnection partners. The following two subsections look at the pricing of peering and transit interconnection respectively.

The Implicit Price of Peered Interconnection

The main difference between interconnection by a transit contract and interconnection by peering is the degree of coverage of the Internet offered by either transit

[4] Transmission links can be leased. Leased lines are priced by their length and by the capacity of the pipe. The larger the extent of the network, the more switches and routers are needed. The costs for employees also rise with the geographical extent of the network.

(complete coverage) or peering (only the direct customers and transit customers of the peering partner are reached).[5] Furthermore, peering generally involves no monetary compensation for using the peering partner's network while in a transit relationship one party pays the other party for delivery of its data traffic from and to the rest of the Internet.

There is, however, an "implicit price for peered interconnection" (Elixmann and Scanlan 2002: 47), namely the cost of providing the reciprocal service for one's peering partner. In order to understand which interconnection services ISPs consider equal, one must understand how traffic exchange among peering partners is organized. The practice in question tellingly has been called "hot potato routing" (Kende 2000: 5ff.). Peering partners generally interconnect their networks at several dispersed geographic locations. For any data transmission, traffic is passed on to the peering partner at the nearest point of exchange to the origin of the communication.[6] The bits of data are then transported to the receiving user on the receiving user's network.

When the geographic extent of the networks of two ISPs are comparable, and when the end-users connected to the ISPs are similar with respect to the data flows they initiate and receive, then ISP 1 and ISP 2 will carry roughly the same amount of traffic for roughly the same distances as a result of a peering agreement. It is interesting to note, that under these circumstances the number of users connected to the ISPs is irrelevant.[7] If, however, ISP 2 had a network of smaller geographic coverage than ISP 1, then ISP 1 would have to carry the traffic further on its own network before having the opportunity to hand the traffic off to ISP 2. ISP 2 would then profit disproportionately from the peering agreement. Furthermore, if ISP 2's customers had more outbound than inbound traffic flow, for instance if ISP 2 had many content servers on its network which receive only small packages containing content requests but send out large volumes of data, then ISP 1 would carry a larger data volume on its network on the return trip than ISP 2 had carried for the content requests. ISP 1 would then need to invest more into the bandwidth of its network without compensation by ISP 2. Again, ISP 2 would profit disproportionately from a peering agreement.

---

[5] For an overview of transit and peering see also Laffont et al. (2001: 287ff.).

[6] This convention also makes sense, considering that the physical geographic location of the receiving host is known only to the home network of the receiving host.

[7] If ISP 1 had more Internet-users than ISP 2, then traffic flows between the two networks would still be balanced, when the probability of communication between all users is the same, and when the geographic extent of the networks is the same (Economides 2005: 381). Consider, for instance, the following example: Suppose a network with 1,000 attached users interconnects with a network with 100 attached users. If every user corresponds once with every other user, then the smaller network transmits 100 × 1000 contacts to the larger network, amounting to 100,000 contacts. The larger network transmits 1000 × 100 contacts to the smaller network, therefore also 100,000 contacts. Thus, if the data volume that the users send to one another is roughly equal, then the traffic carried by the large and the small network is the same, as long as the types of users are the same across the networks and as long as the operators have networks of similar geographic extent.

These examples illustrate that a change in the relative geographic extent of the networks or in the product portfolio of the peering partners (which would attract different types of customers) can result in an unequal distribution of the advantages from a peering contract and lead the party which profits less by the arrangement to terminate the contract. This shows that the observation that an ISP is terminating peering agreements does not suffice as evidence of anti-competitive behavior. If termination of a contract were not allowed (as some ISPs have demanded from the competition authorities), infrastructure investments would degenerate at the rate at which some ISPs would practice "backbone-free-riding"[8] at the costs of other ISPs. If competition policy forbade positive settlement fees in interconnection contracts, then this would lead to under-investment in network infrastructure (Little and Wright 2000). In conclusion, ISPs will enter into peering agreements only if their prospective peering partners have a network of similar geographic extension and have invested into comparable network bandwidth that can guarantee an equivalent level of quality of service. Furthermore, ISPs generally require traffic flows to be roughly similar. For this it is not important to have the same number of customers, only the same type of customers.

The Price for Transit Interconnection

Transit can be bought from transit givers at any available point of network interconnection with transit fees covering at least the costs of the network resources into which a transit provider has invested to be able to offer transit services and the interconnection fees. In addition, a transit giver will try to cover some of its overhead costs by a mark-up on the incremental costs of providing the transit service. In practice, transit fees are typically two-part tariffs. A flat-fee is charged, which varies depending on the bandwidth of the pipe connecting the two networks and the arranged peak throughput of data on this pipe. A variable fee is charged for traffic in excess of this agreed level, generally charged on a Mbit/s basis. The transit giver therefore has the opportunity to price-differentiate in the market for Internet backbone services. A transit taker will pay a lower average price if more traffic is sent via a particular interconnection and if the amount of traffic sent over this interconnection is correctly predicted beforehand. For inelastic demand, often characterized by a short-term need to shift traffic to a new transit provider, the average price paid will be higher. Yet, such price differences cannot be taken as evidence of significant market power by the transit giver. The need to cover the substantial overhead costs in this market force the transit giver to find ways of implementing surcharges on marginal costs, that can cover the overhead costs of production.

The above analysis shows that a transit interconnection requires far less investments into network infrastructure as well as human resources than peering does.

[8] This term was coined by Baake and Wichmann (1998: 2).

Since a transit contract also offers universal connectivity, whereas peering offers only limited coverage of the Internet, a smaller ISP will often find it less costly to pay for transit services in order to reach universal connectivity than to meet the network requirements necessary to peer with several ISPs of higher hierarchy levels. Peering is therefore not always preferred to transit interconnection, even though it generally involves no monetary compensation for the exchange of traffic. Transit fees are justified by the fact that transit givers invest more into their network infrastructure than transit takers.

Dominance at the Tier-1 Level

The preceding section focused exclusively on the decision on whether to interconnect via a peering or transit agreement. It was shown that the differences in the terms for peering or transit do not necessarily reflect discrimination between ISPs operating on different levels of the network hierarchy. The decision to interconnect either via a peering or a transit agreement is not driven by the number of IP-addresses an interconnection partner offers access to. Rather, factors such as the type of customer mix and the relative geographic extent of the two networks were shown to be important. In contrast, the focus of the following analysis is the decision on whether to interconnect at all. In the course of this decision the network reach provided by a potential interconnection partner is of fundamental importance, because the ultimate goal of network interconnection is to provide universal connectivity. All ISPs not active on the highest level of the Internet hierarchy need at least one transit agreement with a Tier-1 ISP or with an ISP that has such a transit interconnection. Therefore the question arises whether it is likely that a merger on the Tier-1 level of the Internet hierarchy could negatively impact competition in Internet backbone services in the sense that a Tier-1 ISP may have an incentive to discriminate lower level ISPs.

As was discussed above, the demand for Internet backbone services on the logical layer of Internet service provision is a derived demand from the end-user demand for universal connectivity on the retail level of Internet service provision. In the retail market, universal connectivity signifies that all other end-users and content providers on the Internet can be reached via one's home ISP. In the Internet backbone services market, universal connectivity signifies that an ISP can send and receive data to and from all IP-addresses allocated to public uses in the Internet.

The literature on Internet backbone services does not differentiate clearly between universal connectivity on the applications layer and universal connectivity on the logical layer of Internet service provision. The difference is, however, of importance when, as is often the case, the number of "customers" attached to ISPs is used as the measure for the Internet coverage the ISP provides. This is a concept relevant on the applications layer of Internet service provision. On the logical layer a customer of an ISP can, however, be either an end-user, representing only one of millions of Internet-Protocol addresses (IP-addresses) or another ISP, representing an important fraction of all registered IP-addresses. For the purposes of measuring Internet coverage on the logical layer of Internet service provision it is therefore more meaningful to speak of the coverage of IP-addresses which this ISP can offer as a peering partner. Transit services, by definition, offer universal connectivity.

Economists have developed models that try to capture the interconnection incentives of ISPs. Theoretical models are of particular relevance in the context of merger policy because competition authorities cannot look at actual market conduct for their analysis. Policy makers depend on predictions derived from economic modeling to understand whether efficiency considerations or attempted exclusionary conduct are at the core of proposed mergers. The model that was influential in the merger proceedings surrounding the MCI and Worldcom merger in 1998 and the attempted merger of the resulting firm MCI/Worldcom and Sprint in 2000 offered initial interesting insights into the interconnection incentives of ISPs with asymmetric installed customer bases. Since then, the literature on interconnection incentives of ISPs has refined this model considerably. The following two subsections shall review the theoretical debate on the interconnection incentives of ISPs in more detail.

The Crémer, Rey and Tirole Model

The reasoning that led the competition authorities to impose severe conditions on the merger of MCI and Worldcom in 1998[9] was based to a great extent on one of the earliest theoretical models, which tried to capture the strategic interconnection decision of ISPs. From this model by Crémer et al. (2000) the conclusion was drawn that an ISP that is dominant in terms of attached customer base in the retail market, would have the means to dominate the market for Internet backbone services. It would either refuse to peer with smaller rivals or price-squeeze them out of the market (Kende 2000: 22–23).[10]

The model by Crémer, Rey and Tirole builds on the Katz and Shapiro (1985) model of network externalities. As Katz and Shapiro, Crémer, Rey and Tirole model the number of firms in the market as exogenously given and assume that there is no product differentiation. Consumers exhibit different basic willingness to pay for the service but show no technology preferences and express the same evaluation of the network effect.

In a first scenario, Crémer, Rey and Tirole focus on interconnection decisions in an asymmetric duopoly situation. The existing users of the two networks are assumed to be locked-in. The networks compete à la Cournot over the addition of new customers to their networks. The choice of the quality of interconnection between the networks is introduced as a strategic variable. In the first stage of the game the

[9] MCI had to divest it's Internet operations before a merger with Worldcom was approved (European Commission 1998).

[10] Crémer, Rey and Tirole argue that a customer in this model can be either an end-user or an ISP. They do not differentiate between the two.

quality of interconnection is determined by the network which sets the lower quality level. Given the interconnection quality, the networks then choose their capacity and prices. In equilibrium, the network with the larger installed customer base prefers a lower level of interconnectivity than the smaller rival because it can expect to dominate the market for new customers.

Two effects determine the equilibrium outcome. Firstly, lower connectivity levels lead to an overall demand reduction in the market, which negatively impacts all firms. Secondly, reduced interconnectivity introduces an element of quality differentiation between the firms, which in this model, can only differentiate among themselves along the dimension of network size. The network with the initially larger locked-in customer base profits from this quality-differentiation effect because it can offer more benefits from network externalities to new users. The bigger network trades off the negative effect of the demand reduction against the positive effect of the quality differentiation. The incentive to choose a lower level of interconnection quality is the more positive the stronger the network externalities are and the greater the difference in installed bases is. A differential analysis shows that the incentive to increase the level of interconnection quality may rise when the number of locked-in customers is already very large, because then the demand expansion effect triggered by a larger network becomes so important, that good quality interconnection is preferred. This equilibrium solution to the model has been the basis for arguing that a dominant Tier-1 ISP would have an incentive to refuse or degrade interconnection with rivals, especially in dynamic markets with high growth potential.

In a second scenario, Crémer, Rey and Tirole (ibid. 456ff.) analyze a market initially consisting of four equal sized ISPs. As long as all four have the same size, all are interested in a good quality interconnection because all profit equally from a demand expansion effect. The elicitor of a quality degradation would suffer the same negative demand reduction as its three rivals without a compensatory gain from a positive quality-differentiation effect. The authors then show how the incentives to interconnect change when two of the ISPs merge and the resulting market of three ISPs then includes one firm with an installed base of at least the size of the combined installed bases of the other two firms. In this scenario the largest firm is generally not interested in deteriorating the quality of interconnection with both of the rival networks, although, in some circumstances, it can profit from a targeted degradation strategy, in which it refuses good quality interconnection with one of the smaller rivals while it continues good quality interconnection with the other rival. This conclusion depends on the non-targeted firm not offering transit services to the targeted firm.[11] The positive quality-differentiation effect will then result in the targeted firm not attracting any new customers while the dominant firm and the non-targeted firm gain more customers (even though the non-targeted rival profits more from the quality-differentiation

[11] Crémer, Rey and Tirole (ibid. 458) argue that the dominant firm can limit the capacity of the interface with the non-targeted network to such an extent that the capacity is only sufficient to provide good quality interconnection for the traffic of the non-targeted network but would result in very bad interconnection quality if the traffic should grow to encompass also the traffic of the targeted network.

effect). It was especially this result that competition authorities relied upon in their decision on the merger by MCI and Worldcom in 1998.

Critique of the Crémer, Rey and Tirole Model and Alternative Modeling

The results of the model by Crémer, Rey and Tirole depend critically on the additional assumptions besides the network effects included in the modeling set-up. It is these assumptions which lead to the result that the largest firm prefers a lower level of interconnection quality compared to its smaller rivals. Below it is discussed whether these assumptions are relevant for the market for Internet backbone services.

Market Entry Conditions

First, consider the assumption of a fixed number of firms in the market. This assumption does not correspond well to the thousands of active ISPs observable in reality. If at all, then this assumption may apply to the market for Tier-1 ISP services in which only five to ten ISPs are active. But whether this market has structural barriers to entry, which justify the assumption of a fixed number of firms, is what is trying to be proved. To start with this assumption distorts the analysis of the effects of network externalities on competition in this market.

It can be shown that the equilibrium results of the model by Crémer, Rey and Tirole change dramatically when the number of firms in the market is endogenized (Malueg and Schwartz 2006). Consumers do not necessarily choose the firm with the initially larger installed base. When this firm chooses not to be compatible with it's smaller rivals,[12] and when smaller rivals in sum have a minimum initial market share and choose to remain compatible among themselves, then, for a large set of parameter values, new consumers will sign on to this network of smaller compatible firms in the expectation that in a dynamic market setting this network will eventually incorporate more contacts than the single-firm network of the initial market leader.[13] If payments for interconnection were introduced, the parameter values for which the initially larger

[12] The targeted degradation scenario is not considered by Malueg and Schwartz. In a related working paper (Malueg and Schwartz 2002: 37) the authors argue that the parameter values that make targeted degradation profitable to the dominant firm imply unrealistic values for price relative to marginal cost and the consumer surplus of the median subscriber.

[13] Even when the dominant network's installed customer base is larger than the combined installed customer bases of its rivals, there are parameter regions in which the rivals will be more successful in adding new customers to their networks (Malueg and Schwartz 2006: 9). This is due to the customer's expectations of market evolvement in dynamic market settings, in which networks are expected to have a high growth potential. This conclusion is comparable to the results by Economides (1996) for a monopolist that prefers inviting market entry.

firm would choose autarky would be even more limited because smaller firms could share their gain from increased connectivity by offering payments to the larger firm.

That the smaller rivals will remain compatible amongst one another and will have a significant network reach through the interconnection of their networks is very realistic for the Internet backbone services market. The presence of many ISPs at Internet exchange points and the availability of standardized contracts together with the fact that market conditions for transit services are transparent facilitate interconnection agreements. The subscribers of the interconnected networks on the lower hierarchy levels can reach all users of these networks. Considering that many subscribers of Internet services are multi-homed (i.e. subscribe to several networks) and that all those customers of the dominant firm that are multi-homed can be reached via an alternative network, it becomes clear that the Internet reach provided to the customers of the lower-level ISPs can be increased significantly by coordination on the lower hierarchy levels.

Product Differentiation

Secondly, consider the assumption that customers do not have individual preferences for technology characteristics of the network they subscribe to. This assumption does not correspond well to the reality of a large degree of product differentiation observable among ISPs. On the Internet backbone services market, ISPs offer their services to other ISPs, to web-hosting services, to large business users and to private end-users. They offer different service levels according to their customers' needs and they offer their services at diverse locations, again according to their customers' needs. An ISP that would hope to make the market tip in its favor would have to cater to all customers in the market. This may not be the most profitable market strategy in a world of customer heterogeneity. ISPs that focus on particular customer groups have comparative advantages in supplying the types of services that these customers prefer. In this case, the proper theoretical reference model may be that ISPs are supplying components of systems rather than competing systems. In such markets, compatible products (as, for instance, interconnected networks) cater to the needs of particular customers. Competition between the products is not as strong as in a market of competing systems because the possibility to make profits is often increased by compatibility (see Economides 1989; Einhorn 1992). When product differentiation is introduced into the model by Crémer, Rey and Tirole it can be shown that in any shared market equilibrium both firms profit from a higher interconnection quality because competition becomes less aggressive when the firms can offer the same positive network externality effect to their customers (Foros and Hansen 2001).[14] General analysis on the compatibility incentives of providers of differentiated network goods come to comparable results (Doganoglu and Wright 2006).

[14] In this model there is also no installed customer base. This fact of course also has an important impact on the results of the model. This aspect is in the focus of a model structure by Economides (2005) which is discussed below.

Switching-Costs

There are other critical assumptions in the Crémer, Rey and Tirole model which do not correspond to the characteristics of the Internet backbone services market. Firstly, consider the assumption that installed bases are locked-in. In reality, switching ISPs is not difficult for end-users or ISPs. Only the cancellation period of their contract may delay the reaction for some weeks. Larger customers such as firms and ISPs are often multi-homed, that is, they connect to more than one ISP at any given time. This is important for the ISP to be able to guarantee its contractual service level vis-à-vis its customers. It is also a signal that traffic can be diverted quickly from one ISP to another without large transaction costs involved. The fact that switching is relatively easy increases the competition between Internet backbone service providers. When the assumption of a locked-in customer base is relaxed, it can be shown that the initially dominant network has an incentive to keep up a high quality of interconnection (Economides 2005, Appendix). A degradation of interconnection quality with one of the smaller rivals would lead to a loss of universal connectivity that would result in a severe demand response by the installed customer base as well and therefore to revenue and profit loss.

Collusion on the Tier-1 Level

Only Tier-1 ISPs can guarantee universal connectivity without relying on a transit offer. The preceding section showed that one Tier-1 alone cannot successfully refuse interconnection with other ISPs or raise interconnection prices in the hopes of ousting competitors from the market. The transit offers of Tier-1 ISPs are perfect substitutes. Absent any collusive practices there is intense competition in this market. This fact provides the Tier-1 ISPs with a motive to collude on the market for transit services. If all Tier-1 ISPs acted simultaneously in increasing prices for transit services, then lower level ISPs would have no alternative transit provider from whom to buy universal connectivity services. And no new provider of universal connectivity could enter the market as long as the Tier-1 ISPs would successfully foreclose this market by not entering into any new peering agreements. The question analyzed in the present section is whether Tier-1 ISPs can organize a stable collusion in the wholesale market in order to collectively raise the price of transit services?

There is a literature on two-way access in telecommunications markets which analyzes whether cooperation on the wholesale level can help enforce collusion on the retail level.[15] A two-way access scenario is given when customers connect to only one network, such that the two networks reciprocally need access to each other's customers on the wholesale level. Termination in this scenario is comparable to a monopolistic bottleneck. The application of this literature has mostly been to

[15] The seminal articles in this research are Laffont et al. (1998a, b).

voice telephony markets, for instance, mobile telephony or reciprocal international termination. Considering that ISPs have a termination monopoly whenever customers exclusively connect to their network only, the models may, however, also be applicable to the market for Internet backbone services. If a large fraction of end-users are connected to only one network, then ISPs may have the possibility to collude on the retail market.

The assumptions that are necessary for successful collusion in a market with reciprocal termination are:

- There is no free market entry.
- There are no capacity limitations.
- Every customer connects to only one network.
- The calling party pays for the connection. The receiving party does not care about the price the caller has to pay to reach him.
- There is no price differential for calling customers on the same network (on-net calls) or on another network (off-net calls).
- Access charges for call termination are set reciprocally.
- Both networks have the same costs of production.
- The probability of a call is independent of the home-network of the calling parties. This implies that given same marginal prices for on-net and off-net calls, the share of calls that originates with network 1 and terminates with network 2 will be equivalent to the market share of network 2.[16]

It can be shown, that when the reciprocal access charge set by the firms is not too high compared to the marginal costs of termination, and when the substitutability between the two networks is also not too high, then there exists a unique equilibrium to this model (Laffont et al. 1998a: 10). In this equilibrium, the retail price is increasing in the access charge for termination. The firms can therefore use the access charge to enforce a higher retail price than would be the outcome of competition.

The intuition behind this result is that if access charges are set at the level of the actual marginal costs of terminating a call, then the marginal costs of producing either an on-net call or an off-net call are the same for the originating network. If the access charges are above the marginal costs of termination, then the costs of producing an off-net call are higher than those of producing an on-net call. The higher the access charge, the higher the marginal costs of producing an off-net call. This mechanism can be used to raise the rival's costs of production and put pressure on retail prices.

For the collusion to be stable the access charge must not be too far above the marginal costs of termination and the substitutability between the networks must not be too high. When the access charge is set well above the marginal costs of termination, a firm has an incentive to increase its market share and avoid paying

[16] This so-called "balanced calling pattern assumption" has important implications for the model. It implies that "…for equal marginal prices, flows in and out of the network are balanced – even if market shares are not." (Laffont et al. 1998a: 3). When wholesale access charges are set reciprocally, this assumption implies that the wholesale interconnection payments cancel each other out.

termination fees.[17] When the substitutability between the networks is high, attempts to increase one's own market share by luring the customers of the other network to switch networks will more likely be successful.

The incentives to compete rather than collude in the retail market are further intensified by allowing for more complex price structures in the retail market besides identical linear prices for on-net and off-net calls. Firstly, consider the possibilities offered by non-linear pricing structures. When charging two-part tariffs the firms can use a lower fixed fee to increase market share while keeping the unit price on the collusive level such as not to induce a quantity expansion effect. As a result of the higher market share, the firm will have less off-net traffic and less termination charges to be paid. With non-linear pricing in the retail market competition is intensified and collusion, again, becomes more difficult (Laffont et al. 1998a: 20ff.). Secondly, consider price discrimination in the retail market. In a companion article, Laffont, Rey and Tirole show that collusion is destabilized when retail prices differentiate between on-net and off-net calls (Laffont et al. 1998b). A defecting firm can use low on-net prices to increase its market share but keep off-net prices on the collusion level such as not to induce a quantity expansion effect which could produce an access-deficit.

Application to the Market for Internet Backbone Services

The model above shows that while collusion via wholesale access charges is possible it is only stable under very restrictive assumptions. Given this information, what can be learned with respect to the market for top-tier Internet backbone services? Is it likely that Tier-1 ISPs can use their wholesale agreements to stabilize higher transit prices?

Some of the assumptions of the model set-up by Laffont, Rey and Tirole fit relatively well with the market characteristics of the Internet backbone services market, at least when only the highest level of the Internet hierarchy is in focus. For instance, for Internet interconnection via peering it is true that there is no price differential between on-net and off-net connections. Furthermore, Tier-1 ISPs, as peering partners, generally set their access charges reciprocally (albeit to the level of zero). Tier-1 ISPs can also be considered to have a similar cost-structure for terminating each others connections. Lastly, the assumption of a balanced calling pattern between Tier-1 ISPs is fitting, given that they are peering partners and can therefore be assumed to have a similar customer structure.

Other assumptions of the model by Laffont, Rey and Tirole, however, do not correspond as well to the market for Internet backbone services on the highest hierarchy level. As these assumptions are essential to the stability of the collusion

[17] Even when the net payments between the two networks are zero with reciprocal access charges and balanced calling patterns, they perceive the access charge as a marginal cost of production and will want to avoid them.

equilibrium the fact that they do not correspond to the market in question is an indication that collusion in the market for top-tier transit services is difficult to maintain. Firstly, consider the assumption that every customer is connected to only one network as prerequisite for the termination monopoly. This assumption is too strong for the market for Internet backbone services, as many small ISPs and many business customers are multi-homed. Therefore, the termination monopoly in Internet interconnection is not as stable as assumed in the model by Laffont, Rey and Tirole. Next, consider the number of players in the market. It can be argued that market entry into Tier-1 Internet service provision is not free because any new entrant must reach a peering agreement with all other Tier-1 ISPs. None the less, there are already several active firms on the Tier-1 level of Internet backbone services which increases the number of potential substitutes and destabilizes any collusive agreement.

Furthermore, the assumption that the receiving party of a connection does not care about the costs the calling party has to pay for the connection is not appropriate in the context of Internet interconnection. Businesses offering content and information on the Internet care very much about the costs their targeted customers face for reaching this content. The costs of being reached are a significant factor in their decision where to place their content on the Internet. The access charge is therefore not only indirectly but also directly a strategic element in the competition over end-users.

Decisive for the stability of any collusion are the level of the access charge and the substitutability of the network offers. Between Tier-1 ISPs the access charge is generally set to the level of zero. It therefore corresponds to the prerequisite that is should not be too far above marginal costs of termination. However, for collusive purposes a termination fee would need to be introduced where there was none before. This may be more difficult than an incremental increase of an existing termination charge. Furthermore, the degree of substitutability between the transit offers of Tier-1 ISPs can be considered to be very high. This fact makes collusion interesting, but at the same time it represents a high risk of instability of any collusion because any of the Tier-1 ISPs could hope to increase its market share by offering a lower transit charge than its competitors.

Lastly, consider the price structures in the market for transit services provided by Tier-1 ISPs. Transit prices generally are not differentiated according to the destination network. However, non-linear prices for transit services are the norm in the transit market. In general a transit taker will pay a fixed fee that depends on the bandwidth by which the two networks are connected plus a variable fee for traffic exceeding a previously defined threshold. The ability to compete in two-part tariffs is a further hindrance to stable collusion in the transit market. To summarize, the prerequisites for a stable collusion are not fulfilled in the market for Tier-1 backbone services.

Conclusions

The purpose of this paper was to analyze the strong network externalities associated with Internet services from a competition-policy perspective. It was argued that in the market for Internet services network effects are so important that an ISP needs to be able to offer universal connectivity in order to survive in this market. The demand for universal connectivity on the logical layer is a derived demand from the demand for universal connectivity on the applications layer. To reach universal connectivity, new entrants to the Internet backbone services market will need to establish a direct or indirect transit agreement with at least one Tier-1 ISP.

Tier-1 ISPs enter into peering agreements only when the benefits from the interconnection are roughly similar to both parties. The fear that a single Tier-1 ISP could be able to abuse a dominant market position in a transit agreement with lower-level ISPs was not substantiated by the analysis. Competitive forces in the market for top-tier Internet backbone services are strong. Tier-1 ISPs compete with product differentiation tactics. Customers frequently multi-home and can relatively conveniently switch their home network. As a result, Tier-1 ISPs cannot benefit from refusing to interconnect respectively from deteriorating interconnection quality with lower-level networks.

In principle, some market constellations are conducive to collusion on the retail level, stabilized via cooperation on the wholesale level. A collusion between Tier-1 ISPs to collectively raise prices in the transit market is not likely to be stable because the prerequisites for a stable collusion are not fulfilled in the market for top-tier Internet backbone services. Most importantly, the assumption of a termination monopoly is not fulfilled. To summarize, the discussion in this paper has provided strong support that competitive forces in the transit market are working and can effectively hinder Tier-1 ISPs to discriminate ISPs on lower levels of the Internet hierarchy.

References

Baake P, Wichmann T (1998) On the economics of internet peering. Netnomics 1:89–105

Blankart C, Knieps G (1992) Netzökonomik. Jahrbücher für neue politische Ökonomie 11:73–87

Buccirossi P, Ferrari Bravo L, Siciliani P (2005) Competition in the internet backbone market. World Competition 28(2):235–254

Crémer J, Rey P, Tirole J (2000) Connectivity in the commercial internet market. J Ind Econ XLVIII:433–472

Doganoglu T, Wright J (2006) Multihoming and compatibility. Int J Ind Organ 24:45–67

Economides N (1989) Desirability of compatibility in the absence of network externalities. Am Econ Rev 79:1165–1181

Economides N (1996) Network externalities, complementarities, and invitations to enter. Eur J Polit Econ 12:211–233

Economides, N (2005) The economics of the internet backbone. Majumdar S et al. Handbook of Telecommunications Economics, Vol. 2. Amsterdam: North Holland

Einhorn M (1992) Mix and match: compatibility with vertical product dimensions. RAND J Econ 23: 535–547

Elixmann D, Scanlan M (2002) The Economics of IP Networks-Market, Technical and Public Policy Issues Relating to Internet Traffic Exchange. wik-Consult Final Report, Bad Honnef

European Commission (1998) Commission decision of 8 July 1998 declaring a concentration to be compatible with the common market and the functioning of the EEA agreement (Case IV/M.1069 '– WorldCom/MCI). Official Journal of the European Commission L116:1–35

European Commission (2000) Commission decision of 28 June 2000 declaring a concentration incompatible with the common market and the EEA agreement (Case COMP/M.1741 – MCI WorldCom/Sprint)
Foros O, Hansen B (2001) Competition and Compatibility Among Internet Service Providers. Info Econ Pol 13(4):411–425
Jennings M, Landweber LH, Fuchs IH, Farber DJ, Adrion WR (1986) Computer networking for scientists. Science 231:943–950
Katz M, Shapiro C (1985) Network externalities, competition, and compatibility. Am Econ Rev 75:424–440
Kende M (2000) The Digital Handshake: Connecting Internet Backbones. OPP Working Paper 32. FCC, Washington, DC
Knieps G (1997) Phasing out sector-specific regulation in competitive telecommunications. Kyklos 50:325–339
Knieps G (2006) Sector-specific market power regulation versus general competition law: Criteria for judging competitive versus regulated markets. In: Sioshansi FP and Pfaffenberger W (eds.), Electricity Market Reform: An International Perspective. Amsterdam: Elsevier
Laffont JJ, Rey P, Tirole J (1998a) Network competition: I. Overview and nondiscriminatory pricing. RAND J Econ 29:1–37
Laffont JJ, Rey P, Tirole J (1998b) Network competition: II. Price discrimination. RAND J Econ 29:38–56
Little I, Wright J (2000) Peering and settlements in the internet: An economic analysis. J Regul Econ 18:151–173
Malueg D, Schwartz M (2002) Interconnection Incentives of a Large Network. Georgetown University Department of Economics Working Paper 01-05, August 2001, revised January 2002
Malueg D, Schwartz M (2006) Compatibility incentives of a large network facing multiple rivals. J Ind Econ 54:527–567
Marcus S (2006) Framework for Interconnection of IP-Based Networks: Accounting Systems and Interconnection Regimes in the USA and the UK. wik-Consult Report, Bad Honnef
Rogers JD (1998) Internetworking and the politics of science: NSFNET in internet history. Inf Soc 14(3):213–228
Vanberg M (2009) Competition and Cooperation among Internet Service Providers: A Network Economic Analysis. Baden-Baden: Nomos

http://www.springer.com/978-3-7908-2081-2

Telecommunication Markets
Drivers and Impediments
Preissl, B.; Haucap, J.; Curwen, P. (Eds.)
2009, XIV, 526 p. 90 illus., Hardcover
ISBN: 978-3-7908-2081-2
A product of Physica-Verlag Heidelberg