Unclassified **DSTI/ICCP/TISP(99)14**

Organisation de Coopération et de Développement Economiques
Organisation for Economic Co-operation and Development

OLIS : 04-Nov-1999
Dist. : 05-Nov-1999

English text only

DIRECTORATE FOR SCIENCE, TECHNOLOGY AND INDUSTRY
COMMITTEE FOR INFORMATION, COMPUTER AND COMMUNICATIONS POLICY

DSTI/ICCP/TISP(99)14
Unclassified

Working Party on Telecommunication and Information Services Policies

INTERCONNECTION AND THE INTERNET: COMPETITION AND REGULATION ISSUES AT LOCAL ACCESS AND BACKBONE LEVELS

The following paper, prepared by Professor Martin Cave from Brunel University, United Kingdom, is provided as a background document for the TISP Workshop on Interconnection Policies and Frameworks to be held on 11-12 November, 1999 in Venice, Italy.

Contact: Mr. Dimitri Ypsilanti, Tel: (33 1) 45 24 94 42, Fax: (33 1) 45 24 93 32,
Email: dimitri.ypsilanti@oecd.org

83848

English text only

TABLE OF CONTENTS

Tables

Figures

MAIN POINTS

1. Whereas all aspects of the telecommunications industry have traditionally been closely regulated, the Internet in its early days was relatively free from regulation. As it has become commercialised and access to it has grown, regulation and competition issues relating to interconnection have become significant at two levels – the means by which subscribers gain access to the Internet through the PSTN, and interconnection among Internet Service Providers (ISPs) at all levels, including particularly Internet Backbone Providers (IBPs).

2. In relation to access to ISPs, regulators have traditionally set both retail prices (directly or indirectly) for services provided by the incumbent telecommunications operator, interconnection rates for call origination and call termination (as competition has developed).

3. In many countries, regulators have sought to introduce reduced retail charges for Internet access, and this has implications for the related wholesale rates.

4. In some countries, the regime for regulating Number Translation Services (NTS) has had the effect of enabling ISPs to acquire a proportion of call termination revenue. This has enabled them to develop "free ISP" services on the basis of appropriating some of the profits in telephone charges, which are then recycled to subscribers via the free service.

5. However, these arrangements are vulnerable to regulatory changes, particularly reductions in the price of Internet access via the PSTN. If these changes occur, the "free ISP" business model will become increasingly dependent upon revenue from other sources, particularly advertising and e-commerce.

6. Relations among ISPs were traditionally based upon peering, which involves the free exchange of traffic and is akin to the sender keeps all (SKA) interconnection regime adopted, in the past, for some international telecommunications traffic.

7. International Backbone Providers (IBPs) have increasingly sought to confine their peering arrangements to relationships with organisations of a similar size, requiring other ISPs to interconnect with them via more conventional paid "transit" arrangements.

8. This development has raised the issue of whether the co-existence of peering and transit arrangements is anti-competitive. This depends essentially upon the degree of competition among IBPs for transit business. If they set prices collusively, discrimination with respect to the payment regime could create barriers to entry and lead to excess profits by IBPs.

9. There is no evidence at present of such effects, but there is a reasonable basis for competition authorities to look with suspicion upon mergers who significantly increase the level of competition in the IBP market.

INTRODUCTION

10. It would be hard to think of activities with more dissimilar regulatory traditions than telecommunications and the Internet. The former has been continuously subject to detailed regulatory intervention, in many countries initially through the medium of public ownership, although latterly this method of control has been replaced by delegation to regulatory authorities with varying degrees of independence. In keeping with the Internet's origins in the research community, its brief but powerful tradition is, by contrast, one of very limited regulation, although questions of content regulation are subject to increasing controversy.

11. However, the two regulatory traditions have to confront one another, not least because the telecommunications networks and the 'Internet Cloud' are interconnected physically. This physical connection is seen in figure 1, which gives a schematic outline of the structure of the Internet. The key point is that most subscribers are connected to it via a local exchange carrier which links them, either directly or by interconnection with another carrier, to an independent access or service provider (IAP or ISP). Normally outside the PSTN, that ISP is linked to others all over the world via Internet backbone providers (IBP), which themselves interconnect either directly or through common interconnection points. Inevitably, these interconnection arrangements among ISPs have close parallels with interconnection among traditional telecommunications operators. The tendency to apply similar models of commercial relationships and the temptation to adopt similar forms of regulation are strong.

Figure 1. The structure of the Internet

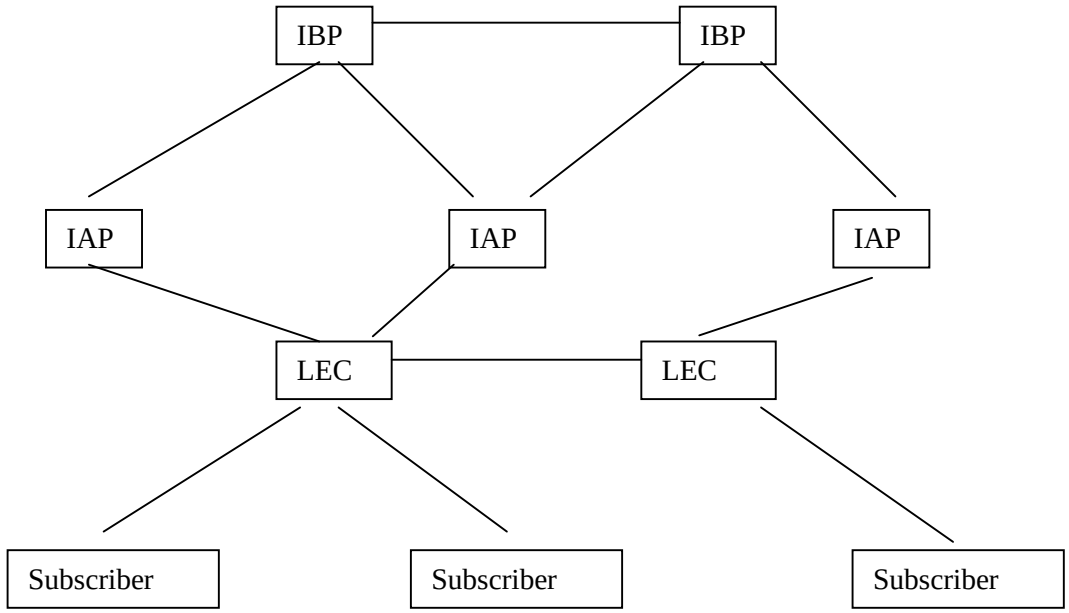

IBP – Internet backbone provider
IAP – Internet access provider
LEC – Local exchange carrier

12. The aim of this paper is to analyse competition and regulatory issues arising in connection with interconnection at two levels. First, how subscribers' access to ISPs is mediated through telecommunications tariffing and interconnection arrangements which were initially developed in a rather different context. The question here is: how do those regulatory interventions influence the strategy and structure of the ISPs?

13. The second issue concerns interrelations among ISPs. Most ISPs, in the United States in particular, are localised organisations with a relatively small number of customers – remember that in the United States there are approximately 7,000 ISPs, and the vast majority of them provide service in a small number of adjacent counties, or on a regional basis (see Bailey 1999). The activities of these ISPs, often called independent access providers (IAPs) are complemented by a group of much larger organisations, sometimes known as internet backbone providers or IBPs. These have some directly connected customers, but they also undertake the task of transmitting data throughout the world using fibre-optic cable, picking up traffic from other IAPs and exchanging it with other IBPs. The question to be addressed here is: given the likelihood of a hierarchy of ISPs, consisting of a small number of IBPs, a very large number of small IAPs and a group of ISPs of intermediate size, how will interconnection among them be organised, and to what extent is it likely that one firm, or a group of firms, will acquire or abuse market power?

14. Sections 1 – 2 of the paper deal with a first issue, and Sections 3 – 4 with the second. But before beginning the discussion it is useful briefly to recapitulate two traditional models of interconnection in telecommunications.

15. In the former, based upon monopolistic international telecommunications the regime relies upon a 50:50 split of some notional cost (the accounting rate) between the originating and terminating operators. When traffic is balanced, no net payment occurs. It would, however, be erroneous to infer from this that the level of the accounting rate makes no difference. To the extent that operators factor the settlement rate they have to pay to the terminating operator into the costs of providing an international service, their level will influence retail prices or collection rates (see Cave and Donnelly 1996).

16. Another mechanism, acceptable under ITU recommendations (although hardly used for PSTN traffic), is 'sender keeps all' (SKA). Under this regime, termination carries zero cost for the originating operator, and prices will be set accordingly, in accordance with whatever objective the operator is pursuing. Clearly this could in principle lead to an inefficient allocation of resources, as terminating costs are absent from the originating operator's pricing equation. In practice, this is not a problem, as retail prices are too high rather than too low.

17. In the context of international calls, there is limited scope for varying the point of hand-over of the call, since the originating and terminating operators are typically responsible for their own domestic transmission and for 50% of the cost of the international circuit. In other circumstances, however, SKA would naturally encourage hand-over at the first available opportunity.

18. The second context for interconnection in telecommunications arises in a competitive environment. Typically, an entrant competes for business with an historic monopoly operator. Interconnection with that operator is essential for the entrant, which at the start of competition has relatively few subscribers, and hence must pay the incumbent to convey and terminate a high proportion of calls. In such circumstances a refusal to interconnect or an excessive interconnection price would destroy the entrant's business. In view of the clear conflict of interest between entrant and incumbent, regulatory authorities typically have to intervene to set the terms and conditions of interconnection including price. Hence the development of the framework for setting interconnection charges contained in the US Telecommunications Act 1996 and EU legislation on interconnection. The dominant pricing rule adopted is based on long run incremental cost (LRIC), often with the inclusion of some mark-up to cover common costs. In the EU and elsewhere, accounting separation for network services is imposed in cases where an operator has significant market power, to prevent discrimination.

19. It is now established that regulators can abstain from regulation where there is sufficient competition. Thus in the UK, for example, more than 50% by value of BT's network services (not including access) are not subject to price regulation, Oftel having taken a view in 1997 that there was sufficient competition for this no longer to be necessary. This unregulated component includes long distance conveyance, the wholesale telecommunications service most akin to those provided by IBPs, where entry is relatively straightforward.

20. It is also recognised, however, that the mutual exchange of interconnection services among two or more operators creates a potential instrument for pre-commitment to collusive behaviour vis-à-vis a group of customers. Happily for consumers, price collusion often breaks down because the parties have an incentive to cheat, in the form of stealing business from one another. If, however, they have entered into interconnection agreements which commit them to the purchase from one another of inputs at above cost, these contracts will lead to higher retail charges and do not require any side payments from one to the other, if sales and purchases are symmetrical. This has the further advantage for them of hampering entrants, whose reliance on purchase network services would be initially be very great. Note, however, that SKA is apparently the opposite of this practice.

SECTION 1.

ACCESS TO THE INTERNET THROUGH THE PSTN.

21. In most countries, both the retail price of local calls supplied by the dominant operator and the wholesale price of call origination and call termination are subject to price regulation. This is considered necessary because of the dominant position held by the historical operator in the provision of local telecommunications service (including call origination) and the supposed bottleneck property of call termination. As a result, when PSTN subscribers seek access to their ISP, they do so within a framework of telecommunications regulations.

22. The simplest form that this access can take is when a customer on the dominant operator's network gains access to an ISP via that network and all call revenues are retained by the network operator (see **Figure 2a**).

Figure 2. Modes of Internet access

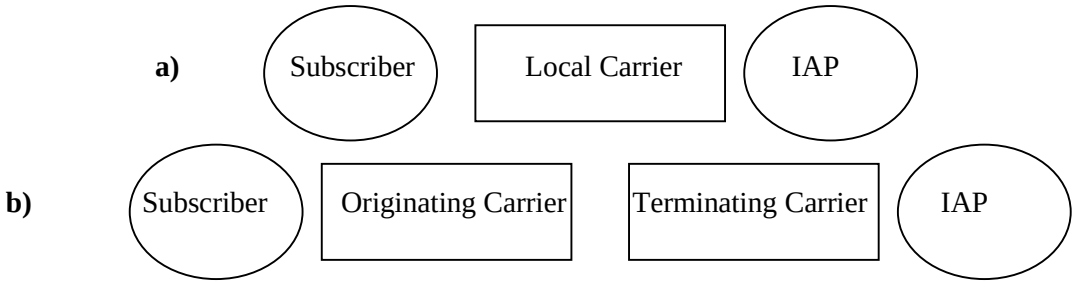

23. The charges levied by the network operator will be a standard rate mediated by whatever discounts may apply to the subscriber in question and by such things as minimum call duration. This retail price is probably subject to regulation, but the simplicity of the arrangement relies upon the fact that there is no interconnection with another exchange carrier, and hence no necessity to divide the revenue between two (or more) operators.

24. The situation that arises when this simplification is removed is illustrated in **Figure 2b**.[1] In this model, the subscriber's network operator passes the call to the ISP's operator, which provides a termination service. As before, the retail price is subject to regulation, but the further question arises of how that regulated call revenue is divided among the parties (the originating operator, the terminating operator and the ISP). The key issue is whether the originating operator (as would be the case with a normal call to another subscriber) keeps the revenue and pays a call termination charge to the operator to

which the call is passed, or whether the roles are reversed so that the revenue accrues to the terminating operator which then pays an origination charge to the subscriber's operator.

25. The second option may come into play in Internet access, because access to ISPs is often achieved via Number Translation Services. These are a range of specially tariffed services primarily used for telemarketing, including free phone services, local call fee access, national call fee access and premium rate services. Customers are able to recognise from the relevant number (0800 etc.) the charge they are paying for the call in question, whatever the location of the called party.

26. Much of the debate about Internet access has arisen as a result of the method for dividing revenue associated with Number Translation Services adopted by Oftel in 1995. The formula is as follows (Oftel 1999) :

- The originating operator retains: P – D + C
- The terminating operator receives: D – C

Where:

- P is the actual retail price charged by the originating operator to the customer;
- C is the pence per minute charge for conveyance over a single tandem segment of BT's network, including an uplift to allow for retail costs incurred by the originating operator in handling the calls;
- D is the Deemed Retail Price for the call. In the case of free phone services this is 0. In the case of local and national call fee services, it is the retail price adjusted by discounts and any effects associated with BT's minimum call charge.

27. The key difference between the two procedures concerns the identity of the residual legatee of the retail charge. The issue acquires significance because call charges will typically embody a significant surplus over marginal cost (where that marginal cost includes, as appropriate, the origination or the termination charge). The formula chosen thus determines the disposition of that margin.

28. A variant of this model arises where the ISP adopts a freephone number. The originating operator is entitled to retain the call origination charge determined by the regulator. The ISP then negotiates a payment to the terminating operator. However, the ISP can also derive revenue from its customer by imposing a per minute charge for use of the ISP services. This charge becomes a proxy for the call charge.

29. Two other arrangements also deserve mention. In the first, the dominant operator establishes a special retail price, and possibly a special code, for access to the ISP. Where the call is handed over to the ISP's operator for termination, the termination charge may be similarly reduced. If, in addition, the dominant operator itself operates a free ISP, implicitly subsidising it from the termination rate, it may come under pressure to make equivalent payments to other ISPs.

30. Finally, subscribers may gain access to the ISP through standard indirect access arrangements. This would mean that ISPs (either directly, or through a terminating operator) would have carrier-select access codes with the originating access operator, and corresponding interconnection agreements. The ISP would bill the customer directly (like any indirect access operator) and recompense the originating network operator. To the extent that the NTS call origination rate includes an uplift over LRIC for marketing costs,

indirect access may represent the cheaper option for an ISP, although the ranking in any particular case may be affected by the way in which each charge is averaged geographically.

31. The key regulatory feature that causes these various options to have different outcomes is the interaction between the regulated levels of the retail price and of interconnection charges for call origination and call termination. If retail prices are too high, rents are created which can then be redistributed through competitive or bargaining processes. The UK NTS system effectively assigns these rents in the first instance to the terminating operator. The process by which they are then transferred throughout the system is illustrated by the emergence of the free ISP business model, the viability of which depends upon the relative levels of the retail rate and the origination charge. Section 2 describes this model and identifies its points of vulnerability.

SECTION 2.

INTERNET ACCESS AND THE FREE ISP

32. In late 1998, Freeserve, a fully owned subsidiary of Dixons, a large electrical retailing firm, made available for the first time in the UK a free ISP service. Customers could simply take a free disk from Dixons stores or elsewhere, and load it into their PC.

33. This then gave them access to Freeserve via a local call charge (0845) number. Although Freeserve levied no charge for its basic ISP service, there was an additional charge for consultancy and advice, via a separate premium rate telephone line. Its earlier competitors, such as America On Line UK, a subsidiary of AOL Europe (AOLE), charged a monthly fee.

34. Freeserve quickly acquired over a million customers, and became the largest ISP in the U.K. Other ISPs were then set up, following its lead. These included services provided by newspapers, bookstores, supermarkets, consumer magazines and telecommunications operators such as BT. In August 1999, AOLE adopted the same model under a Netscape Online brand name.

35. In the Summer of 1999, a minority interest in Freeserve was floated on the Stock Exchange at a full market capitalisation of £1.6bn. After enjoying early rises, the share price subsequently fell, three months later, to below the issue price.

36. The underlying logic of the Freeserve model relies upon the appropriation of some of the rents available to the terminating operator under the UK NTS charging system described above.[2] After deduction of the originating operator's share of the local call charge (adjusted for discounts etc. as described above), the terminating operator receives about 40-70%, depending on the time of day. This is more than adequate to cover the costs of call termination. An opportunity is thus created for an ISP to negotiate with a number of potential terminating operators for a share of the termination charge. Following such a negotiation, Freeserve agreed that Energis would be its terminating operator, and it is reported that Energis would pass to it a proportion of the termination charge. The level of that termination charge is itself a variable percentage of a retail price which varies by time of day and day of week. Industry sources suggest that free ISPs are in a position to gain 15-25% of the retail rate in the UK. This would yield the following illustrative figures:

Table 1. Free ISP gains

	Daytime	***Evenings***	***Weekend***	
Retail prices (Eurocents)	6.0	1.5	1.5	
Terminating Operators Share	68%	53%	41%	
Termination Charge (Eurocents)	4.1	0.8	0.6	
Free ISP's share (Eurocents)	0.9– 1.5	0.2 – 0.4	0.2 – 0.4	
The numbers imply that the following hours (per month) of network access are required to generate for the ISP a monthly revenue of 5 or 10 Euros.				
Revenue				***Average***
5	6-9	22-37	22-37	11-18
10	12-18	44-74	44-74	22-36

* Usage in equal proportions in the three periods. This compares with average usage rate in Europe of about 10 hours per month.

37. The free ISP thus has the advantage of an additional income stream to make up for the lack of charges – the revenue from its terminating operator. However, it is clear that at plausible usage levels, that revenue will cover only a proportion of a standard monthly ISP charge. Hence new revenue sources are required,[3] and a free ISP, by attracting a large volume of subscribers, is better placed to increase its revenues from other sources, notably from advertising and commissions paid by firms selling goods via e-commerce. Clearly, however, these last two advantages accrue only to large ISPs of which there can only be a limited number, likely to be made up particularly of first movers in the market.

38. Given the short life of free ISPs to date, it is too early to speculate, except in very general terms, about their long-term survival or the resulting market structures. That future is, in any case, dependent upon regulation. The particular question addressed here is, accordingly: to what extent is it likely that changes in regulatory arrangements might undermine the free ISP's revenue from the terminating operator, and, if this were to occur, what impact would it have upon the Freeserve business model?

39. The previous section has argued that the free ISP business model, as originally developed in the UK and subsequently adopted in many other counties, is essentially a regulatory artefact: there are three regulated prices – the retail price and two wholesale prices for call origination and call termination. The first should, in principle, equal the sum of the latter two, plus a retailing cost. If it did, there would be little scope for rent-seeking behaviour by ISPs. As it does not, and in particular because the share of the retail price accruing to the terminating operator excludes that terminating operator's costs, there is an opportunity for ISPs to appropriate the rent.

40. However, competition in the ISP sector forces them to re-cycle it back to customers in the form of a zero ISP charge. From a consumer's point of view, this is obviously preferable to an alternative state of the world in which the rent remained either with the telecommunications operator or with an ISP. However, there are a number of ways in which changes in regulatory arrangements might offer a greater variety of benefits to consumers in respect of telephone and ISP charges, and these regulatory arrangements might undermine the free ISP model, although whether they do so will depend ultimately upon consumer preferences.

41. In the first instance, the regulator can allow greater flexibility in the wholesale termination rate. At present, the logic of setting the regulated price for call termination derives principally from the concern that the terminating operator controls a bottleneck, with the consequent risk of exploitation of monopoly power.[4] However, it is clear that the free ISP model operates by creating competition among operators to terminate the calls to an ISP. In these circumstances, provided the competition is effective, there is little justification for regulating termination charges for Internet access. The concentrations of termination business on a few ISP numbers is likely to make the business far more attractive than is the case with other PSTN calls, and hence attract competitors.

42. Within the UK NTS framework, this change could be implemented by offering the terminating operator a larger number of price points at which to set the retail price to the subscriber, leading to an equivalent change in termination revenues. By choosing a lower (even a zero) retail price the terminating operator would be re-cycling the rents associated with access to the ISP at the standard rate back to the consumer directly, rather than indirectly via the ISP. An alternative way of achieving the same outcome is via direct regulatory intervention. The regulator could simply reduce, or alter the structure of, either the retail price for Internet access, differentiating it from the retail price of voice or cutting all local call charges. The incumbent might even choose to do this of its own accord.

43. A straightforward reduction of retail prices for access to the Internet has been introduced in a number of member states of the European Union, and, implicitly, in the United States where the FCC has exempted ISPs, viewed as a special class of enhanced service providers, from paying access charges to local exchange carriers for use of their network.

44. In addition to the proposition that enhancing Internet access is a public policy objective, a number of more or less persuasive cost based arguments can be made for a lower price for Internet access, compared with voice. The first of these is based upon the lengthy nature of Internet access calls. Call origination rates, and retail prices, are often charged on a per minute basis, ignoring the fact that a significant component of the cost of call origination is associated with call set up. The data from the Oftel cost model suggest that call set-up costs per successful call amount to 1.22 Eurocents, while the per minute traffic cost is 0.83 Eurocents. Simply factoring the set-up cost into the per minute charge has the effect of over-recovery of costs on long calls and under-recovery on short calls.

45. Second, Internet access is susceptible to transmission by an alternative technology than voice – on a less expensive data network. It can also be argued that voice and Internet calls cost constitute different economic markets. In these circumstances, it might be legitimate to compute the cost of origination or conveyance of Internet calls on the basis of treating those calls as the increment, rather than the totality of calls.

46. Third, the structure of retail tariffs by time of day and day of week may not be optimal (in the sense that capital costs are recovered, according to peak-load pricing principles, at peak periods, and only operating costs at other times). If off-peak charges are too high, this is likely to bear down heavily on residential Internet users. In addition, if Internet calls are relatively price elastic, there is a case for a lower proportion of common costs to be recovered from them, as against other calls.

47. As far as amendments to the structure of retail prices is concerned, many regulators have promoted optional tariffs embodying substantial quantity discounts for calls, including Internet access calls, in off peak periods. These typically take the form of offering up to a specific number of off peak hours per month at a flat rate. Some operators have gone further by offering pure flat rate charges confined to off peak periods such as late evening and night or weekends. These charges can be construed as implicitly recognising that the retail peak/off peak tariff gradient does not adequately reflect costs. Although these developments do not go as far as the North American model of unmetered local service,

they are likely to have the effect of reducing the element of surplus in local call charges, and hence of undermining the free ISP model.[5]

48. These considerations suggest that the future is likely to hold a greater variety of charge structures for the combination of access to the Internet and services provided by the ISP. These are illustrated in **Figure 3**.

Figure 3. Tariff structures for Internet access and service provision

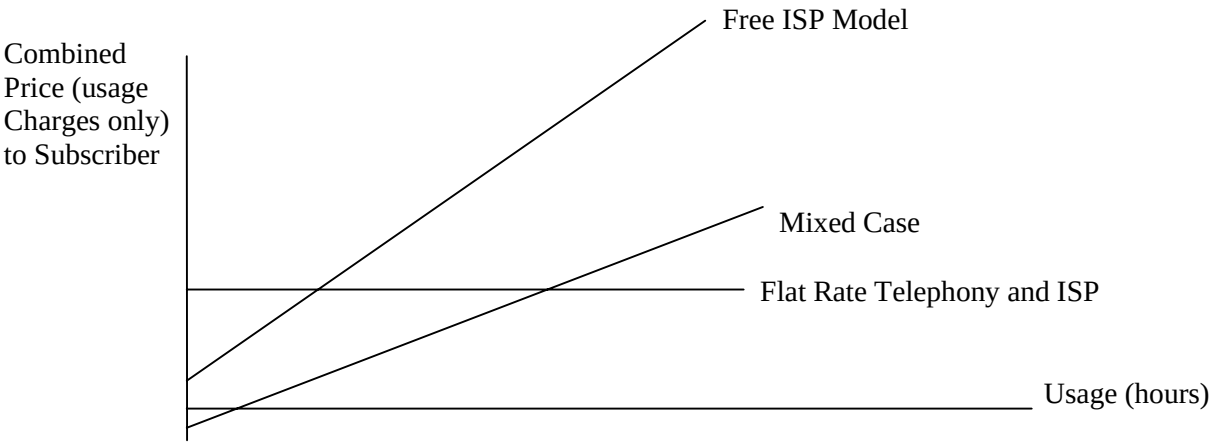

49. The extreme variants are the North American regime of flat rate for local calls combined with flat rate charges by the ISP. The opposite is the free ISP system, where the only charge is the per minute telecommunications charge. Intermediate variants involve a combination of flat rate and per minute charges. The figure fails to do justice to the full complexity of the situation because it does not take account of different charging structures for different times of day, or of the purchase of access in blocks (e.g. X hours per month for up to Y Euros per month).

50. When confronted with such alternatives, users are likely to sort themselves into groups, choosing the tariff that minimises their costs. Clearly low users will be particularly attracted to a charge based solely on a per minute basis. The crucial issue for the survival of the free ISP model is thus whether, taking into account its likely clientele, its cost of operation and revenues from alternative sources, such as advertising, it is a sustainable model. This depends crucially upon the proportion of ISP costs that are independent of minutes of access to the ISP. The greater these are, the harder it will be for a free ISP to survive on the basis of revenues (from termination charges, advertising and e-commerce) which are effectively minute-driven. The per minute charge will limit online time and restrict e-commerce and advertising revenues, as well as revenues from termination.

Conclusions

51. The key point to emerge from this discussion concerns the impact which regulation has upon the pricing structure for Internet use, including the telephony charge and the ISP payment. In particular, it has become apparent that a combination of high local call charges and the UK NTS regime drives charges in the direction of the free ISP model, which lacks transparency and limits price competition.

52. The free ISP model obviously appeals to some customers, but regulation should be directed to encouraging a range of pricing models, including some form of capacity charging for the telecommunications service. A public policy objective of encouraging Internet access can then be superimposed upon a range of alternative cost-reflective charges.

SECTION 3.

BACKBONE INTERCONNECTION ARRANGEMENTS

53. The second issue this paper discusses concerns traffic exchange among Internet Service Providers (ISPs). One of the key questions which has emerged is what continuing role there is for the system of peering among ISPs - an arrangement whereby traffic is exchanged among them without corresponding payment. The issue came to a head in recent years when the largest providers of the Internet backbone withdrew from their peering arrangements with smaller ISPs, replacing them with a regime for charging for access similar to that adopted in voice telecommunications. (For a detailed study of Internet traffic exchange and discussion of policy conclusions, see OECD 1998).

54. The resulting differences in treatment of ISPs have generated calls for regulation at the sector level, and concerns about anti-competitive conduct. Issues of competition and concentration in backbone provision have also come to the fore in the consideration of mergers between backbone providers, notably the MCI-WorldCom merger approved in 1998. The proposed take-over of Sprint by the merged entity has made the issue topical again.

55. **Figures 4** and **5** demonstrate two possible alternative communication structures for ISPs. In the former, or centralised model, communications are fully centralised: local IAPs feed into regional/national IAPs which themselves send data upwards, or receive it downwards, from global IBPs. These last communicate with one another entirely through a shared central node or nodes. As a result, most communications have to go all the way up before coming all the way down (the analogy with information flows in a centrally planned economy is instructive).

56. The opposite extreme (like an ideal-type perfectly competitive market) is one in which all relationships are horizontal. This would mean every ISP having a relationship with every other. This would, however, require n(n-1)/2 connections which would hardly be economic. As a result Figure 5 illustrates an intermediate organisational structure characterised by the following:

- At the backbone level, IBPs communicate with one another directly (private peering), as well as through the central node.
- Among the ISPs, some decide to communicate directly with one another, rather than via the backbone (when this relationship takes a particular form, it is known as secondary peering).
- More generally, smaller ISPs can circumvent the backbone by a sequence of lower level communications.

57. In practice, the organisation of the Internet is currently characterised more accurately by Figure 4 than either by Figure 3 or by the wholly horizontal option, but a distinction between top tier IBPs and other ISPs remains. One possible candidate for drawing the line between IBPs and ISPs is the type of interconnection agreement employed.

Figure 4. Centralised variant

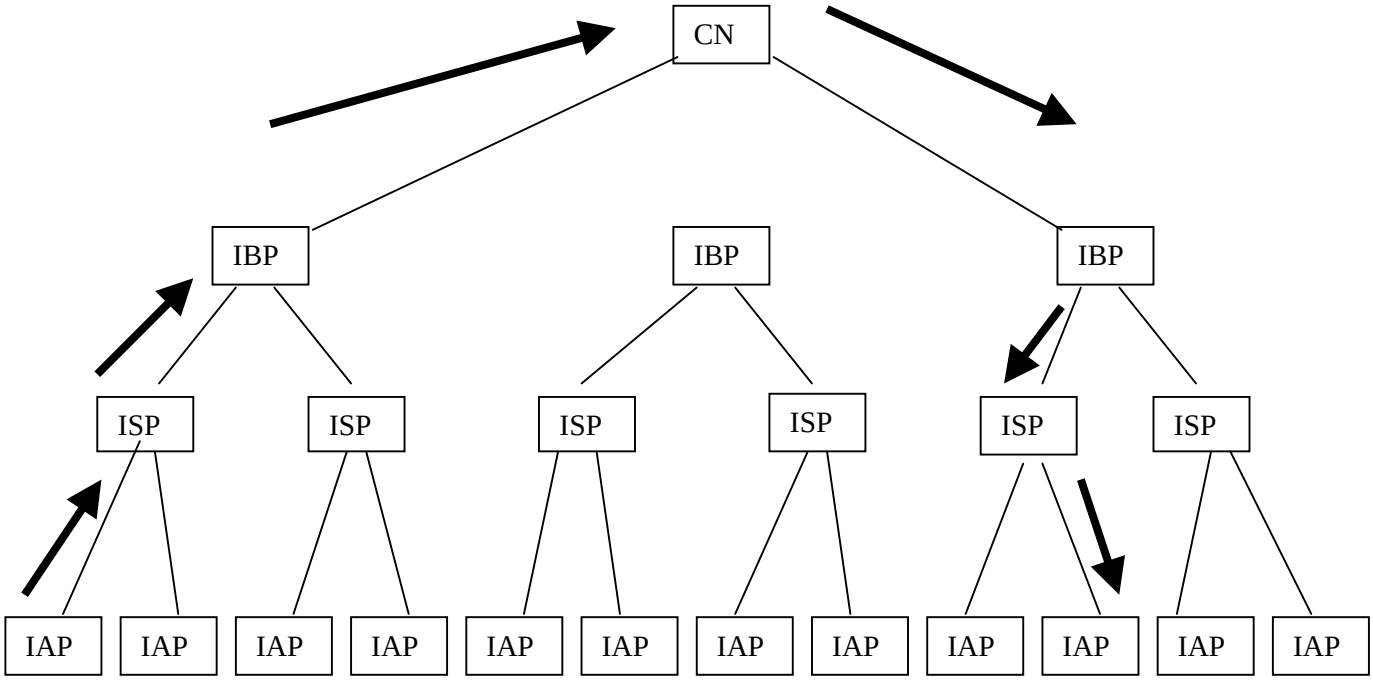

Key: CN = Central Node

IBP = Internet Backbone Provider

ISP = Internet service provider

IAP = Internet access provider

Figure 5. Looser variant

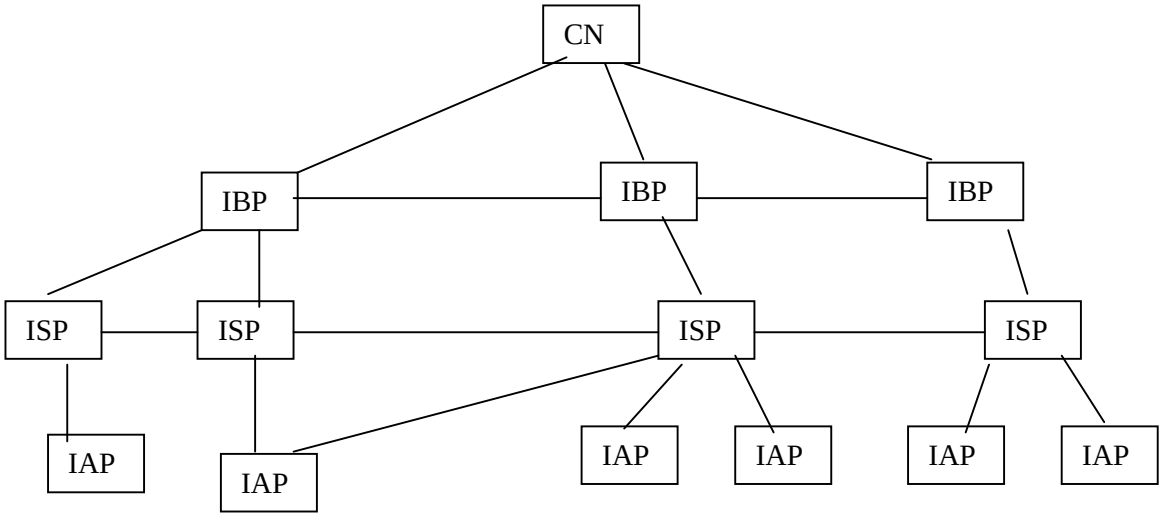

Method of interconnection among ISPs

58. In order to gain an understanding of the different forms of interconnection in the Internet, it is useful to have a simplified account of Internet routing (see Milgrom et al 1999). The Internet consists of a set of nodes connected by physical links with routers

59. at each node switching packets in accordance with routing protocols. At each node the router reads the destination address and forwards the packet onwards in accordance with the ISP's routing table. A routing table can vary from great complexity, including virtually all nodes in the Internet, to relative simplicity, including only provision for packets addressed to adjacent nodes, with a default destination to which all other packets are directed. Interconnectivity of the system as a whole is assured through the existence of core routers which contain information about all destinations and do not point to a default router. Core routers communicate with one another by exchanging routing tables and determining optimal routes. Where many core routers are involved, this is a highly complex interactive process.

60. ISPs can determine their relationship with one another by designating preferred routes for their outgoing traffic and refusing to accept certain incoming traffic. The options one ISP has in relation to another include 'peering' (a mutual exchange of traffic with an equivalent ISP) and 'transit' (the sale of conveyance to another ISP on a common carriage basis). I now define these two terms more precisely.

61. The equivalent in Internet interconnection of a sender keeps all type of arrangement is known as peering. Cukier (1998 p. 3) describes peering as 'an interconnection of two public networks that provide connectivity to hosts whose routes are advertised on the global Internet, on a settlement free basis that allows customers of one network to exchange traffic to customers directly on the second ISP's network'. Kende and Oxman (1999, p5) elaborate as follows:

> "As a system of interconnection peering has a number of distinctive characteristics. First, peering partners only exchange traffic that originates with the customer for one backbone and ten minutes with the customer of another backbone. As part of a peering arrangement, a backbone would not act as an intermediary and accept the traffic of one peering partner and transit this traffic to another peering partner. Second, peering

> partners exchange traffic on a settlements free basis, also known as bill and keep or centre keeps all. The only cost involved in peering is that each partner pays for the equipment and transmission capacity needed for the two peers to meet at each peering point. Third, peering partners have adopted what is known as 'hot potato routing'; they exchange their customers traffic destined for the other backboner at the nearest possible exchange points. Finally, the recipients of traffic only promise to undertake 'best efforts' to transmit data to its destination rather than guarantee any level of performance in delivering packets received from the peering partner."

62. The alternative to peering is a transit arrangement:

> "There are two main differences between transit and peering. The first is that, in a transit arrangement, one backbone pays another backbone for interconnection and therefore becomes a wholesale customer of the other backbone. In return, unlike peering, the backbone selling the transit services will read traffic from the transit customer to its peering partners as well as its other customers."

63. Elaborating these arrangements more fully, Bailey (1997) distinguishes four forms of Internet interconnection – peer to peer bilateral, hierarchical bilateral, third party administrator and co-operative agreement.

64. The *peer to peer bilateral* model typically involved two Internet networks of equal size forming a contract to exchange services with one another, to their mutual benefit. The gains each make in terms of network externalise are symmetric, because each offers the others an equal customer base and equal extension of services. For this reason, peer to peer bilateral relationships are often conducted on the basis of a barter or sender keeps all arrangement. Peering directly with an equal has the advantage of eliminating the use of intermediaries, as illustrated by the contrast between Figure 4 and 5 above. The traditional criteria for similarity among ISPs are: backbone capacity, number of access points, volume of traffic, quality of service and network reliability, (for examples, see OECD 1998, pp 21-22)

65. In a *hierarchical bilateral* relationship, the parties involve differ in terms of size and experience. Because the gains are not equal, the contract involves payment from one party to another, the payment possibly taking many forms, including capacity based, usage sensitive or even flat rate charging.

66. The *third party administrator* model involves more than two networks exchanging packets at a point operated by a third party. The third party routes traffic among the interconnecting networks and has to be trusted. The administrator charges the networks for its services.

67. A *co-operative agreement* also involves several parties, but is operated by the parties themselves in concert rather than independent firm. This was the historic model of the Internet, in the period when the participants were not a profit-making organisation. The agreements are harder to sustain, however, in the presence of competition.

68. Until recently, the norm for Internet interconnection has been third party administration or co-operative agreements, although bilateral peering was also common. There were, however, significant developments in 1996 as Cukier (1998) reports:

- During the winter of 1997, UU-net, MCI and BPN left the CIX (third party administrator) router.

- Between March and May 1997, UUnet told around 15 ISPs that their paying arrangements would be terminated within a few months, and that new bilateral transit agreements must be struck.

- Throughout the year, large ISPs converted their peering arrangements away from large exchange points in favour of direct interconnection with other networks.

- Although IBPs were eliminating many peering arrangements, peering grew more prevalent among smaller networks and was boosted by the creation of more local network access points.

- A national network called Vario Inc. built up by merging with regional networks became a major backbone provider.

- Vertical integration was also occurring within the industry, widening the margin between IBPs and other ISPs. GTE bought PPN and UUnet's parent company WorldCom agreed to merge with MCI. In November 1997, according to Cukier, the US's four largest networks (UUnet, MCI, BPN and Sprint) controlled between 85% and 95% of total backbone traffic.

69. Unfortunately, the precise state of affairs is shrouded in a certain amount of mystery because IBPs typically impose non-disclosure agreements with the ISPs for whom they provide a service. But peering among equals and transit payments from smaller to larger ISPs is the norm.

SECTION 4.

PUBLIC POLICY OBJECTIVES INTERCONNECTION ARRANGEMENTS AMONG ISPS AND REGULATION

70. Almost all OECD governments have as one of their major objectives the development what is referred to as an information society or knowledge economy, relying for the generation of wealth upon intangible capital, including information and human capital. One of the key performance indicators employed is the penetration of Internet access, which in Europe and elsewhere lags behind the United States. The ambition to catch up has led to considerable debate about flat rate charging for telecommunication services (a matter which is normally addressed by national regulatory authorities and the national industry) and about the consequences of the US representing the dominant repository of information. In particular, there have been sharp complaints (not further discussed here) about arrangements for sharing the cost of international Internet transmission, with non US operators arguing in favour of what they regard as more equitable cost sharing.

71. Underlying this quest for competitive advantage is the notion of network externality now familiar in the context of telecommunications and other sectors. (Note that many of these could in principle be internalised, but there are often severe obstacles to achieving this in practice.) A direct externality arises in telecommunications through increased connectivity. A subscriber coming on to the network benefits both herself and previous subscribers who can now contact her. An indirect form of network externality arises when a subscriber purchasing a service (say, a bank account) increases the incentive for firms to provide a complementary service (say, a denser network of automated telling machines). The Internet clearly exhibits both direct and indirect network externalities. The first is illustrated by email. The second by the burgeoning of websites as more people get on the net.

72. Network externalities are enlarged through interconnectivity. In telecommunications, this is typically imposed through a licensing regime which requires any operator of whatever size to interconnect with any other operator.

73. The private interest of firms is not, however, always in favour of extending externalities. Thus a large bank may not wish to make its ATMs inter-operable with those of smaller rival at zero cost, just as a large telecommunications operator will not wish to offer an entrant's subscribers access to its network, as doing so eliminates an important quality difference. There is thus a significant risk that the largest firms will choose to drive out, rather than interconnect with, smaller rivals. The logical extension of this process may be a monopoly in which all customers are interconnected to the single surviving network. However, monopoly pricing will limit access to or use of the network.

74. The incentives for IBPs to interconnect with IBPs in the context of oligopolistic competition are more ambiguous. On the one hand, each large IBP will have a desire to increase the cost of providing service in order to increase profits. On the other, large firms may wish to cut transit prices in order to steal business from one another. On top of this, however, is another factor peculiar to the Internet. As Kende and Oxman (1999 pp 11-12) put it: 'In negotiating peering, one important bargaining chip is the number of customers that a backbone brings to the table. Therefore the larger

backbones will compete with one another to win transit customers in order to be able to negotiate peering relationships with other backbones.

75. The regulation of competition is discussed in the next section. First, however, it is useful to consider whether the form of Internet interconnection agreement – peering or transit – has any relevance, independently of competition issues, to other policy objectives. A recent paper by Milgrom, Mitchell and Srinagesh (1999) has concluded that, for transactions cost reasons, a hierarchical structure in which a few core ISPs peer with one another but sell transit services to smaller ISPs may be desirable. The chief reason is that universal peering encourages free riding, in the forms of default routing (achieving a premature and inefficient hand-over of packets from one ISP to another) and reluctance to invest to overcome problems of congestion (relying instead upon competitors providing alternative, uncongested routes). Limited peering arrangements can, however, be bolstered by relationships of trust - and/or relatively economical monitoring).

76. The events described in Section 3 above, involving the termination of peering arrangements by the largest backbone operators, have focused attention upon the possibility of anti-competitive conduct by IBPs vis-à-vis smaller ISPs. In particular, it has been argued that denial of peering distorts competition and places IBPs in a position where, individually or jointly, they can abuse their market powers. Proposals to deal with these perceived abuses actual or potential include:

- An obligation to peer imposed upon all ISPs or, alternatively, an obligation to negotiate peering with all ISPs if a peering agreement is concluded with any ISP.
- Regulatory setting of prices for transit.
- Prohibition of vertical integration on the part of IBPs into related activities.
- The prohibition of exclusive dealing by a vertically integrated IBP.
- Transparency (i.e. compulsory disclosure) of transit peering agreements.

77. In addition to these proposed behavioural constraints, arguments relating to the structure of the IBP industry have arisen in the context of investigations into mergers between telecommunications, companies, heavily involved in backbone provision.

78. The analysis will assume that Internet backbone provision represents an economic market. In other words, it assumes satisfaction of the standard test for an economic or anti-trust market, which is to pose the following two questions in relation to a group of services:

> "If a single firm controlled the supply of those services, would it be able to raise the price above their competitive level by 5% to 10% for a non-transitory period - say six to twelve months? If yes, is there any smaller sub-set of the services to which the same proposition applies? If no, then the set of services is an economic market."

79. Crémer, Rey and Tirole (1999, pp14) note that there are five major backbones but 'an assessment in another source to the effect that 'somewhere between 6 and perhaps 30 other ISPs could be viewed as backbone ISPs'. They then go on to speculate how IBPs' customers would react to a 10% price increase. These customers would include both ISPs and directly connected customers such as websites, universities and business customers with dedicated services to the backbone. On the basis of their conclusion that there is little demand substitution at the end user level, and that it would be difficult for non-backbone ISPs to replicate backbone provision through secondary peering, they conclude that

there is a separate backbone market. The same conclusion was reached by the European Commission (1999 paras 58-77). The Commission also concluded that the geographical market was a global one.

80. Defining as a top-level network any network which peered, at a minimum, with the four biggest ISPs (WorldCom, MCI, Sprint and GTE/BPN), the Commission identified a total of 16 top-level networks. It estimated WorldCom's market share as falling between 35% and 45%, and MCI's as lying between 5% and 15%, and the share of its two nearest competitors as being between 15% and 25%. In combination this left about 30% for the remainder. This method, the Commission stressed, offered a relatively conservative estimate of levels of concentration in the market.

81. These concentration levels have to be evaluated in the light of barriers to entry. Although regulatory barriers to entry are largely absent, potential entrants might be reluctant to make the sunk investment in capacity in circumstances where they expected that the post entry price might fall to a competitive price which might condemn entrants at least to transitional losses, if they enjoyed fewer economies of scale than the incumbents. However, in an expanding market it is questionable whether investments are truly sunk, since they could be sold on to one of a number of surviving competitors.

82. But would backbone providers, even if partly insulated from entry, be able to raise prices above competitive levels? Clearly this depends upon the ease with which they can engage in parallel behaviour and their incentives to do so, which themselves depend upon a variety of traditional considerations, such as similarities in cost structure and the transparency of pricing in the market. In the latter connection it is worth noting that IBPs typically require ISPs purchasing transit services from them to enter into a non-disclosure agreement. From a competition point of view this may be advantageous as it enables IBPs correctly to cut their prices and increases the likelihood of 'cheating' on an agreement, open or tacit, to price in parallel.

83. Other concerns which have been raised relate to vertical relations. It has been suggested that an IBP could purchase a significant number of local IAPs and thus develop and exploit a monopoly. The same effect could also be achieved contractually by the IBP entering into exclusive dealing arrangements with local Internet access providers. But with the ISP market as apparently competitive as it is, and open to further entry, it is doubtful whether such a strategy would succeed. (See Bailey (1999) and Downes and Greenstein (1999) for an account of entry into US ISP markets.)

84. These arguments do not necessarily imply that competition authorities should be indifferent to increases in concentration of non-IBPs as the result of mergers. Mergers are both readily preventable and introduce a step change of competitive interactions which may have a major affect upon the nature of competitive interactions.

85. One argument which the Commission heard against the merger of WorldCom and MCI's backbone assets was that a dominant backbone provider might drive out its rivals through a targeted degradation of its quality of service. This would involve the dominant firm putting limits on its interface capacity with each of its rivals in turn, successively driving them out of the market or weakening them (see Crémer, Rey and Tirole 1999). Within the framework of EU competition law, the Merger Regulations provide a better basis for guarding against joint dominance than the application of Article 82 to an oligopolistic industry (Lasok 1999).

86. While the discussion above has considered the risk of anti-competitive conduct on the part of IBPs in general, it is useful finally to address the question of whether there is a particular feature of discriminatory peering which has an adverse affect on the competitive process. In particular, does the practice of the largest IBPs in confining peering to their own number either promote collusive behaviour or create a barrier to entry in respect of excluded ISPs? Clearly peering with other large IBPs involves a

higher degree of collaboration than might be necessary with an arms' length transit relationship. However, peering agreements are increasingly bilateral rather than multilateral. As far as preferential terms and conditions for peers rather than transit customers are concerned, the crucial issue is whether transit customers are able to purchase services at competitive cost based prices. The considerations listed above suggest (though not conclusively) that they are likely to be able to do so, so that peering IBPs would enjoy no long-term advantage.

87. If, however, IBPs did through tacit collusion or by other means succeed in raising prices above competitive levels, peering arrangements amongst them might mask price discrimination directed against ISPs which were not operating at the universal or top level.

CONCLUSIONS

88. This section has reviewed some of the recent debates concerning the need for intervention by either competition or regulatory authorities in the arrangements which have developed for Internet backbone provision. During the Internet's pre-commercial days, this was not a problem. But now that the Internet backbone services are provided by a number of large commercial firms, with combined billings of many billion dollars per year, it is natural that question of regulation should come to the fore.

89. The data services in both US and the EU have, in contrast with the PSTN, been left virtually unregulated. Indeed it is often argued that this lack of regulation is one of the reasons that the Internet has been able to develop so quickly and freely. In both economic and legal terms, therefore, it is natural to address the question of the need for regulation from the standpoint of competition law. This involves the traditional stages of defining markets, examining the cost and demand conditions in those markets and analysing the operation of the competitive process. In the case of the Internet backbone market, the coexistence of alternative interconnection arrangements - peering among the largest providers, transit arrangements between these and smaller ISPs - forms an important element in the competitive process.

90. It has been argued that universal peering, if left unmonitored would lead to free riding and opportunistic behaviour. If, on the other hand, peering were restricted to the few largest players, it might have desirable features in terms of reducing transaction costs.

91. If this form of discrimination is not inherently against the public interest, the next question is whether, in practice, Internet backbone providers are able to achieve or maintain and jointly dominant position in the market, and if so, whether this ability relies upon peering arrangements.

92. The evidence given above does not permit a conclusive answer to the first leg of this question. However, there is little persuasive evidence of parallel pricing, entry deterrence or excess profits. If, however, the answer to this question were yes, then it is possible to imagine ways in which confining peering to members of the jointly dominant group would be a method of achieving price discrimination and deterring entry.

93. There is thus a case for keeping the Internet backbone market under review, and in particular for taking a sceptical view of mergers which lead to step changes in concentration, especially where they involve the largest firms.

REFERENCES

Baake, P & T Richmann (1999), 'On the Economics of Internet Peering'. *Netnomics* Volume 1 No. 1 pp 89-105.

Bailey, Joseph 1997, 'The Economics of Internet Interconnection Agreements'. pp 155-168, Lee W McKnight and Joseph P Bailey (Eds), *Internet Economics*, 1997.

Bailey, Joseph (1999) *The Industrial Organisation of the US Internet Service Provider Industry.* Paper presented to 27th Annual Telecommunications Policy Research Conference.

Cave, Martin and Mark Donnelly (1996). 'The Pricing of International Telecommunications Services by Monopoly Operators'. *Information Economics and Policy*, Vol. 8 pp 107-123.

Crémer, Jacques, Patrick Rey and John Tirole (1999), *Connectivity in the Commercial Internet.* Institute d'Economie Industrielle, document de travail 87, July.

Cukier, Kenneth Neale Peering and Fearing: ISP Interconnection Regulatory Issues. http://www.ksg.havard.edu/iip/iicompol/Papers/cukier/html.

Downes, Thomas and Shane Greenstein (1999). 'Do Commercial ISPs provide Universal Access?'. In Sharon Gillett and Ingo Vogelsang (Eds). *Competition, Regulation and Convergence*, pp 125-212.

European Commission (1999) 99/287/EC: Commission decision of 8 July 1998 declaring a concentration to be compatible with the common market and functioning of the EEA agreement. *Official Journal* L 116,04/05 1999, pp 1-35.

Frieden, Rob (1999), 'Last days of the Free Ride?: The Consequences of Settlement- based Interconnection for the Internet'. *Info*, Volume 1, No. 3, pp 225-238.

Huston, Geoff (1999), *Interconnection, Peering and Settlements*. http://www..telstra.net/gih/peerdocs.

Jew, Bernadette and Rob Nichols (1999), *Internet Connectivity: Open Competition in the Face of Commercial Expansion*. Paper presented at the Pacific Telecommunications Conference, Hawaii, January.

Kende, Michael and Jason Oxman (1999), *The Information Interchange: Interconnection on the Internet*. Paper presented to 27th Annual Telecommunications Policy Research Conference.

Lasok, Paul (1999). Collective Dominance and the Telecommunications Sector, Mimeo

Milgrom, Paul, Bridger Mitchell & Padmanabhan Srinagar (1999), *Competitive Effects of Internet Peering Policies*. Paper presented to 27th Telecommunications Research Conference.

OFTEL (1999) Consultation Paper on the Relationship between Retail Prices and Interconnection Charges for Number Translation Services, Nash OECD (1998) Internet Traffic Exchange : Development and Policy DSTI/ICCP/TISP (98) 1/FINAL.

Rapp, Lucien (1999), 'Competing for the Internet: Reciprocal Access, Interconnection Agreements and Economic Control of Backbone Infrastructures'. *Communications & Strategies* No. 34, pp 71-107.

Srinagar, Padmanabhan (1997), 'Internet Cost Structures and Interconnection Agreements'. Lee W McKnight and Joseph P Bailey (Eds), *Internet Economics*, 1998.

NOTES

[1] In practice the distinction between the two models is blurred because In some countries the IAP is treated as a telecom operator and therefore has a right to obtain termination charges or there is direct negotiation between the incumbent and the IAP to determine how much of call revenue is given to the IAP.

[2] Any ISP, including one which charges a monthly fee, can (and does) make the same bargain with the terminating operator. The revenue is, however, more central to a "free ISP".

[3] In Freeserve's case, revenues are split almost equally between a share of termination charges and other sources.

[4] The incumbent may also abuse its market power on negotiating termination rates with other operators.

[5] Where Internet access is achieved by ADSL (which has a different cost structure), flat rate charging is the norm.