Working Paper No 43/02

Competition and regulation strategies in the Internet

by

Øystein Foros
Hans Jarle Kind

SNF project no 1125
Distrikts- og konkurransepolitiske utfordringer knyttet til IKT
(Information and communication technology: Challenges
for regional and competition policies)

The project is financed by the Research Council of Norway

SIØS - Centre for International Economics and Shipping

INSTITUTE FOR RESEARCH IN ECONOMICS AND BUSINESS ADMINISTRATION
BERGEN, SEPTEMBER 2002
ISSN 1503 - 2140

SIØS – Centre for international economics and shipping – is a joint centre for The Norwegian School of Economics and Business Administration (NHH) and The Institute for Research in Economics and Business Administration (SNF). The centre is responsible for research and teaching within the fields of international trade and shipping.

International Trade
The centre works with all types of issues related to international trade and shipping, and has particular expertise in the areas of international real economics (trade, factor mobility, economic integration and industrial policy), international macroeconomics and international tax policy. Research at the centre has in general been dominated by projects aiming to provide increased insight into global, structural issues and the effect of regional economic integration. However, the researchers at the centre also participate actively in projects relating to public economics, industrial policy and competition policy.

International Transport
International transport is another central area of research at the centre. Within this field, studies of the competition between different modes of transport in Europe and the possibilities of increasing sea transport with a view to easing the pressure on the land based transport network on the Continent have been central.

Maritime Research
One of the main tasks of the centre is to act as a link between the maritime industry and the research environment at SNF and NHH. A series of projects that are financed by the Norwegian Shipowners Association and aimed directly at shipowning firms and other maritime companies have been conducted at the centre. These projects include studies of Norwegian shipowners' multinational activities, shipbuilding in Northern Europe and the competition in the ferry markets.

Human Resources
The centre's human resources include researchers at SNF and affiliated professors at NHH as well as leading international economists who are affiliated to the centre through long-term relations. During the last few years the centre has produced five PhDs within international economics and shipping.

Networks
The centre is involved in several major EU projects and collaborates with central research and educational institutions all over Europe. There is particularly close contact with London School of Economics, University of Glasgow, The Graduate Institute of International Studies in Geneva and The Research Institute of Industrial Economics (IUI) in Stockholm. The staff members participate in international research networks, including Centre for Economic Policy Research (CEPR), London and International Association of Maritime Economists (IAME).

Competition and regulation strategies in the Internet

By

Øystein Foros and Hans Jarle Kind

Abstract:
The purpose of this article is to provide a simple introduction to the Internet's value system and historical development from an economic point of view. One of the central themes that we discuss is whether increased user and service heterogeneity requires new allocation mechanisms to secure an efficient utilization of the Internet's capacity. Thereafter we discuss whether dominating network firms may have incentives to foreclose smaller rivals that operate at the same level of the hierarchy, and whether vertical integration may imply that upstream firms in control of essential inputs find it optimal to foreclose competitors in the downstream market. Finally, we argue that the growth of the Internet may require changes in the regulation of the telecommunication sector.

1 Introduction[1]

In the telecommunications and the Internet industry we see that vertical and horizontal convergence are key features of the market development. Vertical convergence implies that the borders between telecommunication, media and IT are eroding. Horizontal convergence implies that networks or platforms that were earlier limited to distribute one given type of services, now can be used as a distribution channel for several services. In the present paper we discuss the following topics:

- Whether increased user and service heterogeneity requires new allocation mechanisms to secure an efficient utilisation of the Internet's capacity.
- Whether dominating network firms may have incentives to foreclose smaller rivals that operate at the same level of the hierarchy.
- Whether vertical integration may imply that upstream firms in control of essential inputs find it optimal to foreclose competitors in the downstream market.
- Whether the growth of the Internet may require changes in the regulation of the telecommunication sector.

2 The history and characteristics of the Internet

In the early 1960s the National Science Foundation (NSF) of the USA initiated the development of the technology and infrastructure behind what we today know as the Internet. As a consequence of this effort, some leading academic institutions in the USA became interconnected through an electronic communication network (NSFNET) in 1986. The NSFNET communication technology, which was invented by the American Ministry of Defence, was based on a so-called Internet Protocol (IP) that has become the standard for distribution of data bits from sender to receiver. At the end of the 1980s commercial firms like IBM and MCI wanted to connect to the

[1] This article is partly based on Foros (2002).

Internet, and in 1993 NSF developed a plan for commercialisation and privatisation of the Internet. Two years later NSF withdrew from NSFNET.

In the early years of the Internet both the users and the services were relatively homogenous. The majority of the users were found at universities and research institutions, and the dominating services were transfers of data files and electronic mail. A common denominator for this kind of users and the applications is that they are relatively “patient” with regard to delays. First, these user groups typically have a relatively low willingness to pay in terms of money compared to time. What we mean by this is that students and researchers in many cases are more likely to accept a delay than to pay a few dollars for an immediate transfer of a data file. Second, services like transfers of data files and e-mails are intrinsically insensitive to delays, since they typically do not require any active real-time cooperation between sender and receiver.

It is a general trend that a large share of new user groups and new applications are more impatient or sensitive to delays than what was the case earlier. New users in the private business sector regularly prefer to pay money in order to progress in the queue rather than to wait. Moreover, we have recently observed a large growth in the number of interactive real-time applications. Examples of such applications are interactive video and telephony over the Internet. The required transfer capacity also varies a lot. World wide Web (www) and real-time video require significantly higher transfer capacity than, for instance, purely text based electronic mail.

The present Internet architecture is based on connectionless packet switching (see below), where data packets are served according to the first come, first served principle. This architecture is not particularly appropriate for serving impatient users or for handling real-time applications. Unless price signals can be used to sort and segment users, it will probably become increasingly difficult to offer reail-time applications and to serve impatient users over the open Internet. Thereby impatient users and time sensitive applications may *de facto* be excluded from the open Internet This may lead to a process where the Internet becomes segmented into several independent networks instead of a process of further convergence.[2]

[2]A discussion of the development and history of the Internet is offered by Mackie-Mason and Varian (1997) and Werbach (1997), while Cave and Mason (2001) give an overview of the Internet with a focus on regulation and the competitive environment.

2.1 Layered and hierarchical structure

Within telecommunications there has traditionally been a close connection between services and the underlying distribution system. Introduction of new services typically requires modifications of the infrastructure, for instance through upgrading of the software in the networks' switches. The basic principle in the Internet is different, since there is a clear separation between the underlying infrastructure, applications and content. Common protocols between the basic infrastructure and applications imply that it is not necessary to change the infrastructure when new applications are introduced. This has made it very simple to introduce new applications and services on the Internet, and this has presumably been a central factor behind the success of the Internet.

The Internet is often described as having a layered network structure as described in Figure 1. As an illustration of the importance of this structure, it should be noted that the present killer-applications in the Internet, like www, were developed long after the underlying IP technology.

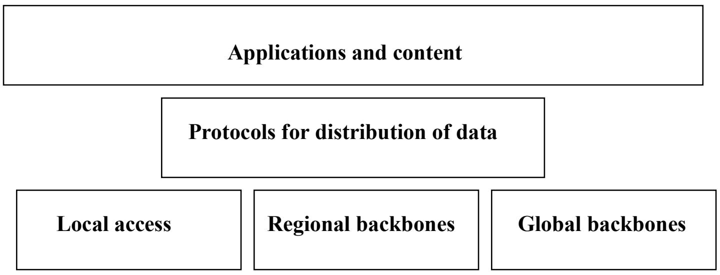

Figure 1: The layered structure of the Internet.

In the bottom layer of the Internet structure we have the physical infrastructure, where local access is an essential component. It should be noted, though, that the total quality of the infrastructure or distribution system does not depend on the quality of local access alone. For instance, there is little reason to upgrade local access to handle broadband applications if the quality of the regional and global backbones implies that the speed of data transfers over the Internet will not increase. A chain is not stronger than its weakest link, and local access is only one out of several components of the distribution system that must be upgraded in order to get high-speed Internet.

In the higher layers of the Internet structure we find applications and content. In traditional telecommunications it is difficult to offer new applications unless one has a close relationship to the firms that control the underlying infrastructure. An example of this is the introduction of a service like number identification. A few years ago it became technologically possible even for users of analogue telephone systems (PSTN) to identify the caller on a display. However, it was not possible to implement number identification unless the local telephone infrastructure was upgraded. Thus, the service could not be introduced independent of the firm that controls the local access network – and this is typically the telecommunications incumbent, like British Telecom in the UK and France Telécom in France. As mentioned above, the situation is completely different within the Internet, since the common protocols between infrastructure and applications/content make it possible to offer new network services independent of the firms that control the infrastructure. As Shapiro and Varian (1998a) put it: "Any idiot can establish a Web presence – and lots of them have."

Since the layered Internet structure means that anyone can introduce new services, incumbents like British Telecom and France Telécom may lose much of the control that they have had within telecommunications. This is one reason why the dominating position of the established telecommunication firms may erode over time.

2.2 Economies of scale and scope

In economic theory it is common to distinguish between increasing returns on the supply side and on the demand side, as illustrated in Figure 2. There are increasing returns on the supply side if it is cost efficient to produce several different products within one and the same firm (economies of scope) or if the average costs fall when each single product is produced in large series (economies of scale). Due to the fact that it has been difficult to separate services from the underlying infrastructure, economies of scope has historically been important in telecommunications. This kind of economies of scope is presumably smaller within the Internet, due to the properties of the layered Internet structure. However, there may exist other kinds of economies of scope, e.g., due to technological convergence that makes it possible to offer services over a common distribution platform rather than through separate and unrelated networks.

While it is uncertain how important the economies of scope on the supply side are, there is little doubt that the economies of scale are significant. This is true with respect to investments both in infrastructure and in development of new services. For instance, there are large fixed costs involved in development of the prototype of a new software, while the subsequent copies are almost costless to produce. The cost structure for content and applications may therefore be similar to the one we have for the physical network infrastructure.

Economies of scope on the demand side, which is placed in the lower right-hand corner of Figure 2, is commonly described by the term complementarity. Two goods or events are complements if they mutually reinforce each other. If increased sales of component A increase the sales of component B, which in turn increase the sales of component A, these two goods are complements. For our purpose we can say that two goods are complements if a price reduction or a quality improvement on one of them increases the demand for both goods. A lower price on Internet browsers, for instance, is likely to stimulate sales of operative systems, and vice versa. This is an insight that Microsoft has taken advantage of.

In this section we will focus on economies of scale on the demand side, which is placed at the lower left-hand corner of Figure 1. This kind of increasing returns is commonly named as *network effects*, and takes place when the unit value of a product or a system is increasing in the number of users.[3]

	Scale	Scope
Supply side	Decreasing average costs	Gains from joint production
Demand side	Network effects	Complementarity

Figure 2: Increasing returns on the supply and demand side.

Let us illustrate the implications of network effects with the following example, which is based on Shapiro and Varian (1998b). Suppose that there are 1000 persons in the market for a given service, and let v be the reservation price for person v, where $v=1,\ldots,1000$. The price is p, and the number of users that value the good at price higher than p is $1000-p$. In a traditional market we will then have a downward-sloping demand curve as in the left-hand side panel of Figure 3. If the good is supplied in a competitive market with constant marginal costs equal to c, we will have a unique equilibrium with $p=c$ and quantity equal to $n=\hat{n}$. Suppose, however, that we

[3] It is common to distinguish between direct and indirect network effects. For instance, we have direct network effects between owners of telephones; the more people that have installed a telephone, the greater its value (this is also labeled a *real* network). An example of indirect network effects is that between users of PCs; a large number of PC users implies that there will be a large demand for PC compatible software. This in turn tends to generate a large variety of PC software, which increases the user value of the PCs (this is an illustration of a *virtual* network).

consider network services like telephony, text messages or e-mail, where it is reasonable to assume that the value for each user increases in the number of other users. The easiest way to incorporate this property is to assume that the value for person *v* of the service is *vn*, where *n* is the number of users. This will qualitatively change the demand curve. By combining *p=vn* and *n=1000-v* we find in fact that the demand curve can be written as

$$p = n(1000 - n).$$

The right hand panel of Figure 3 illustrates this demand curve graphically, and we see that it has a shape that is fundamentally different from the traditional demand curve, since it is at first upward sloping.

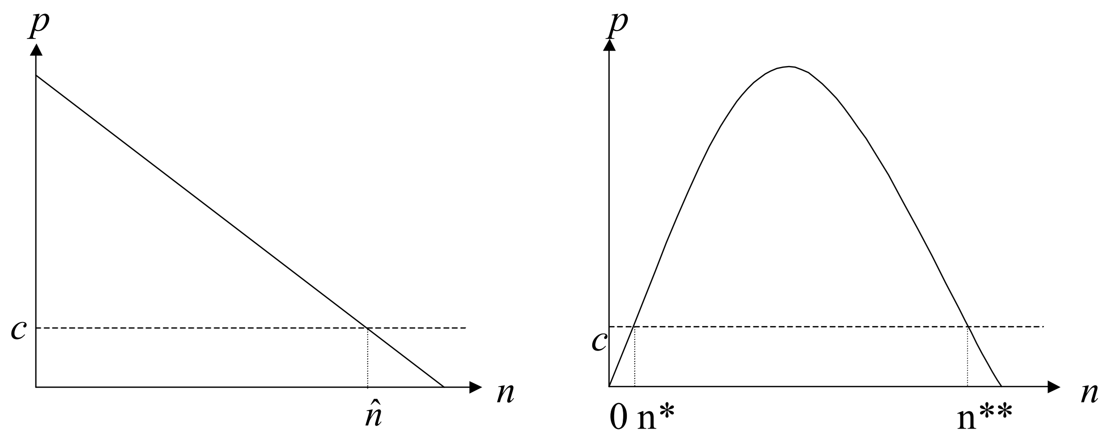

Figure 3: Demand in a traditional market (left-hand side panel) and in a market with network externalities (right-hand side panel).

The intuition behind the shape of the demand curve in the right-hand panel of Figure 3 is as follows. Other things equal, the first few consumers that possibly connect to the network have a low willingness to pay. The reason for this is simply that they have few people to communicate with. However, the willingness to pay increases as more consumers connect to the network. This is what gives rise to the upward-sloping part of the demand curve in the figure. Nonetheless, the figure shows that the marginal willingness to pay decreases if a sufficiently large number of consumers is connected to the network. The reason for this is that those that value the service the most are already connected to the network.

From the right-hand panel of Figure 3 we see that there are three possible equilibria. If no one connects to the network ($n=0$) we have that the willingness to pay is equal to zero, $p=0$. No one uses the network, and therefore no one will pay to get connected to the network.; $n=p^*=0$ This will typically be the result if the potential users do not expect that the system will take off. If instead the consumers believe that at least a few consumers will connect to the system, we may end up in an equilibrium with $n=n^*$ and $p^*=c$. Finally, if a large number of consumers enter the system we end up at $n=n^{**}$ and $p^*=c$. Here the price is low because, as explained above, the marginal consumer has a low willingness to pay for the service.

We thus have three equilibria in the Figure, but it should be noted that the equilibrium denoted by n^* is unstable. The reason for this is that if just a few more consumers – actually, one more consumer is sufficient – enter the system, then the demand curve is above the supply curve. Thereby the willingness to pay is higher than the price, and new consumers will enter until we have reached n^{**}. If, on the other hand, one or more consumers leave the system at n^* the demand curve will be below the supply curve. In this case also all the remaining consumers will exit as well, and we end up at an equilibrium with $n=0$. Consequently, we have two stable equilibria; $n=0$ and $n=n^{**}$.

The n^{**} consumers with the highest willingness to pay will be better off if they all enter the network, and pay a price equal to c, than if no one enters the network.[4] The equilibrium where $n=0$ is thus obviously inferior. Each consumer that enters the network imposes a positive externality on the others, since she increases the value of the system. The problem is that no single consumer has any incentive to enter the network unilaterally. What determines whether the network reaches a critical mass, i.e., a point to the right of n^*, where the system grows and becomes a success?

Whether a network system consisting of for instance distribution and content will reach a critical mass depends, to a large extent, on non-strategic actors like small content providers, small software producers and, not least, on small end-users.[5] All of

[4] The last consumer that enters the network at $n=n^{**}$ has a willingness to pay equal to the price c. All the other connected consumers have a willingness to pay that is higher than this connection prices. Thus, to be precise, we should say that all connected consumers except for the last one will observe a strict welfare gain from the network if $n=n^{**}$.

[5] Historically, we have seen that the most important content providers in communication networks are the end-users; within telephony it is obviously those who make phone calls who are the most important content providers. Likewise, it is the end-users who are the most important content providers of e-mails, perhaps the most important "killer application" within narrowband Internet, and we find a similar relationship for text messages in the mobile phone network.

these market participants must be convinced that the system will actually become a success. If they believe that the system will actually become popular, then this is likely to be a self-fulfilling prophecy. If they fear that the system will not become popular, then this is also likely to be a self-fulfilling prophecy. In short, the *expectations* of non-strategic actors may be decisive for whether the system eventually reaches $n=0$ or $n=n^{**}$. This implies, for instance, that even competing (potential) content providers on the network will have a common interest in convincing the market that the system will become successful.[6] The incentives to co-operate in influencing the market expectations are certainly stronger between firms that offer complementary goods, such as infrastructure and content.

Over the last hundred years a large number of network systems have apparently reached a critical mass and become successful. Examples of this are railroads, conventional telephones, fax machines, e-mail, and – more recently – the i-mode mobile telephone system in Japan. A number of other services have experienced a different destiny, and not reached a critical mass. These services are for obvious reasons less well known, but the mobile Internet WAP in Europe seems to be a system that will vanish. The picture telephone is another example of a service that some industrialists predicted would be owned by most people, but that instead suffered a silent death.

It is important to note that the intrinsic quality of the services or products need not be decisive for whether the system will succeed. The most famous example of this is the different destinies of the competing video system VHS and Beta. Even though Beta perhaps was initially the best system, this system has in most countries disappeared to the advantage of the non-compatible VHS. Likewise, it may be argued that Mac was (and is) a better system that the non-compatible IBM PC-system, but Mac has a hard time surviving. In both these cases the dominating system has arguably been more successful in convincing the non-strategic agents that their system will become a success. These examples provide an illustration of *positive feedback;* the strong becomes stronger and the weak becomes weaker. Figure 4

[6] The co-operation between firms like Sony, Philips and Toshiba to promote DVD is an example of how competitors co-operate to launch a system with indirect network externalities. Indeed, Sony and Philips even co-operated in inventing the system. Later, they have also teamed up with suppliers of complementary goods (e.g., the content provider Time Warner). Note that the co-operation between the hardware producers implies that in the future they will compete *within* the standard. Philips and Sony also had the option to invent each their incompatible "DVD technology", in which case we would presumably have observed competition to *become* the standard.

pictures the positive feedback mechanism between two similar, but non-compatible networks, graphically; at first the two systems *A* and *B* have approximately the same market shares. At time t_1 system *A*, be it by chance or through a better marketing strategy, grows at the expense of system *B*. This makes system *A* more attractive than system *B*. The latter will lose customers, and may at time t_2 have lost so much of the market that it is below the critical mass and eventually dies out. If this happens, the positive feedback mechanism has lead to a *winner-takes-it-all* market. Though this is certainly an extreme outcome, we often observe that the system that captures the larger share of the market becomes highly profitable, while the other system struggles to survive financially.

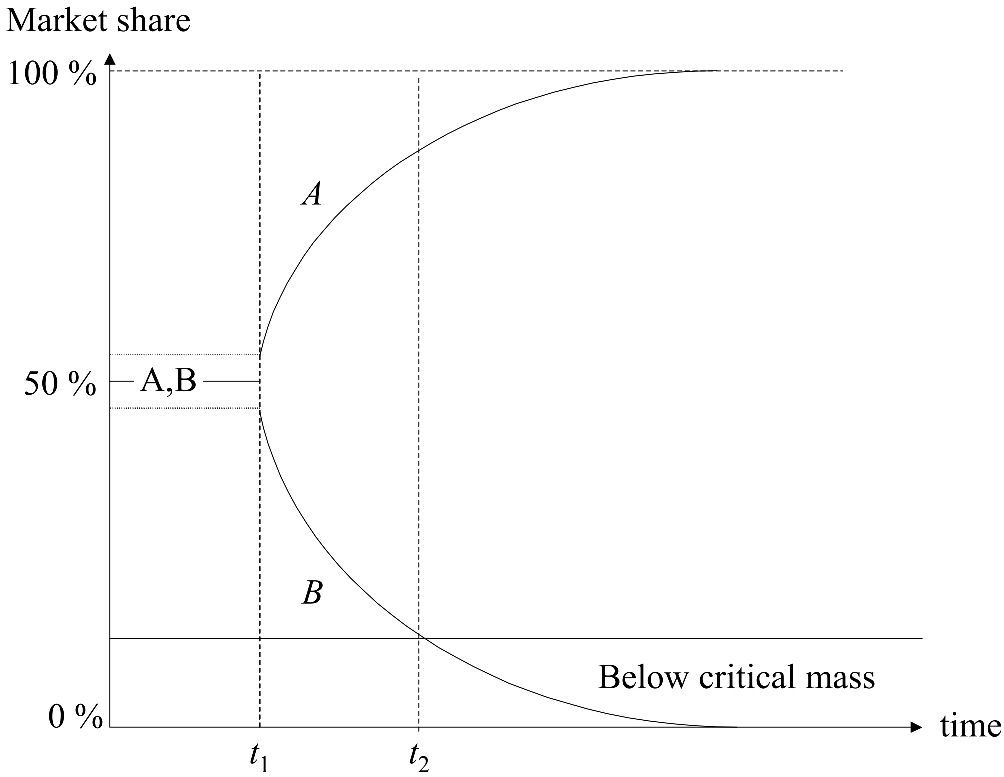

Figure 4: Positive feedback & winner-takes-it-all

2.3 Technology and distribution structure

With regard to the basic physical lines the Internet by and large uses the same infrastructure as traditional telecommunication. This is true both for local access into each single home, where the majority of the consumers uses the traditional telephone line (analogue modem, ISDN, or xDSL), and for the major transmission channels in the regional and global backbones. The local access lines can be considered as short cuts to the Internet, and as such they are not part of the Internet itself. Indeed, local Internet access through the telephone lines uses the same switching technology as traditional telephony – circuit switching. Before the user makes a conventional telephone call, or connects the telephone line to the Internet, an end-to-end connection with a given capacity is established (56 kilo bites per second with an analogue modem, and 64-128 kilo bites per second with ISDN).[7] This capacity is dedicated to the user as long as the conversation (connection) lasts, and for traditional phone calls this line switched technology is used independent of distance. Thus, a continuous end-to-end connection is set up whether one calls one's neighbor or a person on a different continent. Hence, the circuit switching technology is connection-oriented. The Internet, on the other hand, uses packet switched technology, where for instance an e-mail is broken down into several smaller data packets that are independently sent from sender to receiver. Thus, as we discuss below, the present Internet standard implies that the packet switched technology is connectionless.

The ex ante advantage of setting up a continuous end-to-end connection with a given capacity is that it is protected from possible third-party interruptions. A disadvantage is that the utilization of the capacity is poor if the capacity requirement varies over time during the connection. This will typically be the case within the Internet world, for instance when a user downloads a web page, and then reads it before a new web page is downloaded. For this kind of use connectionless packet switching is more effective than an end-to-end connection, since it allows others to use the free capacity. The disadvantage, of course, is that this may cause interruptions and delays if there is congestion.

The Internet is a network of networks that connects decentralized computers all around the world. Each single computer (host) connected to the Internet has a so-called IP address, which has clear similarities with an ordinary postal address. The IP-address identifies the computer (host id) and which sub network (net id) the computer is connected to. Communication between different computers on the network takes place by sending data packets from one computer to another, and each data packet has an address that identifies the receiver. When the packets have reached the receiver, they are sorted and assembled such that they together constitute for instance the e-mail that the receiver sees.

---

[7] Broadband access through the telephone line (ADSL) or cable-TV has capacities of 400 kilobits

The distribution of data packets from sender to receiver does also take place by using computers. These computers are termed routers (analogous to switches in the telephone network) and, as indicated by the name, have the overview of the route that the data packets shall follow. Each router thus operates a routing table. Most of these tables contain only a limited number of addresses, and data packets with unknown addresses are sent away from the router as unknown (default routing) to routers with a larger routing table higher up in the hierarchy. Only a few *core routers* have complete routing tables with an overview of all addresses in the Internet. Standardized rules or Internet protocols (IP) specify how exchange of data takes place between each single computer and between independent networks.

Let us proceed to compare the Internet distribution system to the distribution system for postal services. Local post offices make only a rough sorting of letters. At most, they sort the post intended for households belonging to the given post office. The same principle applies also within the Internet world. Routers in local networks have an address overview only for directly connected host computers. Communication between two hosts connected to the same local router can therefore be distributed directly through this router, but all other data packets are sent further up in the system as unknown (default routing). Routers at the next level have a somewhat broader overview of addresses, in the same manner as regional post distribution offices. Those who sort post in these offices hardly know the exact location of each household, but they have an overview of many local post offices. The same is true for routers higher up in the hierarchy in the Internet.

Figure 3 provides a more detailed illustration of how the addressing happens “regionally” within the Internet. Internal traffic within a local net, for instance between customer 1 and customer 2 in network A, goes via router 1. Traffic from customer 1 and 2 directed to users in any other networks is sent from router A to router C, which has a routing table with an overview of all host computers that are served by router B in local network B (default-free routing). Traffic from customer 1 to a host computer in local network B (customer 3 or customer 4) is sent from router C to router B. All other traffic is sent as default (default routing) from C into the “Internet cloud”.

per second and more.

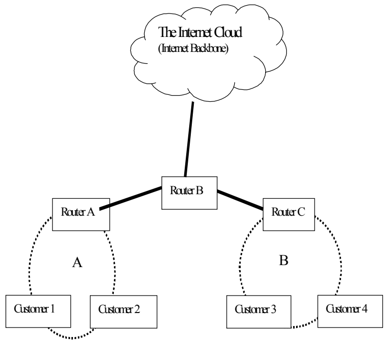

Figure 3: The Internet Cloud.

A hierarchy like the one we have described above needs a top level that does not send away data packets as unknown (default routing). In other words, the core routers at the top of the hierarchy must have complete routing tables with an overview (directly or indirectly) over all the networks further down in the hierarchy. Otherwise, some packets may end up going in indefinite loops. All core routers must be able to communicate with each other, and they must be more or less continuously updated. A small number of such core routers secure complete routing tables, and it is these core routers that define the number of addresses that can be reached over the Internet. A large number of routers with more limited routing tables are in the next round connected to the core routers. Thereby the Internet has a vertical or hierarical address and distribution structure that can be used as inputs for those that operate local and regional networks.[8]

Milgrom et al. (2000) argue that it may be cost efficient that just a few firms control the core backbones and address system in the Internet. Pure cost considerations may therefore indicate that it is optimal to let the central Internet Backbone Providers (IBPs), who control the core routers, limit the number of routers that is allowed to enter “the good company”.

3 The value system of the Internet’s market structure

In Figure 4 we provide a very simplified illustration of the value system of the basic distribution system in the Internet. The Internet Service Provider (ISP) sells access to the Internet to the end-user, and the function of the ISP is to act as a kind of portal to

[8] The core routers never send away packets as unknown, as distinct from the routers with more limited routing tables. If the core router receives a packet with an address that it does not recognize, the packet will be thrown.

the global Internet. The end-user either buys Internet connection from the ISP and local access directly from a telecommunication company (Figure 4a) or he buys both services “bundled” from the ISP (Figure 4b): The former model used to be the most common one earlier, but lately it has become more common to buy the bundled variant consisting of both Internet access and local access. This is particularly true for high-speed (broadband) Internet access. Access to the global backbone is in any case an input that local ISPs must buy directly or indirectly from those who control the top level of the Internet (the IBPs[9]), and with the bundled variant the ISP must also buy local access as an input.

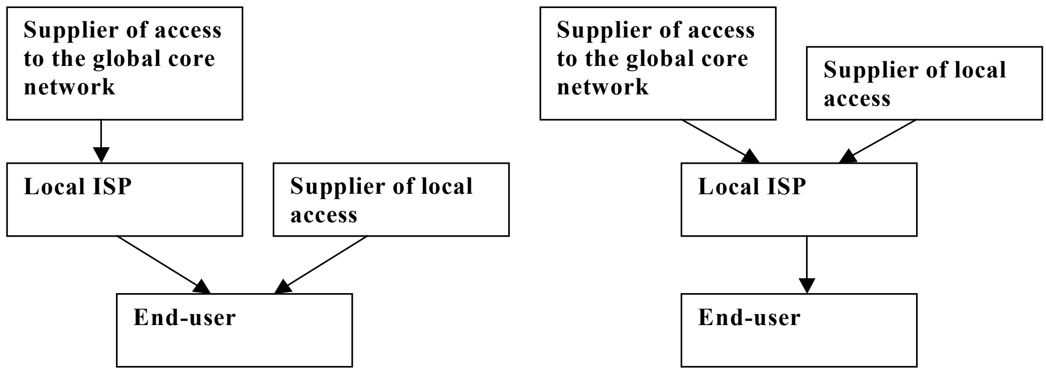

Fig. 4a: Separate supply of local access and Internet access

Fig. 4b: Local access and Internet access bundled

Today it is apparently firms with market power that control the top level, i.e., the global backbones, of the Internet. We also find firms with significant market power in the segment for local access (where the dominating telecommunication firms have large market shares). For the ISP segment the situation is different. In this segment there is a large number of firms, and entry barriers are seemingly small compared to the local access segment and the global backbone segment.

3.1 Local ISPs

Usually local ISPs operate their own local data network, but these networks are to a large extent based on leased lines in a market with relatively strong competition. The profit opportunities for independent ISPs have proven to be relatively small, since there are low entry barriers. Cave and Mason (2001) argue that a main reason for this is the prevailing regulation regimes in telecommunications. However, they also argue

[9] Internet Backbone Providers.

that we may observe increased market concentration in this segment along with increased penetration of broadband technology in the local access.[10]

3.2 Suppliers of access to the global core backbones

The fact that the addressing within the Internet takes place within a strict hierarchy has immediate implications for the market structure. Those who control the top level of this distribution system and control the core routers are in possession of an input that all the other agents down in the system must have access to in order to sell Internet access to end-users. This top level consists of a few American giants, with MCI WorldCom at the forefront. Since in addition these firms control much of the basic transmission networks, both in the USA and across the Atlantic Sea, one may argue that these companies control the global infrastructure in the Internet – denoted the global backbone in Figure 4. These firms (four to five in number) constitute what has been labelled Tier-1. Retailers that sell Internet access are dependent upon buying access to the Internet's global infrastructure as inputs from these firms.

The American giants now almost stand in a pure wholesaler-retailer relation to smaller agents, and one of the largest cost components for independent Internet suppliers is access to the global infrastructure. This is true partly because they have to pay for capacity on the transatlantic lines, and partly because they have to pay for access to the infrastructure in the USA. We thus see a clear asymmetry here between Europe and the USA. It should further be noted that Tier-1 firms like MCI WorldCom have begun to orient themselves towards Europe and integrate into the retail market for Internet access (the ISP segment) and local access. Consequently, they are no longer just wholesalers. On the contrary, they will to an increasingly larger extent also operate as retailers, and sell Internet access to end-users in competition with the retailers to whom they sell inputs.

3.3 Providers of local access

The firms that sell Internet access to me as an end-user must have a physical connection to the outer wall of my house. This is what is labelled local access. For private users it is not reasonable to believe that anyone will find it commercially profitable in the near future to build new cables into private homes (Clark, 1999).

[10] Many ISPs also offer content, but here we concentrate on access to the infrastructure for the end-user.

Thereby private users will at most have two alternatives to choose between, namely the copper cable for telephony and cable-TV. The majority of the households in Europe use the telephone line (through modem or ISDN) to reach the Internet Service Provider (ISP). Thus, the alternatives are limited with regard to local access, and the firms that control the local access network are in possession of a central component. Moreover, the dominating providers of local access are also to a large extent vertically integrated into the ISP segment. In Norway, for instance, the incumbent telecommunication firm (Telenor) controls the most important local access network in the country (through its copper network). Telenor is also the largest cable-tv provider. At the same time, Telenor is the largest retailer of Internet access.

The market power of the dominating telecommunication companies should not be exaggerated, since they are subject to comprehensive public regulation. This will be discussed in detail below. Noteworthy is that only the telecommunication companies are mandated to sell local access as an input to independent retailers. Cable-TV companies do not face the same requirement, and interestingly they have chosen not to sell local access as an input to independent ISPs. Consequently, in this case broadband Internet access has to be bought directly from the network owner. Hasuman et al. (2001) analyze the implications of this asymmetric regulation of telecommunication and the cable-TV network with focus on the USA.

4 Congestion pricing, interconnection, and access pricing

In this section we look at the following issues:

- Allocation of scarce capacity in the Internet
- Interconnection incentives between different networks
- Access price regulation, vertical integration and foreclosure incentives

These issues are closely related, not least the issues of interconnectivity and access pricing. However, for simplicity we discuss them separately.

4.1 Allocation of scarce capacity

The incremental resource costs of sending a data packet through the Internet are close to zero if the capacity is not fully utilized. However, there are large fixed costs involved in developing and maintaining line- and routing capacity. The particular problems connected to pricing of services that have low variable costs and high fixed costs will not be discussed here (see Laffont and Tirole, 2000, and Varian, 1998a, for thorough discussions). Instead we shall discuss how price signals may be used to allocate capacity when this is a scarce resource. There is no doubt that capacity has indeed at times been a scarce resource in the Internet. Overloading may take place at several network resources, for instance in the capacity of the transport network, routers and popular servers.

As for other scarce resources, my use of a scant Internet capacity will crowd out others. In absence of price signals to allocate the transfer capacity in the Internet, the users will have to pay by accepting delays or that data packets are dropped. The overloading that I create with my use, imposes a negative externality on other users.

Ideally, the price should give a signal such that the buyer takes into consideration (internalizes) both positive and negative externalities imposed on other users. My use of the Internet will not impose any kind of costs on other users if the capacity is not fully utilized, in which case allocation efficiency calls for a price equal to zero.[11] If the capacity is fully utilized, on the other hand, I should pay a price on my use that mirrors the costs that I impose on others because they are delayed or foreclosed. On some highways toll money is implemented to solve the problem of congestion.[12] The consumer then has a simple choice: she can choose not to use the road, or she can choose to use it at the price of a ticket for one car. However, the queue problem may become considerably more comprehensive and complex within the Internet, since the heterogeneity among users and applications has become so large. For instance, in a data network each user can send a few bits of data as an e-mail or several megabits per second in true-time applications where speed is essential (see Hallgren and McAdams, 1997, and Mackie-Mason and Varian, 1997).

Despite the fact that we have observed overloading and congestion problems in the Internet, neither the end-user nor the Internet Service Providers has faced prices that on the margin depend on the volume of the data flows they initiate. Firms, universities and public institutions are usually connected to the Internet via leased lines from telecommunication companies. As long as they do not use more capacity than dedicated, they will not face a price that depends on the actual use of scarce

[11] A positive price may, nonetheless, be necessary in order to cover fixed costs, or be the outcome of profit maximization.

[12] However, more often the toll money is used to raise income, whether we consider private or public road owners. It is well known from economic theory that using toll money in absence of congestion tends to reduce efficiency.

capacity. Most private users are connected to the Internet via their subscription on telephony (modem or ISDN), and therefore pay a time dependent price. However, the trend is that also private users pay a fixed fee, and this appears to become the dominating business model for the growing broadband market. So while the broadband opens up for more capacity demanding applications, the users are not likely to face a price on neither volume nor time. From some quarters it is argued that this will create larger queue problems and more delays "backwards" in the network (see Cave and Mason, 2001, for a discussion).

Today's Internet Protocol (Internet Protocol version 4, Ipv4) offers so-called *best effort services*. All users and applications are served in the same manner, *one size fits all*, and if there is an overload, the allocation of capacity takes place through the first come, first served principle. Consequently, delays and dropouts strike the users accidentally, independent of their willingness to pay or what kind of applications the data packets are part of. Thereby the users have to pay in terms of delays and dropouts. A data packet that is part of a real-time application, which is very sensitive to delays, has the same probability of being delayed as a data packet that is part of an e-mail, where a short delay may have only small or no consequences for the users. Likewise, the present system does not distinguish between a bank with a high willingness to pay and a teenager who sends a real-time video to some friends.

If both services and users were relatively homogenous, which was the case in the early years of the Internet, a uniform offer like *best effort* would function quite well as an allocation mechanism. With heterogeneous services and/or users, on the other hand, both profit maximizing and welfare-maximizing (regulated) firms typically find it optimal to offer differentiated menus of quality and price.

It is worth noting that price is not the only way to solve potential congestion problems. Actually, in the Internet one has tried to solve the congestion problem by custom and usage procedures and by overinvestments:

- One way of preventing a user from imposing excessive costs on other users is to establish disciplining rules for custom and usage. In the childhood of the Internet a set of rules were developed that intended to create self-justice among the users. Such norms could act as a disciplining tool in relatively small groups with a common set of norms. As the number of Internet users, and not least the heterogeneity among these, increases, it is unlikely that a system based on norms will be an effective way of avoiding or significantly reducing negative externalities.

- Some argue that the congestion problems within the Internet can best be solved by over-investing in capacity. This kind of approach to the problem, however, is expensive and inefficient with regard to flexibility and resource utilization. Some suppliers have tried to invest in aggregate capacity large enough to always guarantee that demand is equal to supply. The background is that the traffic volume is unpredictable and highly volatile. MacKie-Mason and Varian (1997) compare this with a bank that always keeps 100% of the deposits in ready money in case all the customers come on the same day to withdraw their money.

There exist technological solutions where the data packets may contain information about prioritizing (Ipv6).[13] A data packet can then be addressed with high priority, such that it is prioritized ahead of a packet with low priority. However, in order to create *incentive compatibility* it is presumably necessary to implement user dependent prices, such that high-priority data packets are charged higher prices than low-priority data packets whenever there is a queue. If not, it is reason to believe that an unnecessary large share of the users will set a high priority on their packets. As Mackie-Mason and Varian (1995a) put it: "… without pricing it is hard to imagine how priority schemes could be implemented. What is to stop an e-mail user from setting the highest priority if it costs nothing?"

Many proposals have been advocated for how the price mechanism could be implemented to solve the congestion problem in the Internet, and the best-known proposal is "the smart market" introduced by Mackie-Mason and Varian (1995a, 1995b). *The smart market* is an application of a Vickery auction.[14] In a Vickrey auction the winner of the auction – the one with the highest bid - pays the next-highest bid. This solves the incentive compatibility problem, since it implies that no one will have a motive to bid above or below his or her true willingness to pay (truthfulness is a dominating strategy). Applied to the *smart-market* model of Mackie-Mason and Varian the Vickrey auction is organized as follows. Each packet contains a "bid". In a

[13] The addressing room is extended in Ipv6 compared to Ipv4. Additionally to opening up for priority of data packets Ipv6 increases the access to IP addresses. Furthermore, Ipv6 also introduces better possibilities to use multi-media services, and it has better security mechanisms for e-commerce. The new addressing structure also simplifies the routing procedures.

router with congestion the packets will be given priority according to the bids in the addresses. The "winners" then pay the bid on the packet with the *lowest priority that is accepted* in the network.[15]

An auction like *the smart market* is not trivial to implement. This is also emphasized by MacKie-Mason and Varian. First, a packet may take many ways through the network, and it is difficult to guarantee that it will take the most cost-efficient way. The degree of congestion may vary on the different paths that the packet may follow. The user's value depends on how rapidly the complete data packet passes through the system; he or she is not interested in each single packet or what happens at each single router. To make an auction that extends the bid mechanism to encompass the complete data packet's way from sender to receiver will be very complex (Shenker, Estrin and Herzog, 1996).

Second, within the Internet it will be problematic to implement the general principle "sender pays", which is common for most other communication services, among them telecommunication. The value of using an application will sometimes be at the hand of the sender and sometimes at the hand of the receiver. In such networks it will therefore be important to have sufficient flexibility to be able to vary between whom you bill.

Third, there will be significant implementation problems when independent networks have to coordinate an auction. The Internet consists of a very large number of different suppliers, and the contract between them will have the same function as interconnect agreements within telecommunication. An auction system requires that flows of payments go through the different networks. Moreover, there has to be an agreement on standards for how to bid and how to set priority on packets. For instance, if the routers use different standards for priority a packet may happen to be delayed or dropped at the advantage of packets with lower priority.[16] Due to the large heterogeneity and the large number of agents in the Internet, it will be difficult to reach binding agreements on such standards. This is one reason why many experts believe that it may be impossible to implement the *smart market* within the open Internet that we know today. At best, it may be implemented within sub-networks.

---

[14] A Vickrey auction is often denoted as a 2. price auction. The background for this is that the bidder with the highest bid pays the next-highest bid. See, for instance, Vickrey (1961).

[15] Alternatively, packets that are refused may be redirected to another network. This may, for instance, take place after some time in order to see if the overload stops.

[16] See Srinagesh (1997) and Gong (1997).

There exist a number of extensions and alternatives to MacKie-Mason and Varian's *smart market* (e.g., Clark (1997), Crémer and Hariton (1999), MacKie-Mason, Murphy and Murphy (1997) Odlyzko (1997), Gibbons, Mason and Steinberg (2000) and Mason (2000)).

4.2 Interconnection incentives

The Internet is a system that consists of a number of discrete networks. Through vertical and horizontal integration and/or interconnection agreements between these networks, a customer that is connected to a small local ISP is able to reach almost any other user connected to the Internet. Other things equal, higher interconnection quality allows each ISP to charge higher prices, since the user value is increasing in the communication quality both within and between different networks. This is true for any system with economies of scale on the demand side.

It is obvious that network owners must cooperate on interconnection and on compatibility while competing for the same customers. Below, we do not distinguish between interconnection quality and compatibility. What is decisive for our discussion is that a seamless Internet world requires that the communication quality between sub-networks is equally good as the communication quality within each sub-network. If this is fulfilled, the customers will, other things equal, be indifferent with regard to the size of the sub-network to which they buy access. If the interconnection quality is relatively poor, on the other hand, the customers tend to prefer the largest sub-network. The same is true if it is more expensive to communicate between than within each network.

The underlying philosophy behind the Internet was that it should be an open network where the users can freely communicate with each other, independent of which sub-network or ISP they are connected to. Physical connection is a necessary, but far from sufficient, condition to ensure that the interface between independent networks and ISPs is invisible for the users. Additionally, a number of other virtual elements must be coordinated, presumably through comprehensive and complex contracts.

Historically, the interconnection agreements between different sub-networks in the Internet were of the form "I bring your traffic if you bring my traffic", with no flow of payments. These agreements worked amazingly well as long as the public sector financed most of the infrastructure and the Internet was characterized by homogeneity both on the supply and the demand side. Additionally, as discussed above, early applications like e-mail and transfer of data files typically tolerated delays. The latter implied that neither users nor services were particularly sensitive to small frictions in the interfaces between different networks.[17]

The 4-5 dominating IBPs at the top level of the Internet still have "I bring your traffic if you bring my traffic"- agreements with each other (*peering agreements*). However, since 1997 these firms have charged smaller IBPs and ISPs for access to the global infrastructure and addressing system in the Internet through so-called *transit agreements*.

It is an important question whether the dominating IBPs have incentives to execute market power in a manner that directly hurts smaller rivals in the same segment, local ISPs and end-users. On the one hand, there are clearly valid arguments that the top-tier firms should be allowed to cooperate on maintenance of the top level of the Internet. Smaller IBPs may, for instance, be tempted to overload other parts of the network rather than to increase their own capacity (Srinagesh, 1997). Therefore it may be optimal to restrict the number of firms that are allowed to enter into peering agreements. Put differently, it may be socially advantageous that small Internet suppliers have to pay for complete Internet access (Milgrom et al., 2000, Besen et al., 2001). Additionally, Varian (1998) argues that cooperation between the top-tier firms helps to secure high quality on the global core network in the Internet. However, Varian (1999) also argues: "The problem with such a board would be the temptation to use it as a device for collusion". So even if individual IBPs do not have a sufficiently dominating position to abuse their market power towards either smaller IBPs or retailers further down the hierarchy, the top-tier IBPs as a group may have the ability to come in such a position.

When MCI and WorldCom applied for permission to merge in 1998, it was questioned whether the new company, as a dominating IBP, would be able to partly foreclose competitors by increasing their costs (e.g., by setting a high price for interconnection) or by lowering their demand (by reducing the quality of interconnection). The most outspoken concern of the other IBPs was that the merged MCI WorldCom would choose the latter strategy; offer an inferior interconnection quality in order to gain a competitive advantage in the competition of selling inputs

[17] See Srinagesh (1997), Kende (2000) and Bailey (1997) for a detailed description of the structure

(transit) to firms further down in the Internet hierarchy. In order to avoid this scenario, both American and European competition authorities set as a precondition for accepting a merger that MCI's IBP activities were sold.[18]

What does the theory tell us about the incentives to strategically reduce the interconnection quality towards smaller rivals? On the one hand, it is clear that the existence of network externalities implies that the consumers' willingness to pay for network connection is increasing in the communication quality between different sub-networks. An increase in the interconnection quality thus generates a positive demand effect for all suppliers. On the other hand, it is also clear that a relatively large network gains a competitive advantage by setting a poor interconnection quality to smaller networks. This trade-off between a positive demand effect and a negative competitive effect (quality differentiation) was first analyzed by Katz and Shapiro (1985), who showed that a large network in general has lower incentives to set a high interconnection quality than have smaller networks.

4.2.1 Interconnection incentives between networks with installed bases

In many markets we observe that competing network firms, not least telecommunication incumbents, have some existing customers that are more or less 'locked in'. How does this phenomena affect the interconnection incentives? Crémer et al. (2000) extend the model by Katz and Shapiro (1985) in order to analyze this question in context with competition between two IBPs that have each their base of installed customers. They show the IBPs will always have incentives to set a high interconnection quality if they have the same size of their installed bases. However, they also show that if one IBP has a larger base of installed customers than the other, then the firm with the larger base may have incentives to reduce the interconnection quality towards the rival.

It should be noted that Cremer et al. assume that new customers have a higher willingness to pay for network connection the larger the effective networks, while the income from the installed bases is independent of the network sizes. In contrast, Foros, Kind and Sand (2002) assume that also customers in the installed bases are charged a price that is increasing in the total network size and in the interconnection

and history behind the interconnection arrangements in the Internet.

[18] In connection with this case, it should also be mentioned that MCI WorldCom planned to merger with Sprint (a major IBP) in year 2000, but that the European Commission stopped these plans.

quality. How does this income structure affect the interconnection incentives for the larger firm?

On the one hand, we may expect that the larger firm will have relatively strong incentives to improve the interconnection quality, because this makes it possible to charge higher prices from customers in its installed base. In line with this, Foros et al show that the larger firm is willing to set a high interconnection quality even if it means that it captures a lower number of new customers than would be the case with a poor interconnection quality. On the other hand, as shown by Katz and Shapiro, the larger firm gains a competitive advantage if it sets a low interconnection quality. This effect is shown to be more likely to dominate the larger the difference between the installed bases and the higher the price the firms can charge from each customer in the installed bases. Furthermore, the larger firm is more aggressive in the competition for new customers the higher the price paid by customers in its installed base, and this may harm the smaller rival even if there is perfect interconnection quality.

The existence of installed bases of customers may have implications for international and regional competition in telecommunications. When a firm like AOL Time Warner enters a regional market in Europe, for instance, they compete with a regional ISP. AOL Time Warner's customer base in the USA may be seen as an installed base or a clientele. Obviously, AOL Time Warner may gain a competitive advantage by offering the regional ISP a low interconnection quality with the customers that AOL Time Warner has in the US. However, it is likely that AOL Time Warner's income from American customers also depends on the interconnection quality with European Internet users that are connected to regional ISPs. Typically, the revenue from the installed base customers will be higher if there are more people to have high quality communication with. Intuitively, the gain from the installed base from a high interconnection quality may well offset the loss due to reduced competitiveness in the new market.

As another example, consider the market for broadband access to residential users. The two main alternatives are offered by telecommunication incumbents (who upgrade their copper network to handle DSL) and by cable-TV providers. In Europe the coverage of the telecommunication network is much larger than that of the cable-TV-networks. Hence, we have a duopoly in some regions (typically in urban areas), while we have a monopoly controlled by the telecommunication incumbent in other regions (rural areas). Suppose that there are strong network effects such that the reservation price of a customer increases with the number of broadband users and with the interconnection quality between DSL and the cable-TV network. Since existing broadband users in rural areas have no alternative access possibilities, they can be seen as an installed base or a clientele for the incumbent. Thereby the telecommunication incumbent has the ability to create a competitive advantage over the cable-TV providers also in urban areas if it degrades the interconnection quality. The degradation may take place by reducing the dataflow capacity between the networks, such that for instance an interactive videoconference between people in rural and urban areas is possible only if they subscribe to the incumbent. However, degrading the interconnection quality reduces the reservation price from the customers in the monopoly area, and this may well dominate the competition effect.

Other analyses that focus on the relationship between IBPs include Milgrom et al. (2000), Besen et al. (2001), Laffont et al. (2001a, 2001b), and Little and Wright (2001). Foros and Hansen (2001) analyze interconnection incentives for local ISPs. They model the competition á-la Hotelling, and show that the ISPs may have incentives to set a high interconnection quality, because this reduces the intensity of the price competition.

In all the models discussed above it is assumed that agents do not charge each other for interconnection. This means that interconnection arrangements take the form of the peering agreements discussed above. Limited possibilities to charge for interconnection may be caused by regulation or problems with writing contracts (see Cremer et al., 2000). But as we have seen, the top-tier firms do charge the smaller firms for access to the core network through transit agreements. Moreover, telecommunication incumbents charge independent ISPs for local access as an input. In the next section we discuss access pricing, and some accompanying regulatory challenges, when the access providers are vertically integrated into the retail market.

4.3 Regulation, vertical integration and foreclosure incentives

The last few years we have observed a trend where firms that control essential inputs like the global core network or local access to households vertically integrate into the retail market and sell Internet access downstream directly to end-users. Thus, they will consider other downstream firms both as competitors and as customers, and there is reason to fear that they will utilize their market power to gain a competitive advantage relative to their rivals. An obvious way to do this is to (partly) foreclose the rivals by reducing the quality or increasing the price of the essential input.

It should be noted that foreclosure strategies are far from new, and this can be illustrated by a well-known example from the early years of the telecommunication industry. About 100 years ago, the Bell System (later AT&T) had less than half of the telephone subscribers in the US, and faced competition from a number of local competitors in the country. However, Bell System was the dominant provider of long-distance calls, and they established a foreclosure strategy by denying local rivals access to Bell System's long-distance network. Thereby Bell System was able to offer a product that telephone subscribers perceived as being better than what the local rivals could offer. This strategy implied that Bell System quickly out competed their local rivals (see Shapiro and Varian, 1998a).

There is a large strand of literature that discusses questions related to access pricing and regulation of the same for essential inputs. Laffont and Tirole (2000) and Armstrong (2001) offer comprehensive overviews on both access pricing literature and on existing regulation regimes. Rey and Tirole (1996) analyze the incentives that an unregulated vertically integrated upstream monopolist has to foreclose downstream rivals, while a number of other articles analyze the incentives for similar regulated firms to circumvent the regulation by reducing the quality of the inputs that they sell to their rivals. Economides (1998a,1998b), for instance, argues that a regulated vertically integrated upstream monopolist will always chose to practice such foreclosure. Foros, Kind and Sørgard (2001), Sand (2002), Sibley and Weisman (1998), Weisman (1995, 1998), Reiffen (1998), Mandy (2000), and Weisman and Kang (2001), on the other hand, argue that whether the firm will actually use a foreclosure strategy depends on how strict is the price regulation.[19]

4.3.1 Current regulation of telecommunication

The end-user market for Internet connectivity is currently unregulated in most countries, while the input segment for local access is regulated both with respect to price and quality. According to Laffont and Tirole (2000) the regulators' decision not to regulate the end-user market builds on two premises. First, if the local bottleneck is

[19] Even though all these papers are motivated by telecommunication and network industries, none of them explicitly take network effects into account. The difference from the models that we discussed in section 4.2 on interconnection, is that in the present context the large firm (i.e., the one which controls

eliminated, then head to head competition in the retail market ensures that there is no need for regulation. Consequently, regulation of the local access bottleneck is sufficient to ensure competition in the retail market. Second, the products and services in the end-user market change very fast, and this makes it very costly to monitor the retail markets compared to the wholesale markets.

The former argument, that it is sufficient to remove the local bottleneck, deserves a comment. First, the fact that the end-user market is unregulated creates an incentive for a vertically integrated provider of local access to discriminate against rivals in the retail segments. Several researchers have examined this issue, see, e.g. Laffont and Tirole (2000). Second, since the market for global access is unregulated, there is also an asymmetry between the two complementary infrastructure inputs, and this will be of particular interest due to the US dominance in provision of the unregulated input (global access).

The prevailing regulation regime of local access in Europe is cost-oriented, which means that the incumbent is not allowed to charge higher access prices than those reflecting its long-run marginal costs.[20] The incumbent controlling the local telephone network often uses three main arguments against cost-based regulation. The first argument is that it is practically impossible to compute the long run marginal cost in an industry involving large joint costs. The second argument is that the local access network for telephony no longer constitutes a bottleneck, because cable-TV and wireless networks are bypass opportunities for residential users. The third argument is that cost oriented regulation will reduce the incumbent's dynamic incentives to invest in infrastructure and product innovation. The danger that regulation creates dynamic inefficiency is an important topic in all technologically advanced industries. We will not go into this discussion (see Laffont and Tirole, 2000), which basically has the same arguments whether we consider the traditional market for telephony services or the market for Internet services. We will instead focus on the interplay between providers of the complementary inputs local and global access and on the timing of the interaction between the domestic regulator and the market players.

The current sector specific cost-based price regulation for local access is often seen

---

the essential input) will always want to impose a foreclosing strategy unless it is able to make a profit on providing access (given that the goods on the retail market are homogenous).

[20] More generally, all local access providers that have a "dominating" market position are subject to regulation. This will invariably be true for the telecommunication incumbent, but may also be true for more recent entrants.

as a "hands-on" *ex ante* approach, while the competition rules are seen as an *ex post* regulation approach. This distinction may be misleading, since the current cost-based sector regulation *de facto* will often appear as *ex post* regulation (see e.g. Laffont and Tirole (2000) and Foros and Kind (2000, 2001)). Below, we compare a situation where the domestic regulator credibly commits itself to a given price policy for local access before the input suppliers choose their wholesale prices with a situation where the domestic regulator does *not* make such a commitment. The former we refer to as *ex ante* regulation, while the latter we refer to as *ex post* regulation. Hence, in our context *ex ante* regulation does not necessarily imply cost-based price of local access. Indeed, it may be argued that cost-based regulation of local access may be inoptimal due to the growth of the Internet.

4.3.2 Does the rise of the Internet change the optimal regulation policy?

An increasingly larger share of the traffic on the telephone network is connected to the Internet, while the share of traditional voice telephony is decreasing. Combined with the fact that most of the Internet traffic goes through the USA this may have some implication for the optimal regulation of local access prices in Europe. In particular, it should be noted that access to the global Internet backbone is an essential input that together with the transatlantic telephone cables is controlled by a few large American companies (c.f., Section 3.2). Local ISPs that sell Internet connectivity to end-users have to purchase access to the global infrastructure as an input, which is *complementary* to other essential inputs. On this background Foros, Kind and Sørgard (2002a, 2002b) argue that a cost-based regulation of local access is possibly detrimental to national welfare in Europe, since it may imply excessive profit shifting to American firms. The reason for this is simply that the American firms may increase the price of global access if European regulators reduce the price of local access. The optimal regulation policy that seeks to maximize national welfare may therefore imply that the regulator commits itself to set relatively high prices on local access, even if this should reduce domestic competition.

The distinction between ex post and ex ante regulation is potentially important when we consider the effects of domestic regulation of the local access price. Ex post regulation may actually reduce welfare compared to market equilibrium. The reason for this is that the foreign firms are aware of the fact that the regulator ex post has an incentive to set the price of local access equal to long-run incremental costs. This in turn gives the foreign firms incentives to set relatively high prices on access to the global backbone. If the regulator can commit itself to set a relatively high price, on the other hand, the foreign firms may have incentives to set relatively low prices. This is due to the fact that local and global access are complements; the higher the price of one of these inputs, the greater the incentives to reduce the price of the other input. This is the opposite of what would be the case between substitutes, where a price increase by one firm typically will be followed by a price increase by the other firm as well.

The fact that a strict price regulation may be detrimental to welfare because it leads foreign firms to set higher prices raises the question of whether there is a need for some kind of supranational regulation of global access prices. The problem, however, is that many of these firms are vertically integrated into the end-user market. Thereby they may have incentives to implement quality-reducing actions towards downstream competitors. In the case of IBPs, for instance, it seems difficult to impose quality restrictions. For example, suppose that a retail subsidiary of MCI WorldCom offers a better Internet service than local competitors. It is then almost impossible for an international regulatory authority to decide whether this is because the subsidiary has a superior technology, or whether MCI Worldcom just discriminates the rivals by degrading the quality of global access.[21] In Foros, Kind and Sørgard (2002a) it is shown that an international regulation may increase welfare, but only if the global access price is set so high that the firms do not have incentives to foreclose the rivals.

4.3.3 Should USO be imposed on the broadband access market?

Historically, telecommunication incumbents have been required to charge uniform prices throughout the country. Thus, the end-user cost of installing and using a telephone has typically been the same in rural areas, where it is expensive to provide the service, as in urban areas. The requirement to charge uniform prices is commonly described as a Universal Service Obligation (USO). Lately, it has been discussed whether USOs should also be imposed on providers of broadband access.

[21] The Microsoft case gives an illustration of the problems in such a context, see, e.g., Economides (1998b).

Broadband access is the last mile of the telecommunication network, and it is an essential component in order to offer broadband Internet connectivity. A key technological feature of this market is that it is considerably more expensive to connect consumers in rural locations than in urban locations. In an unregulated market we may therefore expect that the price of access to the broadband would be higher in rural locations than in urban locations. This is true independent of whether the market is served by a monopoly or by several competing firms. We have thus seen a political concern that peripheral locations will be harmed unless broadband access providers are required to charge the same price for the same service in all locations that they cover (uniform prices). However, even though there may be implicit or explicit political requirements of uniform prices, the actual price level will hardly be regulated. Instead, as in other industries, governments seek to prevent unduly high prices by inviting several firms to compete. Some implications of this policy mix are discussed in Foros and Kind (2002).

First, it should be noted that the socially optimal regional coverage may fall if there is a requirement of uniform pricing. The intuition for this runs as follows: The fact that it is relatively inexpensive to serve consumers in locations with a high population density indicates that also the access price should be low. However, a low price induces too high demand in peripheral locations, where the real costs of providing broadband access are high. In order to reduce the magnitude of the latter effect, it is socially optimal not to serve some of the least populated areas. This clearly indicates that uniform pricing may be a poor regional policy.

Second, increased competition need not improve welfare if there is a requirement of uniform prices. While a monopolist will still have incentives to set the same regional coverage as the social planner, the coverage level decreases if there is competition. Competition reduces prices, but herein lies, in a sense, also the problem: due to the convexity of the cost function, the lower market price makes it less profitable to serve peripheral locations. Competition therefore implies that the regional coverage falls to a sub-optimal level, and this negative welfare effect is more likely to dominate the larger the number of firms that offer broadband access. Consequently, welfare may be lower with free entry than if the market is served by a monopolist.

The fact that it is relatively more expensive to serve rural areas than urban areas is not unique for the broadband access technology. There is a similar cost structure also for, e.g., postal services and third generation mobile telephone systems (UMTS in Europe). In some countries (like France, Norway and Sweden) the governments have specified a minimum regional coverage by the firms that are granted UMTS licenses, and proposals have been advanced to specify similar requirements for firms providing broadband access. In an extension of the basic model Foros et al. therefore assume that the government is able to set a binding coverage requirement prior to downstream competition between the firms, and they show that this has a positive effect on aggregate consumer surplus. More surprisingly, this policy may also increase the profit level of the firms. The reason for this is that the regulator, by acting as a first-mover, solves a co-ordination problem; the oligopolistic firms would prefer the same regional coverage as the one chosen by a hypothetical monopolist, but this does not constitute an equilibrium in a free market economy. Thus, by requiring the firms to build out to larger areas the government may actually be able both to increase the profitability of the firms and to increase the geographical coverage. This suggests that a mere requirement of uniform prices may be bad policy; it should be combined with a requirement of geographical coverage in addition.

5 Conclusion

In this article we have provided a short overview of the history and value chain of the Internet, and argued that increased heterogeneity among users and applications may require new allocation mechanisms to ensure efficient capacity utilization. Moreover, we have discussed how owners of "large" networks may have incentives to foreclose smaller rivals on the same level of the Internet's hierarchy, and how vertical integration may imply that upstream firms in control of essential inputs may find it optimal to foreclose rivals in the downstream market.

We have focused on firms in the infrastructure segments. However, it should be noted that for instance the merger between AOL and Time Warner has attracted much attention on welfare effects of vertical integration between content providers and access providers. Rubinfeld and Singer (2001) have analysed whether a merged AOL Time Warner will have incentives to implement different kinds of foreclosure strategies in the market for broadband access. The merged company has a strong position both in the segment for broadband access (through its cable-TV network in the USA) and in the content segment. First, the merged company can create a competitive advantage for its own access provider by restricting admission to AOL Time Warner's content for other access suppliers. Secondly, it can create a competitive advantage for its own content provider by reducing the quality on external content for its own access customers. On this background, Rubinfeld and Singer (2001) argue that AOL Time Warner will have both the ability and incentive to foreclose competitors. Similar concerns have been put forward with regard to telecommunication incumbents that enter into for instance the newspaper industry. The vertical convergence between media and IT is thus a great challenge for regulators, not least since a strict national regulation may increase the market power of foreign firms.

6 References

Armstrong, M. 2001. "The Theory of Access Pricing and Interconnection", in Cave, M., S Majumdar, and I. Vogelsang (eds): *The Handbook of Telecommunications Economics*, North Holland, forthcoming.

Bailey, J. 1997. The Economics of Internet Interconnection Agreements. In: *Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 155-168.

Besen, S., P. Milgrom, S. Mitchell, and P. Srinagesh. 2001. Advances in Routing Technologies and Internet Peering Agreements, *American Economic Review, Papers and Proceedings,* 91, 292-296.

Cave, M. and R. Mason. 2001. The Economics of the Internet: Infrastructure and Regulation, *Oxford Review of Economic Policy*, 17(2), 188-210.

Clark, D. 1997. "Internet Cost Allocation and Pricing", in: *Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 155-168.

Clark, D. 1999. "Implications of Local Loop Technology for Future Industry Structure", in S. E. Gillett and I. Vogelsang (eds), *Competition, Regulation, and Convergence*, LEA, London

Crémer, J, P. Rey and J. Tirole. 2000. Connectivity in the Commercial Internet, *Journal of Industrial Economics*, XLVIII, 433-472.

Crémer, J. and C. Hariton. 1999. The Pricing of Critical Applications on the Internet, Mimeo, Toulouse.

Economides, N. 1998a. The incentive for non-price discrimination by an input monopolist, *International Journal of Industrial Organization*, 16, 271-284.

Economides, N. 1998b. Raising Rivals' Costs in Complementary Goods Markets: LECs Entering into Long Distance and Microsoft Bundling Internet Explorer, Discussion Paper EC 98-03, Stern School of Business, N.Y.U. available at http://raven.stern.nyu.edu/networks/

Foros, Ø. and B. Hansen. 2001. Competition and Compatibility among Internet Service Providers, *Information Economics and Policy*, 13(4), 411-425.

Foros, Ø. and H.J. Kind. 2002. The Broadband Access Market: Competition, Uniform Pricing and Geographical Coverage, *Journal of Regulatory Economics*, forthcoming.

Foros, Ø. and H.J. Kind. 2001. 'National and Global Regulation of the Market for Internet Connectivity', in Thorsten Wichman (ed.): *Economics and the Internet: Proceedings from the Third Berlin Internet Economics Workshop,* Berlecon Research, 2001, 33-45.

Foros, Ø. and H.J. Kind. 2000. The Internet Market Structure: Implications for National and International Regulation. *Telektronikk,* 96 (2), 45-59.

Foros, Ø., H.J. Kind and L. Sørgard. 2002a. Access Pricing, Quality Degradation and Foreclosure in the Internet, *Journal of Regulatory Economics*, 22(1)

Foros, Ø., H.J. Kind and L. Sørgard. 2002b. 'International Complementarities in the Internet: Should Local Access Prices be Regulated?' Discussion Paper 9/02, Norwegian School of Economics and Business Administration.

Foros, Ø., H.J. Kind and J.Y. Sand. 2002. Do Incumbents Have Incentives to Degrade Interconnection Quality in the Internet?, SNF WP 22/02.

Foros, Ø. 2002. Competition, complementarity, and compatibility in the Internet. Manuscript.

Gibbens, R, R. Mason and R. Steinberg. 2000. Internet Service Classes under Competition, *IEEE Journal on Selected Areas in Communications*, Vol. 18(12), December 2000, 2490-2498.

Gong, J. and P. Srinagesh. 1997. "The Economics of Layered Networks", in*: Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 63-76.

Hallgren, M. and A. McAdams. 1997. "The Economic Efficiency of Internet Public Goods", in*: Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 455-478.

Hausman, J. A., J. G. Sidak, and H. J. Singer. 2001. Cable Modems and DSL: Broadband Internet Access for Residential Customers*, American Economic Review, Papers and Proceedings,* 91, 302-307.

Katz, M. and C. Shapiro. 1985. Network Externalities, Competition and Compatibility, *American Economic Review*, 75, 424-440.

Kende, M. 2000. The Digital Handshake: Connecting Internet Backbones. OPP WP 32, Federal Communications Commission.

Laffont, J.J. and J. Tirole. 2000. *Competition in Telecommunication*, The MIT Press,

Laffont, J. J., S. Marcus, P. Rey and J. Tirole. 2001a. Internet Interconnection and the Off-Net-Cost Pricing Principle, manuscript, IDEI, Toulouse.

Laffont, J. J., S. Marcus, P. Rey and J. Tirole. 2001b. Internet Peering, *American Economic Review, Papers and Proceedings,* 91, 287-291.

Little, I. and J. Wright. 2000. Peering and Settlement in the Internet. *Journal of Regulatory Economics,* 18(2), 151-173.

MacKie-Mason, J., L. Murphy and J. Murphy. 1997. “Responsive Pricing in the Internet”, in*: Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 279-304.

Mackie-Mason, J. and H. Varian. 1997. “Economic FAQs about the Internet”, in *Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 27-62.

MacKie-Mason, J. and H. Varian. 1995a. “Pricing the Internet”, in: *Public Access to the Internet.* Kahin, B. and J. Keller (eds). Englewood Cliffs, N.J., Pretence Hall.

MacKie-Mason, J. and H. Varian. 1995b. “Some economics of the Internet”, in: *Networks, infrastructure and new task for regulation.* Sichel, W. (ed) . Ann Arbor, Michigan, University of Michigan Press.

Mandy, D. 2000. Killing the Goose That May Have Laid the Golden Egg: Only Data Knows Whether Sabotage Pays, *Journal of Regulatory Economics*, 17, 157-172.

Mason, R. 2000. "Simple Competitive Internet Pricing", European Economic Review, Vol. 44(4-6), 1045-1056.

Milgrom, P., S. Mitchell, and P. Srinagesh. 2000. “Competitive Effects of Internet Peering Policies”, in Vogelsang, I. and B.M. Compaine (eds), *The Internet*

Upheaval, The MIT Press.

Odlyzko, A. 1997. A Modest Proposal of Preventing Internet Congestion, Mimeo, AT&T Labs.

Rey, P. and J. Tirole. 1996. A Primer on Foreclosure. Mimeo, IDEI, Toulouse. (Forthcoming in *Handbook of Industrial Organization.* M. Armstrong and R.H. Porter (eds)).

Reiffen, D. 1998. Regulation and the Vertically Integrated Firm: A Reevaluation and Extension of Weisman's Result, *Journal of Regulatory Economics,* 14, 79-86.

Rohlfs, J. 1974. A Theory of Independent Demand for Communications Service. *Bell Journal of Economics*, 5, 16-37

Rubinfeld, D. and H. Singer. 2001. Vertical Foreclosure in Broadband Access? *The Journal of Industrial Economics*, XLIX, 299-318.

Shapiro, C. and H. Varian. 1998a. *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, Boston, Massachusetts

Shapiro, C. and H. Varian. 1998b. Networks Effects, Manuscript.

Shenker, S, D. Clark, D. Estrin and S. Herzog. 1996. Pricing in computer networks: reshaping the research agenda. *Telecommunications Policy*, 20, 183-201.

Sibley, D.S. and D.L. Weisman. 1998. Raising Rivals' Costs: The Entry of an Upstream Monopolist into Downstream Markets', *Information Economics and Policy*, 10, 451-470.

Srinagesh, P. 1997. "Internet Cost Structures and Interconnection Agreements", in: *Internet Economics*. Bailey, J. and L. McKnight (eds). Cambridge, MIT Press, 121-154.

Varian, H. 1999. Market Structure in the Network Age. Available at www.sims.berkeley.edu/~hal/people/hal/papers.html

Varian, H. 1998. How to Strengthen the Internet's Backbone. *Wall Street Journal*, June 8, 1998

Vickrey, W. 1961. Counterspeculation, auctions and competitive sealed tenders. *Journal of Finance*, 16, 8-37.

Werbach, K. 1997. Digital Tornado: the Internet and Telecommunications Policy. OPP WP 29/99, FCC.