

Structure Conduct Performance Analysis of the Internet

By Jorim Schraven

Table of contents:

1 Introduction:

- 1.1 Product:
- 1.2 Supply:
- 1.3 Demand:

2 Structure:

- 2.1 Horizontal concentration
- 2.2 Vertical integration
- 2.3 Product differentiation
- 2.4 Diversification
- 2.5 Barriers to entry

3 Conduct:

- 3.1 Short Run - Long Run
- 3.2 Individual vs Collective
- 3.3 Access policies
- 3.4 Pricing
- 3.5 Quality
- 3.6 Capacity choice

4 Performance:

- 4.1 Private
- 4.2 Social

5 Regulation:

- 5.1 Anti-trust
- 5.2 Special rules

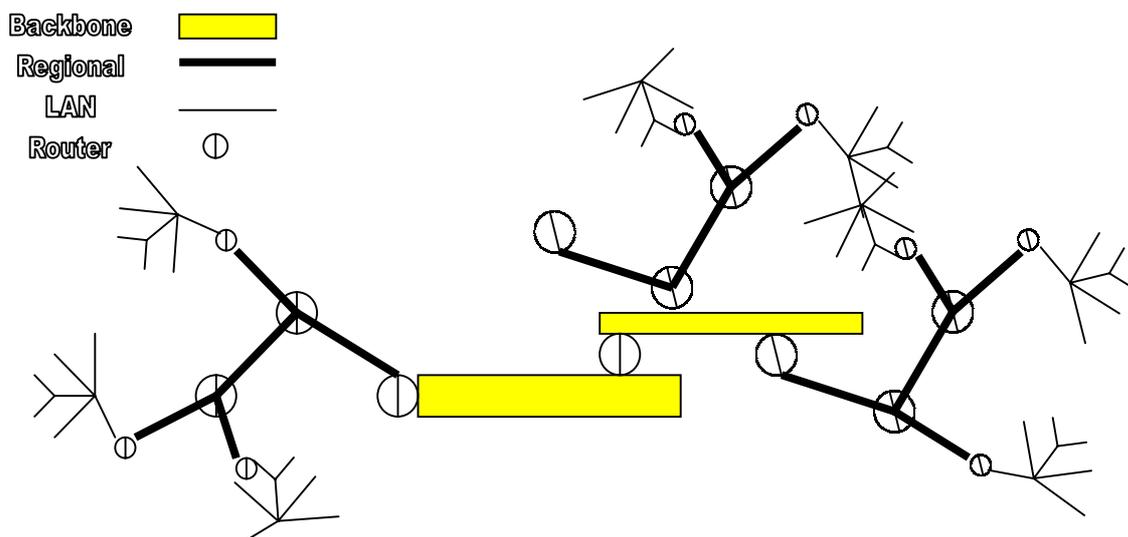
6 Conclusion

7 Bibliography

1 Introduction:

This paper performs a Structure Conduct Performance (SCP) analysis of the Internet industry. Before tackling that task, this section aims to set out the fundamental technological characteristics of the Internet, thus facilitating the economic analysis below. In popular folklore the Internet is often associated with, or even used as a synonym for, the World-Wide-Web. These are, however, different albeit interconnected industries. The World-Wide-Web refers to the content network, whereas the Internet refers to the physical network. The World-Wide-Web can be seen as an encyclopaedia with numerous cross references (hyper-links) in its entries, which constitute an information network. These hyperlinks are the arcs which allow the user move from one entry (node) to another on the network. The Internet is the physical and organisational infrastructure which facilitates exchanges of digital information. When surfing the World-Wide-Web, each click on a hyper-link is equivalent to a request for information, causing digital information packets to be sent and received by the user. In addition to facilitating information exchange via the World-Wide-Web the Internet facilitates information exchange via other means, such as direct file transfer, Gnutetella based private exchange networks (for chatting and files transfer) and Email, which may not be directly linked to the World-Wide-Web network.

The Internet facilitates these exchanges of information through the connection of the work-stations of end-users (end-nodes) to assembly points (interim-nodes), which assemble and re-route information packets to facilitate transfer to the destination node via data lines (arcs), such as optical fibre, telephone lines and television lines. The Internet is said to have a three level hierarchy, namely the backbones, the regional networks, and the Local Area Networks (LAN). The backbones connect regional networks at a national (US: NSFNET, SprintLink) or international (EU: Ebone) level and tend to have the highest capacity for transporting information packets simultaneously (parallel), that is they have the highest bandwidth. The regional networks tend to have lower bandwidth but more nodes connecting numerous LANs and dial-in users directly and



connecting to other regional networks directly and/or indirectly through the backbone(s). The LANs are smaller networks of yet lower bandwidth, generally managed by a university or organisation with several work-stations being directly interconnected.

Although a sizeable part of the Internet is comprised of the telephone infrastructure, its usage of these lines is completely different from a traditional telephone conversation. The latter reserves a dedicated connection between two nodes, thus monopolising the arcs for the duration of the information exchange. This means, that even if the two persons are not talking the full bandwidth of the reserved line is only used for their conversation. The internet uses statistical multiplexing, which entails the break-up of an information exchange into small size packets (generally about 200 bytes), which are then sent via any available combination of arcs and nodes and simultaneously over the same lines as packets from other information exchanges. This is achieved by placing routers at each intermediate node, which have dynamically updated lists of available nodes to which they can distribute incoming traffic. At the final receiving node, packets arrive non-sequentially, and are re-assembled by the work-station. This means that when an information exchange does not require much bandwidth, the excess can be used by other information exchanges. As long as re-routing is cheaper than line-rentals (cost of bandwidth) and sequential arrival of packets within a maximum delay interval is not essential (as with telephone conversations), this tends to yield a more efficient usage of the network capacity in terms of volume and quality of data transfer.

1.1 Product:

"Internet Service Providers (ISPs) offer their customers a bundle of services that typically includes hardware and software, customer support, Internet Protocol (IP) transport, information content and provision, and access to individuals and information sources on the Internet." (Srinagesh, 1996). The actual product offered by the Internet is the transport of specific information packets from one node to another. The current organisation of the Internet, through the TCP/IP protocol, does not allow comprehensive prioritisation of packets at routers. Consequently, there is no product differentiation in terms of relative fast or slow packet transport between nodes. As such, this transport has almost no quality vector even though strictly speaking the service is not homogeneous because each information packet is unique, packets are treated in a homogeneous way at the routing level. However, the capacity of the network, in terms of bandwidth, and the quality of routers, combined with the nature of the interconnection agreements between parts of the network do affect the average transmission speed of packets between nodes. As such, there can be product differentiation in terms of delivery speed and quality between nodes. For example, a intensive user may rent a high bandwidth private line

connection onto a regional network assuring superior connectivity for the 'last mile' relative to a simple dial-up user. This would guarantee that on average packages can be transmitted and will be received at higher speed for the high-end user. However, once their respective packets reach a particular interim node, they will be treated equally.

Nevertheless, by simply defining the product as the transport of information packets between nodes given the network capacity in terms of bandwidth we ignore what these packets of data represent, namely information. The point is that as a product, bandwidth can only be considered from either the sender's or the receiver's perspective. Whereas in practice, both the sender and the receiver may have an interest in the receiver's receipt of the package in case of information (Crawford, 1996). Consequently, we may run into trouble when contemplating a pricing scheme for bandwidth. For instance, if the receiver pays for the bandwidth, how can we be sure that the receiver requested the information sent?

Therefore, there are two things that must be realised when discussing markets for bandwidth. Firstly, even though the Internet itself does not distinguish between different packets of data, there are underlying reasons for which they are sent, crucially connected to their informational content, which prevent a straightforward interpretation of bandwidth as the product. Secondly, "in a commodity network (oil, gas, water, or electricity), the objects transmitted are generic and perfectly interchangeable. In an information network (mail, phone, computer data), the objects sent may be individualized and not interchangeable" (Crawford, 1996). This means, that in addition to the bandwidth aspect of the product, which could be interpreted as the size of the pipeline at the final node, the product is defined by the delivery of a packet to a particular node.

To summarise, the product is the transport of unique data packets between particular nodes at a point in time with a particular (not necessarily controllable) delay.

1.2 Supply:

Supplying the service to transport data between two network nodes requires these nodes to be connected. As described above, there are different types of connection possible between nodes, ranging from slow telephone and LAN lines to the large fiber-optical cables forming the backbones. In practice, most data will pass more than one or even each of these types of lines as it is transported from one final node to another. This is facilitated by access agreements between the backbones, regional networks and LANs (discussed in more detail in section 3.3). As a consequence of these agreements, the marginal cost for the supplier of transporting an additional package tends to be zero. Each additional packet transported does however pose a negative (congestion) externality on the other packets being transported on the network. The nature of costs in the industry is summarised by Srinagesh (1996):

“An estimate based on an analysis of several midlevels suggests that IP transport accounts for 25 to 40% of a typical ISP's total costs. For many ISPs, transport costs are sunk over the business planning horizon. The major cost of constructing fiber optic links is in the trenching and labor cost of installation. The cost of the fiber is a relatively small proportion of the total cost of construction and installation. It is therefore common practice to install "excess" fiber. According to the FCC's Fiber Deployment Update (May 1994), between 40 and 50% of the fiber installed by the typical interexchange carriers is "dark"; the lasers and electronics required for transmission are not in place”.

The per packet transport costs involve minimal electricity expenditure and a Samuelson ice-berg type transport cost in terms of the positive probability of packets getting lost or corrupted in transit.

Supply substitutability of the product differs depending on the type of connection being used. Clearly, when telephone, electricity or television cables are used, capacity can be used for these respective services instead of data transfer without high conversion cost, yielding high substitutability. Conversely, for the industry as a whole, especially for the backbones and LANs which tend to be exclusively built for digital data transfer supply substitutability is low. However, capacity could be used to create a separate network, for instance for security reasons, which could yield product differentiation.

1.3 Demand:

Initially, “government support of the Internet through government agencies such as the National Science Foundation, NASA, and the Department of Energy have helped create a critical mass of users connected to the Internet” (Bailey, 1996). Since the internet is commercially exploited the number of users of the Internet has been growing rapidly, reflecting among other things the increase in availability of high quality information and services to users of the World-Wide-Web and positive network externalities in their subscriptions.

There is a huge heterogeneity in terms of types of users on the internet. They range from private users with low bandwidth mobile or simple dial-up or higher bandwidth (e.g. ADSL) connections to government departments, universities or companies aggregating demand of their entire networks, to even regional networks or backbones. Thanks to the international nature of the internet and the associated time-zones of its users, at the macro-level demand for data transport is variable and not

highly predictable. Congestion dynamics can have a stabilising effect, with some high-bandwidth users (such as file transfers) rescheduling usage for low congestion periods. Additionally, the standard settings of the TCP/IP protocol used by most Internet users reduce demand for bandwidth automatically in the face of congestion.

The substitutes for the Internet are far from perfect. For some of its uses, such as file transfers, telephone, satellite lines or even the surface mail service (using magnetic storage) can be used. However, for most purposes they tend to be slower and more expensive, albeit yielding higher security. For the World-Wide-Web usage, there is no substitute for the internet. Although, many applications on the World-Wide-Web have real-world substitutes, such as television, newspapers, auctions, and advertising. This suggests that price-elasticity of demand would be fairly low at the macro level. One caveat is that at the level of the information content there is much space in terms of quality, which would allow a reduction in quality and file-size if transport prices would rise, thus increasing price-elasticity. An example of this mechanism is that most websites with elaborate Flash-introductions will allow users to skip it to increase download speed in case of congestion. Moreover, most users have many substitutes for each individual provider because by the nature of the internet there are many routes which an information packet can take in its transport between two nodes. Consequently, price-elasticity can be expected to be high for individual providers.

2 Structure:

2.1 Horizontal concentration:

In line with the analysis above horizontal concentration can be analysed for the three hierarchic levels of the internet. Breaking up the analysis to three levels is most informative in terms of competitive potential because the types of investment and service vary at each level and so do pricing mechanisms¹. An added feature of the horizontal concentration is the geographical nature of the product. Considering that we are analysing a single network instead of competing networks, the relevant question to ask is how much choice each customer at each level has in terms of providers. The answer is inherently linked to geography, as the network is physical and hence, given that costs of interconnection outlay tend to increase with distance, the effective choice of a customer is limited by geographical proximity.

At the level of backbones the market tends to be fairly concentrated for the larger economic units such as the United States and Europe. In both markets there is (a remnant of) a publicly funded backbone, such as Ebone and NSFNET. In addition, there tend to be a small number of private sector backbones, such as Altnet, PSInet, and SprintLink in the US (MacKie-Mason and Varian, 1994:4). In smaller countries, which are not part of these economic units, there may exist only one national (public) backbone, or none at all, with all national capacity being supported by lower bandwidth lines such as the telephone network, which is then connected to the Internet via other countries.

At the regional level concentration appears to be higher, with only 12 regional networks in the entire US in 1994. These were mostly state-run or managed by coalitions of academic institutions. Since then, some private regionals have been realised by telephone and cable companies.

At the level of the ISPs horizontal concentration is lowest. In developed internet markets such as Europe and the US most consumers will have a choice between at least 10 ISPs. Most of these will require the user to connect via either the telephone or television cable network and provide a connection a regional from there. Businesses located in business parks tend to have a wider choice of connection method because ISPs lay their own fiber-optic lines there. Clearly, this can be entirely different in under-developed markets in developing countries, where only one national ISP may be available.

The tables below give a sense of the level of horizontal concentration in the UK and US.

¹ For pricing mechanisms read sections 3.3 and 3.4.

Top 10 U.K. ISPs by Subscriber²

data compiled by [John Lewell](#)

Rank	ISP	Subscribers
1	Freeserve	2 million active
2	World Online	1.4 million
3	AOL multiple brands	1.2 million (est.)
4	LineOne	1.1 million
5	Virgin Net	500,000
6	Genie	450,000
7	BT Click	440,000
8	BT Internet	425,000
9	Claranet	350,000
10	Thus	246,000

Top U.S. ISPs by Subscriber: 2001 Year End³

Rank	ISP	Subs. (millions)	Market Share
1	America Online (Dial-Up)	27.7 *	19.4%
2	MSN (Dial-Up)	8.0	5.6%
3	United Online (Dial-Up) [NetZero + Juno Online]	5.6 ±	3.9%
4	EarthLink (Dial-Up)	4.8	3.4%

² John, Lewell, <http://www.isp-planet.com/>

³ Patricia Fusco, 11/02/2002, <http://www.isp-planet.com/>

5	Prodigy (Dial-Up) [Includes SBC Narrowband]	3.6	2.5%
6	CompuServe (Dial-Up) [AOL Owned]	3.0	2.1%
7	Road Runner (Cable) [AOL Owned]	1.9	1.3%
8	AT&T Broadband (DSL)	1.5	1.0%
9	AT&T WorldNet (Dial-Up)	1.4	1.0%
10	SBC (DSL) [Includes Prodigy Broadband]	1.3	.91%
11	Verizon (DSL)	1.2	.84%
12	Comcast (Cable)	.948	.66%
24	Other U.S. ISPs	77.5²	55.0%

HHI-UK: 1462 (Equivalent no. firms: 7)

HHI-US: 595 (Equivalent no. firms: 17)⁴

According to the HHI-indexes the UK market is fairly competitive and the US market is highly competitive.

2.2 Vertical integration

In terms of vertical integration at the internet hierarchy level ISPs span the range. In the US, Sprint is fully integrated offering dial-up services via its telephone lines up to commercial backbone connections via SprintLink. At the next level, ANS provided traditionally inter-regional connectivity to regional networks sponsored by the NSF. There even exist 'reselling' ISPs with no physical infrastructure at all, leasing all capacity from other ISPs, in effect being little more than marketing firms.

Whether this is an equilibrium is doubtful. For the time being it may be sustainable because of the nature of interconnection-agreements. That is, despite cost structures being completely different across different types of ISPs, neither type may have a decisive competitive advantage. For example, Sprint may have a far lower

⁴ To compute this index I assumed that 'Other' is equally divided among ISPs with a market share equal to the smallest reported ISP (number 23, with 0.07%). Also, market shares due to different ISPs with the same owner have been added.

incremental cost for transport provisions than ISPs, who lease lines, because of its vertical integration. However, it also has much higher sunk costs due to infrastructure investments. If access-agreements change the balance may tip in favour of vertically integrated ISPs through their enhanced ability to guarantee connectivity (Srinagesh, 1996).

In terms of vertical integration beyond the Internet infrastructure, it appears to go mainly in one direction. Forward integration into World-Wide-Web content is common if not standard among ISPs, albeit to different degrees. AOL is as much an ISP as a content provider. For telephone companies the direction of integration was technically backward, as most expanded to include ISP services. Further backward integration does not, however, seem common. Notably, Cisco corporation, the world's largest producer of internet infrastructure inputs, such as routers, is not engaged in the provision of data transport facilities. As long as the market for inputs is sufficiently competitive and liquid the benefit of backward integration may be small. Additionally, entry barriers in terms of technological capabilities may be prohibitive for ISPs. Conversely, there are clear advantages in integrating connection and information content services for instance in terms capabilities for advertising and targeted content provision.

2.3 Product differentiation

In practice it has been difficult to differentiate the product as defined in section 1.1 due to technological restrictions. Currently, the TCP/IP protocol does not allow prioritisation of packets in a comprehensive and individualised way. That is, transport services cannot be deliberately changed for an individual packet according to date, location or time or the state of the network. This leaves product differentiation to the specific inter-connection agreements between ISPs and to the 'last mile'. Consequently, there is a host of different possibilities of connecting to the Internet, from mobile and fixed dial-up accounts, to faster ADSL or broadband lines via telephone lines or television cables to high-speed T1-T3 connections offered mainly to organisations. An additional form of product differentiation comes from forward integration into content offering. As most options tend to be available from multiple ISPs or figure as substitutes, this differentiation has not led the market to transform from a (oligopolistically) competitive to a monopolistically competitive market.

2.4 Diversification

As with vertical integration there is great heterogeneity among ISPs in terms of diversification. AOL-Time Warner is an example of a diversified company stretching from internet connections to film production. At the other extreme there are the resellers with

only a single activity. Whether cross-subsidisation has facilitated anti-competitive behaviour is difficult to say but the continuing existence of resellers on the market does not point towards a successful employment of such a strategy.

2.5 Barriers to entry

There are barriers to entry at each level of the Internet hierarchy. Most clearly, at the backbone and regional level the investment costs are enormous. This may explain and justify initial state involvement and strict monitoring of anti-trust policies in case of consolidation of the industry. This involvement should concentrate on the nature of interconnection agreements and access-agreements governing the rules for utilising an existing network by a competitor.

At the same time, there may be a problem of under-provision. At the level of ISPs catering directly to final customers there are two barriers. Firstly, switching costs for potential customers have traditionally been said to be high. As with bank accounts, email accounts are addresses which may be expensive to change. This may have changed however with the rise of non-provider related email addresses, such as Hotmail. At the other end, regional networks and backbones may block traffic from resellers if they are not compensated for it, as tends to be the case under current interconnection agreements.

3 Conduct:

3.1 Short Run - Long Run

At the level of backbones and regionals, capacity is fixed over the short-run because building the physical network is time and resource intensive. However, as mentioned above because the largest element of the cost of laying cables is the labour cost, ISPs tend to commit to excess capacity in terms of bandwidth on the fibre-optic cables. (Srinagesh, 1996). Consequently, the true barrier to increases in capacity in the short-run is router-capacity and the incremental cost of lighting up lines. For the 'last mile' there may be a trade-off with other services such as telephone or television capacity, which limits capacity overall but allows some flexibility in the short-run with respect to bandwidth for data services. Consequently, we may speak of a very short run, in which capacity is completely fixed; a medium run, in which some additional capacity can be generated at relatively low incremental cost; and a long-run in which new capacity can be generated by expanding the physical network.

3.2 Individual vs Collective

Much of decision making on the internet is at the collective level. This is a direct consequence of the fact that we are dealing with a single network comprised of many individual agents. The principal objective of collective interaction is to guarantee connectivity and interoperability. In principal, this can be done via a large set of bi-lateral agreements between backbones, regionals, and LANs. In practice, it is a combination of multi-lateral agreements at the national level, such as in the US, and the international level, and bi-lateral agreements between networks. This need for collective action can generate barriers to entry and may facilitate anti-competitive behaviour. On the other hand, the need to be interconnected for any single operator does yield power to coalitions of operators to restrain the ability of operators to leverage local monopoly power to gain a foothold in other areas of the network. Taken the access-prices as given, pricing decisions and capacity choice are individual choices for ISPs.

3.3 Access policies

Inter-connection agreements form the core of the access-policies on the Internet. As such, they are currently one of the most important factors shaping the structure, conduct, and performance of the Internet industry. This, because they shape the nature of barriers to entry and costs of different types of ISPs.

Fundamental in shaping the structure of the industry is the way in which the parts of the Internet are connected. As should be clear from the discussion above, there can be many routes a packet can take between two nodes. Consequently, connectivity can be guaranteed by bi-lateral agreements only, or even by a single multi-lateral agreement connecting all sub-networks at one central point, or a third party administrator (see Bailey 1996 for details on different types of interconnection agreements). The practice is one of a mixture between the two but clearly there is a question of optimality of the level and type of interconnectivity (see also Srinagesh, 1996). Moreover, the type of interconnectivity affects the cost-structures of ISPs and thus their competitiveness profoundly. For example, a multi-lateral agreement with a central connection point such as the Federal Internet eXchange (FIX) located on the East and West Coasts of the United States, leaves smaller networks with the same connectivity but lower sunk costs than the larger networks. Yet if such an option is unavailable, backbones may have disproportionately large pricing power in bi-lateral agreements.

Traditionally, at least in the US market, there is a mix between bi-lateral and a multi-lateral agreements. In both cases, the system is one of 'bill and keep', which means that at the connection point all traffic is reciprocally routed without charges being applied. For the multi-lateral case, there is a membership fee for the central connection point.

The strategic importance of interconnection agreements can be demonstrated in terms of a model on interconnection quality between rival ISPs. A simple model due to Crémer, Rey, and Tirole set out in Roson (2002:2-3) analyses this issue in a duopoly setting where the user's willingness to pay for access depends on a weighted average of the size own network and the other's network adjusted by an interconnection parameter. The networks compete in Cournot and first set the interconnection quality, with the lower value determining the actual quality. When network sizes differ, the larger network can maximise profits by setting a lower interconnection quality. The intuition is that lowering the interconnection quality reduces the value of the combined network for its users but differentiates the own network from the rival's in a positive way. Whereas for the smaller rival, reducing the interconnection quality makes their network less attractive relative to the other network. There is no reason to believe *a priori* that in face of such strategic behaviour, the correct interconnection quality, from a social point of view, will be selected.

3.3 Pricing

Pricing of data transfer is of great importance if the quality of the internet is to be preserved in light of rapidly growing demand for bandwidth. As Srinagesh (1996) puts it:

The costs of providing IP transport represent a substantial fraction (25-40%) of an ISP's cost. This proportion will fall as less costly fast packet services are more widely deployed. However, the increase in the use of multimedia applications may result in a proportionally greater increase in the need for bandwidth. The tension between satisfying customers with bandwidth-intensive needs and satisfying customers with low-bandwidth applications cannot be efficiently resolved with current technology.

The pricing system must provide incentives for efficient allocation of the network capacity, which means it must yield an optimal level of resale and guide user behaviour in size and timing of data transfer requests. Usage-sensitive pricing schemes, such as a per-bit charge may provide reasonable incentives. Currently, however, flat rate pricing schemes are most common with exceptions of 'time based pricing' dial-up accounts. The problem with usage-sensitive pricing schemes is that the overhead cost is prohibitively high. Bailey (1996) cites research that billing overhead accounts for over 50% of current telephone bills. This percentage would much higher for the internet as each 200 byte package would need to be accounted for. A midway is to price bandwidth and let congestion do the allocation. Although it is efficient to have some congestion on the network (otherwise its capacity would be underused, given that its cost of expansion is positive), congestion is an inefficient way to allocate the capacity. There is, however, currently no technologically feasible way to allocate data transfer capacity in a comprehensive, price based way.

Alternatives to a pricing system are currently employed. Some policing takes place, based on collective or individual action in case of 'abuse' of informal internet codes. In addition, the TCP/IP protocol set its standard configuration such that bandwidth is rationed in an egalitarian way, adjusting the transfer speed of a user to the average speed in case of congestion. This does not however eliminate the congestion, nor are these measures likely to withstand the phenomenal growth of the internet.

So what would an ideal congestion pricing system look like? MacKie-Mason and Varian (1994:16-17) suggest a 'smart-market', operating as a real-time Vickrey auction. In this scheme, when the network is congested users put a bid on each packet reflecting their willingness to pay for immediate servicing. This bid is used, in stead of the current First-In-First-Out protocol, to prioritise packets. Each user pays the bid of the highest bid rejected packet, which yields efficient incentives for truthful revelation. If all revenues are invested in new capacity, this will expand capacity exactly to the point where marginal cost equals marginal revenue. Crawford (1996), suggests to implement a similar system but utilising a two-part tariff, with a fixed element for access and capacity and a variable element at marginal cost. The problem with these schemes is that it requires dynamic updating of prices, whereas in practice prices would have to be cleared

at discrete points in time. Moreover, it assumes a single entry point. In practice, either the auction would have to be at a central point, requiring all packets to move through this point, or there would need to be simultaneous auctions for different routes with some degree of coordination. Moreover, it is not clear who should be billed, the sender or the receiver, making the scheme more costly.

Figure 2: Reduced Capacity Increases Price and Revenue
D. W. Crawford 1995 March

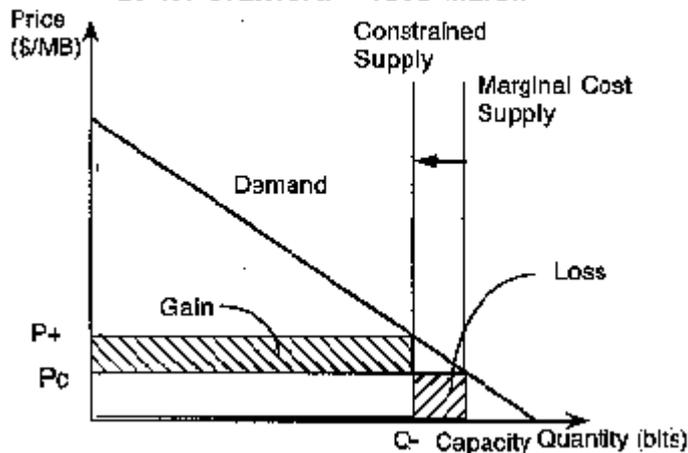
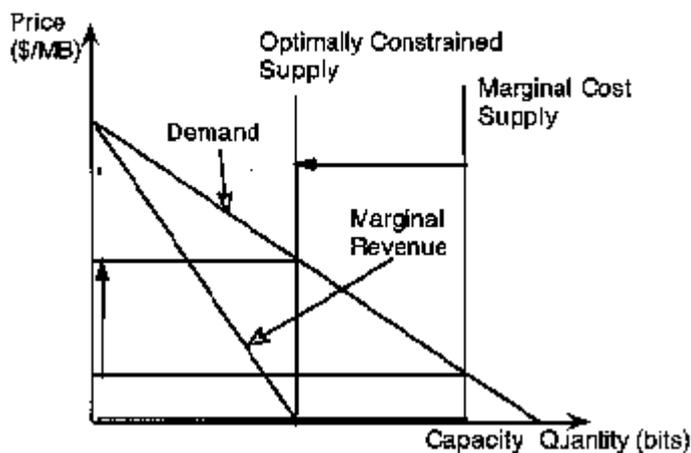


Figure 3: Monopolistic Solution by Constraining Supply
D. W. Crawford 1995 March



Also, the scheme would generate more revenue as capacity is limited, as is illustrated in the diagrams above (due to Crawford, 1996). This provides an incentive for strategic behaviour of the ISPs. Much like a monopolist reducing output, an ISP with market power would have an incentive to reduce capacity (by under-investing or artificially by increasing congestion) and thus increase congestion revenues (Crawford, 1996). This problem could be solved through the two-part tariff scheme in face of competition for bandwidth provision.

Despite their potential problems, these schemes are informative as to what direction the market should take in terms of pricing. Soon, technology will allow protocols to accommodate these types of schemes, and given the rapidly growing demand for

bandwidth usage-sensitive pricing will become necessary to avoid unacceptable congestion of the Internet.

3.5 Quality

As discussed under section 3.3, quality in terms of interconnectivity between rival networks can be a vital strategic aspect in access agreements. Quality is also an important factor in terms of product differentiation at the 'last mile'. For example, the number of available lines for dial-up strongly affects the connectivity of end-users and may yield a competitive advantage/disadvantage for an ISP. Beyond these levels, quality has fewer dimensions and is mostly a collective issue through overall network performance in terms of congestion and package loss.

3.6 Capacity choice

"In a competitive environment with excess capacity, there is a tension between the large sunk costs of physical networks and very low incremental costs of usage. On the one hand, the need to recover sunk costs suggests using price structures with high up-front charges and low (or zero) usage rates. On the other hand, with significant excess capacity present, short-run profits can be increased by selling at any price above incremental cost. Economic theory would suggest that the pricing outcome in this situation might be unstable, unless regulatory forces or other influences inhibiting competition were present.

The consequence of the leased line tariff structure described above for the cost of IP transport is straightforward. Given a high nonrecurring service order charge, ISPs with leased line backbones have an incentive to size their needs over a three to five year period, and commit to a level of purchase determined by projected demand. In a rapidly growing Internet, this can result in substantial excess capacity among ISPs in the short run. The incremental cost of carrying IP packets will be close to zero".

(Srinagesh, 1996)

4 Performance:

4.1 Private:

Growth of users and data traffic has been spectacular over the past years and continues at double digit monthly rates:

"Internet traffic in the US doubled during 2001, according to a new report from RHK. According to RHK, US Internet traffic now stands at around 100 petabytes, (100 million gigabytes) a month. This is more than twice the equivalent long distance voice traffic for the whole of the US.

RHK predicts that while dial-up is still the access method of choice for around 79 percent of US households, the number of residential broadband subscribers in North America, will reach 36.8 million by 2005⁵."

Similarly, ISPs have also grown rapidly:

"The average growth rate of U.S. ISPs between the second and third quarter of this year is 9.9 percent, based on the performance of the eight ISPs that remain. We consider conventional ISPs to include MSN, EarthLink, Prodigy and PeoplePC. The business model of these providers predominately reflects narrowband users in the U.S. Accepting this formula, dialup access grew 4.7 percent on average between the second and third quarters. On average, broadband use in the U.S. grew 14.4 percent during the same reporting period⁶."

In line with the above, the backbone market has grown fast, albeit slower than connections:

"Revenues for the overall backbone connection market will have a compound annual growth rate (CAGR) of 16.82 percent between 1999 and 2004 while the number of backbone connections will have a CAGR close to 23 percent during the same period, according to Cahners In-Stat Group (www.instat.com). The high-tech market research firm

⁵ http://www.nua.ie/surveys/?f=VS&art_id=905357928&rel=true, [Nua Internet Surveys](#), May 2002

⁶ Patricia Fusco

(http://www.isp-planet.com/research/rankings/usa_history_q32001.html, Nov. 2001)

finds that the difference between the growth rates in revenue and connections can be attributed to the continued decline in the cost of bandwidth. UUNet and Sprint still have the largest share of the backbone connection market⁷.”

These figures indicate that the Internet industry is capable of dealing with enormous growth in and changes in the type of demand. Also, it appears that, to a large extent, productivity gains due to technology advances are passed on to consumers in the form of price reductions. This suggests that the market is still quite competitive.

4.2 Social:

By the statistics below, we now have nearly 10% of the world population online:

World total⁸: 544.2 million

US & Canada: 181.23 million

In the US this percentage is far higher:

“According to the [CyberAtlas](#) there were 143 million Americans online in September 2001. Of these U.S. users, about 70 million are reported to be residential users. The study, compiled by the National Telecommunications and Information Administration ([NTIA](#)) , *A Nation Online: How Americans Are Expanding Their Use of the Internet*, estimates that 54 percent of the U.S. population has access to the Internet. This represents a 26 percent increase over August 2000”⁹

⁷ http://www.instat.com/pr/2001/is0007if_pr.htm, 2001

⁸ http://www.nua.ie/surveys/how_many_online/index.html, [Nua Internet Surveys](#), February 2002

⁹ Patricia Fusco, 11/02/2002, <http://www.isp-planet.com/>

5 Regulation:

5.1 Anti-trust:

There are no special regulations with respect to anti-trust cases for the Internet industry. The analysis above does suggest that because of the interdependencies of competitors through interconnection agreements there may be an increased risk of anti-competitive behaviour by dominant players. Additionally, the necessity for these agreements provides a natural forum on which to build illegal price agreements. Consequently, special attention if not special regulation may be in order.

5.2 Special rules:

In terms of taxation the World-Wide-Web appears to benefit from a special exemption from the US authorities. Although similar regulations do not apply to the underlying Internet industry, the general attitude in terms of industry regulation appears to be 'hands-off'. This is changing however, notably on the content side, where an increasing number of countries, including the US, are introducing bills regulating what should and should not be available to what type of user. The onus for compliance to these laws has been on the side of the ISPs but so far has been mostly limited to the obligation to close down sites that are deemed illegal. As the technological capabilities in terms of data identification and tracking improve, this type of regulation is likely to start affecting the underlying Internet market significantly. The potential future need to 'filter' information on a large scale is likely to generate significant increasing economies to scale and can thus be expected to affect the structure of the industry profoundly.

On the other side, the Internet industry has benefited from significant subsidies from the outset. Notably, its former central backbone, the NSFNET was state subsidised. Even now, a significant part of its infrastructure is still (indirectly) subsidised via state institutions such as NASA and academic institutions.

In general the international nature of the internet makes regulation on a national level complex and sensitive. Even if collective price fixing is avoided nationally, some coordination to achieve international agreement between various the national networks will remain necessary.

6 Conclusion:

The Internet is a typical example of a network industry, where both positive and negative network externalities play an important role in its structure, conduct and performance. The purpose of the Internet is to transport data packets between its nodes. This is achieved by connecting private and public parts of the network to each other and to high speed backbones. Internet Service Providers (ISPs) provide the connection between the workstations (final nodes) and the Internet for most final users. In the more developed parts of the Internet, coinciding with the more developed economies, the market for ISPs tends not to be highly concentrated. Vertical integration beyond providing internet services does not commonly extend beyond voice or television throughput or content provision. Together with limited scope for product differentiation the structure of the industry is conducive to competition. Entry barriers can be high in terms of investment costs for laying the network but whether this is a problem in practice depends on access and interconnection policies.

It is thus not in its structure but in its conduct that we find the key to this industry's competitiveness. Most important is the nature of interconnection and access agreements between providers. These may, but do not at present, erect insurmountable entry barriers. Moreover, they determine the cost-basis for ISPs, which determines consequent pricing strategies. Based on the analysis above, we can conclude that for the US market, access agreements are at present conducive to a competitive industry. However, once technological capabilities are enhanced to allow usage-sensitive pricing schemes, the nature of such agreements may change profoundly. At present, pricing tends to be lump-sum non-usage sensitive. In some cases it may involve variable prices for rented bandwidth capacity, although commonly it does not between ISPs. As a consequence of these pricing issues and because of strategic behaviour involved in access agreements, there is no reason to believe that quality and capacity will be set at socially efficient levels. In addition, the collective nature of the industry, although not apparently so at present, may become conducive to anti-competitive behaviour in the future. Therefore, some regulation and supervision may be in order.

Despite its limitations in terms of its resource allocation mechanisms and the network externalities to which it is subjected, the Internet industry has performed well in accommodating break-neck speed growth in terms of connections and volume of data transport without rendering the network unusable due to congestion. In line with the finding above, that the current structure and conduct foster competitiveness, the industry has passed most of its productivity gains on to consumers in terms of falling prices.

7 Bibliography

Bailey, Joseph P. (1996), "Economics and Internet Interconnection Agreements", *The Journal of Electronic Publishing*, May, 1996 Volume 2, Issue 1, <http://www.press.umich.edu/jep/works/BailEconAg.html>

Crawford, David W. (1996), "Pricing Network Usage: A Market for Bandwidth or Market for Communication?", *The Journal of Electronic Publishing*, May, 1996 Volume 2, Issue 1, <http://www.press.umich.edu/jep/works/CrawMarket.html>

Mackie-Mason, J. K., and Varian, H., (1994) "Economic FAQs About the Internet", *Journal of Economic Perspectives*, (Fall, 1994)

Roson, Roberto (2002), "Two Papers on Internet Connectivity and Quality", *Review of Network Economics*, Vol. 1, Issue 1, http://www.rnejournal.com/articles/roson_internet_mar02.pdf

Srinagesh, Padmanabhan (1996), "Internet Cost Structures and Interconnection Agreements", *The Journal of Electronic Publishing*, May, 1996 Volume 2, Issue 1, <http://www.press.umich.edu/jep/works/SrinCostSt.html>