Competition in the Internet Backbone Market•

Paolo BUCCIROSSI*, Laura FERRARI BRAVO** and Paolo SICILIANI***

April 2005

Forthcoming in *World Competition* 28(2): 235-254

This paper presents a competitive assessment of the global market for the provision of universal Internet connectivity (backbone market). We discuss the approach followed by the EC Commission in two important merger cases. The main argument is for a restless evolution of the structure of the market, whereby an highly concentrated US-centric industry, with a strict vertical hierarchy between Internet Service Providers and a neat separation between first-level ISPs and the rest of the market, is going to be superseded by a more horizontally shaped configuration. It is argued that, as the landscape of the industry is subject to continuous change, the approach followed by the EC Commission in assessing the competitive forces that drive the industry is likely to be no longer appropriate. New behavioral strategies, such as differentiation through the introduction of new enhanced Internet services based on the concept of Quality of Service, and, related to that, new competitive threats seem to characterize the foreseeable future of the Internet. We then investigate the competitive concerns that might emerge in the new environment.

JEL Classification L40 (Antitrust General) L86 (Information and Internet Services),

Keywords: Internet, backbone, competition policy

• The authors would like to thank Fabio Manenti for helpful comments and acknowledge EU financing through the Cocombine project (www.cocombine.org) IST-2004-2012.

* Director of Lear – Laboratorio di economia antitrust e regolamentazione.

** Assistant Professor, University of Rome "La Sapienza" and Lear.

*** Junior Economist, Lear - Laboratorio di economia antitrust e regolamentazione.

I. INTRODUCTION

The Internet is a complex industry where local and global players co-operate and compete to offer end users a vast array of services. Some of these services (such as email and websites) are inherent to the Internet; others (such as music, banking, voice communication, video, etc.) are available through the Internet as well as through other channels. The most striking characteristic of Internet services is their global reach that allows users to communicate or conclude transactions with other users located everywhere in the world, reducing or eliminating the need of physical movements of people and goods. Competition in the Internet depends significantly on the availability of universal connectivity that inherently qualifies Internet services. Universal connectivity is provided by Internet Backbone Providers (IBPs) that form the Internet backbone market. The aim of this paper is to provide the essential elements of the economic analysis required for the application of competition law in the backbone market. The point of departure of our investigation is the competitive assessment of the backbone market made by the EC Commission in two merger cases notified between 1998 and 2000, namely: MCI/WorldCom and MCI WorldCom/Sprint.[1] The Commission in both cases found that the proposed merger could negatively affect competition in the backbone market. Therefore it conditioned the first merger to significant undertakings and blocked the second merger. In this paper we discuss the assessment carried out by the Commission and come to the conclusion that the recent developments of the backbone as well as of vertically related markets make that assessment no longer appropriate.

Before going into the details of how the Commission worked out the competitive concerns arising from the notified mergers, it is necessary to succinctly describe the functioning of the Internet. Section 2 presents a brief and simple description of how universal Internet connectivity has been ever since delivered and about the way this seamless interconnection might be put into jeopardy by the development of distinct Quality of Service (QoS) proprietary platforms. Section 3 concisely describes the two mergers assessed by the Commission. In Section 4 we take up the market definition problem and show why the backbone market has to be held separate from those for other Internet-related services as envisaged by the EC Commission. Section 5 goes through the EC Commission competitive assessment of the "backbone market", providing an articulated description of the different economic agents engaged and of the structural dimension affecting the competitive dynamics in the market. The reasons why the EC assessment may no longer fit the current market configuration are treated in Section 6. Section 7 describes the "next thing", i.e. the competitive concerns that could arise in the backbone market in the future. It provides the outlook for a "balkanized" Internet or, even more, a monopolized Internet due to, respectively, the development of several non-compatible proprietary QoS platforms and the overwhelming imposition of that platform possessed and operated by a dominant IBP. Section 8 concludes about possible future directions of competition policy in Internet-related markets.

II. HOW THE INTERNET WORKS

The Internet is a system of interconnected computer networks. Thus, this industry is characterized by strong network externalities, whereby end users, web-content providers and businesses seek ubiquitous connectivity and purchase it from ISPs, which in turn refer to IBPs for the provision of high-band long-haul transmission, routing and interconnection services.

The networks that comprise the Internet are autonomous and self deterministic, and communicate with each other without being controlled by a central authority. The role of each network cannot be predicted in advance, since the Internet is based on a connectionless transmission technology. No dedicated end-to end connectivity is required and no fixed route

[1] Dec. 99/287/EC *WorldCom/MCI,* OJ [1999] L 116 and Dec. 03/790/EC *MCI WorldCom/Sprint*, OJ [2003] L 300.

has to be set up between the sender and the receiver, because the Internet makes use of a packet-switching technology to transmit data across the network. A packet formatting and addressing mechanism as such is independent of any specific characteristic of the individual networks comprising the Internet.

The operation of the Internet is supported mainly by two basic transmission protocols:

i) Internet Protocol (IP) is responsible for routing individual packets from their origin to their destination. Each computer has at least one globally unique identification address (IP address). The IP address contains information on both the network, the computer it belongs to, as well as its location in that network. Each packet transmitted over the Internet contains a "header" where both the sender's IP address and the receiver's IP address are codified.

ii) Transmission Control Protocol (TCP) controls the assembly of data into packets before the transmission and the reassembly of transmitted packets at the destination. TCP is a connection-oriented transmission mode, whereby ISPs can guarantee that all the data will be delivered to the other end in the same order as sent and without duplications. TCP is actually built on IP, adding reliability and traffic control to the Internet.

The best route for transmitting a packet from the origin to its destination is determined at each router-computer that the packet passes on its trip. The router's decision about where to send the packet depends on its current understanding of the state of the networks it is connected to. This includes information on available routes, their conditions, distance and cost. The packets, having the same origin and destination, travel across any network path that the routers or the sending system consider most suitable for that packet at each point of time. If at some point in time some parts of the network do not function, the sending system or a router between the origin and destination detects the failure and forwards the packet via a different route. The conventional IP/TCP packet-handling rule a router implements is the first come-first served (or First In First Out – FIFO).

Given the need for interoperability in order to provide universal Internet connectivity to customers, IBPs have spontaneously achieved seamless interconnection through a system known as peering. Peering agreements present several distinct features:[2] (*i*) Peering partners reciprocally exchange traffic that originates with the customer of one network and terminates with the customer of the other peering network. Consequently, as part of the peering arrangement, a network would not act as an intermediary and accept the traffic from one peering partner and transit this traffic to another peering partner (peering is not a transitive relationship); (*ii*) In order to peer, the only costs are those borne by each peering network for its own equipment and for the transmission capacity required for the two peers to meet at each peering point; (*iii*) routing is governed by a conventional rule known as "hot-potato routing", whereby a backbone passes traffic to another backbone at the earliest point of exchange.

Lastly, it is worth noting that as peering incurs between pairs and does not imply any kind of payment, recipients of traffic promise to undertake "best effort"[3] when terminating traffic, rather than ensuring a level of performance in delivering packets received from peering partners.

---

[2] See J.P. Bailey, "*The Economics of Internet Interconnection Agreements*", in: McKnight and Bailey (eds), *Internet Economics*, (MIT Press, 1997).

[3] In a "best effort" setting, when congestion occurs, the clients (software) are expected to detect this event and slow down their sending rate, so that they achieve a collective transmission rate equal to the sending throughput capacity of the congested point. The rate adjustment is implemented by the TCP. The process runs as follows: a congestion episode causes a queue of packets to build up; when the queue overflows and one or more packets are dropped, this event is taken by the sending TCPs as an indication of congestion, so that the sender can slow down. Each TCP then gradually increases its sending rate until it again receives a congestion signal.

A. QUALITY OF SERVICE

Regarding the interconnection setting, the *status quo*, as managed through the "best effort" conventional rule, has proven to be unsatisfactory in dealing with ever increasing and bursty traffic flows ingenerating congestion at the routing nodes of the Internet. Congestion is particularly concerning given the increasing adoption of new web applications which are strongly demanding in terms of traffic throughput, requested capacity and network affordability, as they are live applications (real time provision and/or interaction). All this has put forward the importance for an IBP to guarantee high standard of quality in terms of connectivity provision (bandwidth capacity, redundancy, affordability, scalability, etc.). These features are usually summed up by a synthetic index of performance: Quality of Services (QoS).[4] QoS refers to the probability of the network meeting a given traffic contract, or, more informally, it refers to the probability of a packet passing between two points in the network. A traffic contract, usually labeled as a Service Level Agreement (SLA), specifies guarantees for the ability of a network/protocol to give guaranteed performance/throughput/latency bounds based on mutually agreed measures, usually by prioritizing traffic. A SLA is an agreement aimed to avoid several transmission hiccups.

A given QoS may be necessary for certain types of network traffic, for example: Streaming multimedia that require guaranteed throughput; IP telephony or video conferencing that require strict limits on jitter and delay; a safety-critical application, such as remote surgery.[5]

There are essentially two ways to provide a QoS guarantee. The first is simply to deploy enough transmission capacity to meet the expected peak demand with a substantial safety margin. However, if the peak demand increases faster than forecasted, this solution could not suffice. Moreover, it is expensive and time-consuming in practice.

The second one is to require people to make reservations and only accept the reservations if the routers are able to serve them reliably. This solution amounts to a sort of priority scheduling, whereby bandwidth capacity allocation among customers is accomplished by creating transmission service classes of different priorities to serve customers with different needs.[6] The way a customer applies a reservation is by negotiating with the ISP a SLA. The contract of SLA will specify what classes of traffic will be provided, what guarantees are needed for each class and how much data will be sent for each class.[7] The definition of priority requires the sender to set the "type of service", and to fill in the IP header according to the class of data, so that better classes get higher priority. There are many ways to split traffic into classes. Special handling may be done in at least two different ways:

i) Preferential forwarding, where more recent higher precedence packets are allowed to jump the queue over old lower preference packets;

ii) Preferential discarding, where buffer space for higher preference packets is allowed to grow at the expense of lower precedence packets which are discarded.

The technical implementation of a SLA contract makes use of a set of alternative transmission protocol modes usually clustered under the definition of "fast-packet services" or

[4] See e.g. H. Junseok and M.B.H. Weiss, *The Economics of QoS Allocation Strategies; An Empirical Study*, 2001, available at http://web.syr.edu/~jshwang/resource/isqe-hwang-paper-mit.pdf. For a technical perspective visit the Internet Engineering Task Force – IEFT web site: http://www.ietf.org/home.html.

[5] These types of applications are called "inelastic", meaning that they require a certain level of bandwidth to function (no less/no more).

[6] For an exhaustive analysis of usage-sensitive pricing schemes, see L.W. McKnight and J.P. Bailey, *Internet Economics*, (MIT Press, 1997.

[7] Traffic requirements are made up of four categories: (1) bandwidth, (2) delay, (3) delay jitter, and (4) traffic loss.

“cloud technologies”.[8] These protocols form one of the “virtual networks”[9] built on top of facilities and layered services provided by telecommunication carriers. The need of an additional underlying transmission protocol is due to the fact that the IP/TCP protocol, as it was originally conceived with its FIFO routing rule, is unable to count for differentiated class of services and, thus, cannot manage a prioritizing allocation of bandwidth resources. Yet, this lack turns out to be the main reason for the success of IP/TCP, since it has effectively provided a minimum common denominator for granting universal interoperability between private operators competing among each others for transit revenues and managing networks with different architectures, routers and switching facilities. Thus, the prioritizing function is handled at a virtual layer just below that of the IP/TCP.

So far, the significance of this issue in the assessment of relevant markets for the provision of universal internet connectivity has been of less relevance than the imperative, for IBPs to be facility based so as to be able to preside different regions and maintain a widespread customer base, since private peering interconnection between IBPs has coped fine with the connectivity requirements implied by the current applications massively carried over the Internet.

III. MERGERS IN THE BACKBONE MARKET

Between 1998 and 2000 two mergers concerning the Internet backbone market were notified to the EC Commission under the EC Merger Regulation. The first operation involved MCI and WorldCom, two large US operators providing the full range of telecommunication services. The second merger was between MCIWorldCom, the entity resulted from the 1998 merger, and the US telecom operator Sprint. The first operation was given conditional clearance with the imposition of structural remedies; for the second one authorization was denied on the grounds that it was incompatible with the common market.[10]

The 1998 merger case was the first occasion for the Commission to examine Internet-related markets from a competition law perspective. In order to assess competition, the Commission identified the products/services on which the parties competed (relevant markets), the presence of other operators and the geographical scope of the markets. The Commission found that the backbone is a relevant market on its own in that second-level ISP who do not manage backbones cannot achieve universal connectivity without purchasing transit from IBPs. The structure of supply on this market had to be derived by the Commission as there were no available official statistics on market shares. The Commission drew a list of 16 actual competitors and calculated that in terms of traffic flow and revenues the merging parties would have had the largest share of the market, with the two nearest competitors enjoying only half the size of the combined entity. The great absolute and relative size achieved by the combined entity after the merger suggested that there was the risk that it could behave to an appreciable extent independently from its competitors and customers. This conclusion was strengthened by the circumstance that the combined entity could also act strategically to maintain or reinforce its dominant position by denying peering requests by potential competitors and by raising costs to actual rivals and/or degrading their connection. The Commission examined also the remedies proposed by the parties to relieve the competitive concerns. The main remedy consisted in the divestiture of MCI Internet business. The Commission felt that the divestiture to an acquirer capable of replacing the departing player would have restored competition in the market.

---

[8] The term “cloud” refers to the geographic area covered by the collection of routes and links between them that delimitates the network area over which the protocol uniformly applies.

[9] J. Gong and P. Srinagesh, “*The Economics of Layered Networks*”, in: McKnight and Bailey (eds) *Internet Economics* (MIT Press, 1997).

[10] In September 2004, the Court of First Instance overturned the Commission decision declaring that the Commission did not have the authority to prevent the merger because the parties had withdrawn their notification. The court ruling was based only on procedural considerations and did not deny validity of the Commission assessment.

Therefore, the merger was authorized conditionally to the divestiture of MCI's Internet business.

In the 2000 merger case the Commission maintained that the backbone market was to be held separate from other Internet-related markets and that the merger would have affected competition in this market as the merging parties were the two largest providers. As in the previous case, an in-depth investigation showed that the merger would have led, through the combination of the merging parties' extensive networks and large customer base, to the creation of such a powerful force that both competitors and customers would have been dependent on the new company to obtain universal Internet connectivity (unilateral effect). The Commission estimated that the combined entity would have gained a market share of between [37-51]% with more than [40 to 80]% of its traffic staying on-net, whereas other networks would have had substantially lower market shares and no more than 32% of traffic on-net. This lead the Commission to the same conclusion as the 1998 case as to the ability of the post-merger dominant operator to obstruct competition by implementing a selective degradation strategy and foreclosing entry to potential competitors. As in the 1998 case, the parties proposed as a remedy to divest Sprint's Internet business. However, this time the remedy was deemed insufficient to relieve the competitive threats from the merger as the efficiency and competitiveness of the divested entity would have been seriously jeopardized by the separation from the underlying Sprint telecommunication infrastructure. Authorization was thus denied.

In both cases the Commission had to conduct a thorough assessment of competition in the markets affected by the proposed mergers before clearing the parties' request for authorization. The line of reasoning set forth by the Commission in these two cases is more deeply described in the next two sections before illustrating from Section 6 onward why the framework depicted by the Commission is no longer appropriate.

IV. MARKET DEFINITION

The market definition for the provision of top level or universal Internet connectivity draws upon the hierarchical structure and topology of the Internet. The hierarchy is a matter of technical constraints, but its main reflections are commercial. IBPs are network operators who own, run and upgrade long-distance transmission networks that together form the global Internet international 'backbone', which connects to multiple countries in more than one region of the world. These infrastructures constitute an essential facility for every ISP wishing to provide its customers with global connectivity. The US-centric topology of the Internet is a historical legacy of its inception.[11]

The hierarchical topology of the Internet is the outcome of the evolution of ISPs' interconnection policies. Since its initial development, the commercial Internet was managed through voluntary peering agreements. Initially, most exchange of traffic under peering arrangements took place at Network Access Points (NAPs), as it was efficient for each backbone to interconnect with as many backbones as possible at the same location.

The ever increasing diffusion of the Internet and the penetration of new traffic-demanding web applications (such as voice over IP, video conferencing, video and music streaming) caused traffic flows to growth steadily. The rapid growth in Internet traffic soon caused the NAPs to become congested, which led to delayed and dropped[12] packets. As a result of the increased congestion at NAPs, many backbones began to interconnect directly with one another. This

[11] For an historical overview of the Internet from its academic take off to the current commercial era, see M. Kende, "*The Digital Handshake: Connecting Internet Backbones*", Working Paper No. 32, FCC-OPP, 2000.

[12] At the router, when the incoming rate exceeds the outgoing rate, packets can be temporarily queued and delayed, and eventually discarded ("dropped").

system has come to be known as private peering, as opposed to the public peering that takes place at the NAPs.[13]

The EC Commission has grasped this evolution, as it is arguable by comparing how the Commission has handled the assessment of the relevant markets for universal Internet connectivity in the two cases in 1998 and 2000 respectively.

In the 1998 WorldCom/MCI case, in order to identify IBPs (also labeled as "top-level, or top-tier, or tier-1 ISPs") and to distinguish them from ordinary ISPs (also labeled second-level, or second-tier ISPs), the EC Commission defined top level operators as those able to grant themselves global reach by means of peering agreements, both public and private, without purchasing "transit"[14] from anyone.[15] In theory, second-level ISPs could reach global connectivity also through peering agreements between each other. The benefits from entering into a peering agreement relative to a transit one, would be lower transit costs and lower latency in traffic delivering. However, since this would imply managing a great number of contractual relationships, it is actually implausible that second-level ISP would make such a choice. Rather, second-level ISPs will prefer to gain global connectivity by entering into a single transit agreement with any of the existing IBPs. Therefore, the products offered by IBPs are differentiated in that the connectivity is supplied entirely by peering agreements between those top-level networks or internally.

In the WorldCom/MCI market definition the Commission deemed both private and public peering as valid proxies to identify IBPs. Therefore, the possession of, for example, peering agreements at public NAPs with all other ISPs might well have guaranteed the ISP concerned the status of a top level network. However, the Commission argued that this was a week definition, necessarily bounded to lose any signaling power in the short term, due to the increasing congestion at public peering points.

In the 2000 WorldCom MCI/Sprint case a more rigorous and strict defining criterion was adopted by the EC Commission for defining relevant markets. As congestion at NAPs had increased and large providers had increasingly began to form their own private peering arrangements at points away from the NAPs, the smaller IBPs, who previously had peered only at the NAPs, were refused settlement-free private peering by the largest networks. As a consequence, these were no longer capable of acting as top-level networks, and, therefore, were dropped out of market definition.

In order to be an operator to peer with at a private level, one needs to demonstrate comparable traffic throughput, flows and geographic scope, since actual top-tier ISPs cautiously settle private peering agreements only with pair-status operators.[16] Actual tier-1 ISPs even require operational dimensions on would-be peering partners so as to prevent the latter from free riding.[17] Traffic asymmetry could cause free riding if, for example, one of the two peering partners has focused its offering toward content providers and the other partner, given the hot-

---

[13] For a game theoretic analysis of the choice for an IBP between private and public peering, see N. Badasyan and S. Chakrabarti, "*Private Peering among Internet Backbone Providers*", Economic Working Paper No. 0301002, WUSTL,2003.

[14] Transit is a commercial service granting access to the Internet for a fee and it is a vertical contract of service provision from an upstream ISP to a downstream customer (ISP, business or consumer). Transit can take three forms: dedicated access (a dedicated line to another network provider or large customer), retail dial-up access (to residential and business customers) or wholesale dial-up access to Internet service providers.

[15] For a comparison between peering and transit, see W.B. Norton, "*Internet Service Providers and Peering*", Draft Version No. 2.5, 2001, available at http://www.ecse.rpi.edu/Homepages/shivkuma/teaching/sp2001/readings/norton-peering.pdf.

[16] Recently, many tier-1 ISPs have published their private peering policies and prerequisites on line. The prerequisites for peering with Tier 1 ISPs vary but generally include a peering presence in four or more regions where both parties have a presence along with sufficient transport bandwidth and traffic volume to warrant direct interconnections. *See* e.g. www.uu.net/peering/; www.level3.com/us/info/network/interconnection; www.genuity.com/infrastructure/interconnection.htm.

[17] For an analysis of the peering decision making process of IBPs, see note 13, above.

potato routing rule, has to carry over its network cumbersome returning flows of web-content in response of content-requests originated with its consumer base. Under these circumstances, the tier-1 ISP will normally refuse to private peer with the ISP specialized in web hosting, or with ISPs with a customer base composed in large proportion of web sites and content providers.

Given the reasons above, the Commission concluded that an IBP should be able to gain universal connectivity solely on the basis of its own network (and its customer's networks) and the networks of its private peering partners (and their respective customer's networks) in order to qualify for a top-tier status.

V. COMPETITIVE ISSUES

The present section outlines the competitive assessment of the "backbone" market, as laid down by the EC Commission. The detailed description of the market is essential to grasp the competitive concerns expressed by the Commission in the two mentioned cases.

A. THE COMPETITIVE FIELD

The Internet is not subject to any sector-specific regulation governing interconnection between IBPs. One consequence of the lack of any regulatory framework is the absence of specific reporting obligations on ISPs about Internet revenues, and, therefore, of a consistent reporting standard. The Commission could not find a reliable publicly available estimate of the size of either the Internet sector as a whole or of any relevant sub-sector, not to mention the possibility to obtain accurate figures about market shares. Also, there was no sort of industry consensus about a preferred and proper unit of measurement of market shares. Therefore, the Commission had to collect data on a variety of key operational dimensions before choosing the one(s) that best described the market. These are:

i) Traffic flows;
ii) Revenues from basic Internet access;
iii) Aggregate capacity in interconnection links;
iv) Number of addresses reachable;
v) Numbers of points of presence;
vi) Actual bandwidth used for traffic exchange.

The lack of a reporting standard and of consistent figures implies the use of conjectures and estimates in order to identify actual competitors and assess their markets shares.[18] As stated before, the minimum identification criteria used by the Commission looked for ISPs who peered at least with all the main renown[19] top-tier ISPs, since the failure, on the part of an ISP, to peer with just one of these dominant tier-1 ISPs would have implied a substantial absence in coverage of the Internet as a whole.

A common feature to both merger cases examined by the Commission is the strong concentration observed in the market for the provision of universal connectivity.

In *WorldCom/MCI*, beyond the two merging parties, the Commission registered the presence of two other large IBPs - Sprint and GTE (ex-Genuity) – and twelve small competitors. Even though the Commission did not provide detailed data, it stated that the combined entity would have held over 50% of the market, however widely defined, and would have been significantly larger than the size of its nearest competitor (Sprint), on either revenue

[18] According to the Commission, there is not a clear-cut dividing threshold between the smallest IBP and the biggest second-level ISP. Therefore there is the need for a conventional proxy in order to identify top-tier ISPs.

[19] In both merger cases, the Commission inquiries resulted in the identification of a cluster of 4/5 big backbone operators and a fringe of minor competitors.

or traffic flow, bearing in mind that the next other competitor, GTE, was about half the size of Sprint.

In *WorldCom MCI/Sprint* the main operators were: WorldCom MCI, Sprint, AT&T, Cable & Wireless and GTE and a fringe of twelve minor operators. The large five IBPs' combined market share amounted to roughly [42-86] %.

Another proxy measure frequently quoted by the Commission to map the competitive field and the IBPs relative position of strength is the percentage of traffic "staying on-net"[20] out of a network's total traffic flow. The Commission argued that these structural figures capture the degree of "strategic" independence of an IBP from other actual and potential competitors.

In respect to potential competitors, the hierarchical structure of the Internet and its peculiar peering interconnection architecture suggest that likely new entrants are actual transit customers striving to grow in terms of both geographic coverage and customer base served in order to become eligible as private peering partners. Thus, the trajectories pursued by potential competitors are necessarily bottom-up and internal (captive) to the ISPs' hierarchy. This condition creates competitive concerns in so far as potential competitors are, at the same time, the main source of revenues for IBPs who would be more reluctant to accommodate a former customer than an outsider new entrant.[21]

This identifies "strategic entry barriers"; incumbent IBPs would likely foreclose new entrants not only because they represent former customers paying for transit but also because, given the imperative for an IBP to preserve its status of peering partner, top-tier ISPs have to maintain their consumer base and, even more, to attract new transit customers, rather than accommodate their entry.

Independently of strategic interactions between IBPs, the backbone market is characterized by other entry barriers. In order to provide universal connectivity, an ISP has to build its own network infrastructure (backbone). Moreover, the would-be top tier ISP has to reach a traffic throughput and geographic coverage comparable to that of actual top tier operators in order to be an operator to peer with at a private level. As the Internet grows hastily and actual top tier ISPs strive to catch up with ever increasing bandwidth capacity requirements, it would be more and more difficult for would-be top level ISPs to gain a pair status that will enable them to be eligible as private-peering partners. The market is thus characterized by the presence of strategic and dynamic entry barriers.

Transit customers could be functionally subdivided into the following categories:[22]

i) Downstream ISPs serving individuals, businesses and even smaller providers. They pay upstream backbone ISPs for connectivity, the price of which depends on the location and amount of data. Potential competitors are more likely to be previous downstream ISPs;

ii) Online service providers, like AOL, who earn revenues by providing Internet access and focusing on content and ease of use. Online service providers lease connectivity from backbones or other upstream ISPs and manage the network points of presence (POPs) that connect dial-up customers to the Internet;

iii) Web hosting companies, like Exodus, who host websites that are accessed by the Internet public. It is important to note that web hosting ISPs create unidirectional traffic, as websites originate a lot of traffic, while not requesting much. As a result, backbone

[20] The term "on-net" refers to traffic which is end-to-end comprised within the boundaries of a network and, thus, is independently carried over by a sole IBP.

[21] In the language of the Peering Community: "Once a Customer, Never a Peer", W.B. Norton, "*The Art of Peering: The Peering Playbook*", Draft Version No. 1.2, 2002, available at http://www.xchangepoint.net/info/wp20020625.pdf.

[22] See note 11,above.

ISPs demand that web hosting providers, which typically do not maintain a national network, purchase connectivity form a backbone or downstream ISPs.

iv) Big businesses, whose ubiquitous global presence and traffic induce a direct relationship with IBPs rather than a mediated one through downstream ISPs. Indeed, the provision of universal connectivity is usually part of a composite service[23] delivered to large undertakings and, therefore, it may be misleading to assess customers' reaction of large businesses and other transit customers alike.

This competitive landscape suggested the EC Commission that the horizontal mergers scrutinized would have created a dominant IBP in the backbone market capable of increasing prices regardless of the constraints on its market power put forward by existing rivals, potential competitors and the countervailing power of buyers (unilateral effects).

B. UNILATERAL EFFECTS THROUGH SELECTIVE DEGRADATION

According to the Commission, each of the two notified concentrations raised serious competitive concerns about the creation of a dominant position in the market for the provision of universal Internet connectivity which would enable the new entity to operate independently from its actual and potential competitors as well as customers. Thus, the economic rationale underpinning the Commission arguments is that of unilateral effects.

In the Commission view, the key tactic a dominant player would implement is that of "selective degradation" (SD). This consists of decreasing bandwidth capacity at private peering points, or refusing to increase it when requested. Since peering points are private and in a peering agreement partners exchange traffic only for termination purposes, this tactic hits directly the selected network operator and affects other top tier ISPs only indirectly and marginally. The tactic of SD is rational when a dominant IBP can act independently of its competitors. Given the strong network externalities in interconnection policies, which provide IBPs incentive to cooperate with its competitors to provide seamless universal connectivity, to be rational the choice of SD has to generate a payoff that outweighs (dominates) that from a seamless interconnection strategy. As an indicator of independence and, thus, of dominance of the SD strategy, the Commission adopted the percentage of traffic "staying on-net" out of one network's total traffic flow, since a dominant IBP relies only marginally on each of its smaller competitors, whereas it will be a major source of connectivity to each of these.[24]

The EC Commission position is largely based on theoretical arguments developed in a seminal paper on backbone competition by Cremer et al.[25] In their model, backbones have some installed base of customers and compete for new ones. The model incorporates positive externality effects of increasing the number of customers. The more customers are attached to the backbone, the better is the quality of service. On the other hand, quality increases with better interconnection among backbones. Demand functions depend on prices and qualities of service. Given this setting, the authors show that in case of backbones of different size, the larger backbone prefers a lower quality of interconnection compared to the smaller backbone. Indeed, even though a higher quality of interconnection expands demand, it also reduces the quality

[23] In *WorldCom MCI/Sprint*, the Commission argued that a separate relevant market for the provision of telecommunication services to multinational corporations (MNC) does exist. Given the complex blend of needs and requirements expressed by MNCs, their procurement decisions refer to a bundle of telecommunication services (Global Telecommunication Services), whereby the provision of universal connectivity is just a basic component.

[24] In *WorldCom MCI/Sprint*, for example, if the merged entity were to degrade the connectivity of one of its four largest competitors, this would only affect about [0 to 10] % of its overall traffic (as implied by on-net traffic percentages), but, it would amount to more than [10 to 20] % of such traffic for any of the largest competitors exchanging with the merged entity (off-net traffic).

[25] J.P. Cremer, J.P. Rey and J. Tirole, "*Connectivity in the Commercial Internet*", Journal of Industrial Economics, vol. 48 No. 4 (2000) pp. 433-472.

differentiation between the two networks with asymmetric customer base. Therefore, the larger IBP prefers to sacrifice some demand expansion in order to preserve (or increase) its quality advantage.

The strategic goal pursued by the dominant IBP through the implementation of SD shall not be necessarily to foreclose competitors. A dominant IBP could also raise rivals' costs, or discipline the market signaling its intention to retaliate against those who deviate from his behavioral prescriptions. More in detail:

i) A dominant IBP could increase the relative price of its customers connections. In doing so, the dominant ISP would not be constrained by the presence of competitors, as these would be exposed to the threat of SD;

ii) A dominant IBP could discipline the market by the mere threat of selectively degrading the connectivity of its competitors. This will allow it to control both actual and potential competitors, as well as customers in the market,

In response to a degradation at a private peering point, an actual competitor could by-pass the damaged node by recurring to multihoming.[26] However, this strategy is virtually unfeasible as multihoming proves effective for out-going traffic but not for returning traffic flows (over which the degraded network could not exercise any significant control).

Furthermore, actual competitors would also have to face the reaction of their customers. Indeed, when comparing the quality of connectivity being offered by the dominant IBP to that being offered by its competitors, customers would find it more beneficial to switch a bulk of their traffic away from the degraded network to the dominant IBP.

As to potential competitors, in addition to the offensive moves against them by the dominant IBP, these would face the very same reactions by the other actual top-tier providers which would strive to maintain their pair status and, consequently, keep the pace of the dominant player (in terms of traffic throughput and geographic coverage).

Lastly, given the importance of being connected to the dominant IBP's network, also its own customers (second level ISPs, and MNCs) would not be able to retaliate to an increase in price or to a degraded connectivity. Unless all customers can act as a unit (and there is no evidence that the customer base is sufficiently concentrated to permit this) no individual customer will take the risk of moving to obtain a possibly inferior service without having any assurance that a sufficient number of other customers would take the same step.

The reasons above suggested the EC Commission that the US dominant player resulting from any of the proposed mergers would have been unfettered in increasing transit fares in the EC market and, thereby, access prices to European Internet final users. This conclusion stemmed also from the fact that at the time the assessment was conducted, European largest ISPs were only classifiable as second level ISPs and, thus, as customers of US IBPs. Their countervailing power was deemed to be negligible since they lacked the infrastructure necessary to gain peer status, thus constraining the dominant IBP.

Recent years have witnessed intense competitive dynamics which have profoundly changed the structure of the backbone market. In light of these changes, the EC Commission competitive assessment of competition in the backbone market may no longer be appropriate. Next section discusses the main directions of change and the new threats to competition that are posed by these recent developments.

[26] The practice of network providers and Internet access providers of being connected to more than one network is referred to as 'multihoming'.

VI. A LESS HIERARCHICAL INTERNET STRUCTURE

The pivotal element underpinning the EC Commission competitive assessment is that the international "backbone" infrastructure, which is the "essential facility" for the provision of universal Internet connectivity, was almost all possessed and operated by the five larger US IBPs, and that within this group, as a result of the two notified concentrations, a dominant IBP capable of exercising substantive market power independently of its rivals would have emerged.

This argument is progressively losing its significance as countervailing centrifugal forces are driving a de-agglomeration process of the Internet traffic, whereby Internet traffic flows are no longer tightly knitted to specific physical places. The result of this evolution is that the original U.S. centric Internet architecture is progressively eroding.

As Giovannetti and Ricuccia[27] argued, these centrifugal forces are: i) the European investment wave in backbone infrastructure; ii) the development of trading marketplaces for IP/transit exchanges; iii) the proliferation of European IXPs; iv) a process of cultural and linguistic differentiation of web contents; v) the application of new technologies and practices: catching, multi-homing and mirroring.[28]

As regards the first point, the process of liberalization in the TLC sector both in the U.S. (with the 1996 Telecommunication Act) and all across Europe, spurred carriers of both side of the Atlantic to deploy "end-to-end" infrastructure on global and national routes. Thus, European large companies have deployed their own "backbone" running throughout the U.S.[29] and, therefore, have avoided paying transit fares for gaining connectivity into the U.S. This achievement was made feasible by a range of various options: building network infrastructure; partnering with a U.S. carrier; purchasing dark fiber capacity and services; swapping capacity; purchasing a U.S. company; sharing a company.[30].

The second point made by Giovannetti and Ricuccia[31] (development of trading marketplaces for IP/transit) refers to trading platforms that gather IBPs in a single marketplace wherein bandwidth trading occurs through a centralized process with transparent prices while providing QoS information. For example, the Band-X trading place provides daily prices for monthly Internet transit at different bandwidths, from its trading floors in London and New York.[32] This increases transparency in the market for the provision of universal Internet connectivity (in the usual bilateral transaction, prices and terms of trade are kept confidential), allowing transit customers to perform a more conscious supply decision and, therefore, decreasing the switching costs they have to incur. Ultimately, this market-mediated transaction mode increases the effectiveness of competitive constraints from both the customer base and the existing competitors of U.S. dominant IBPs.

The third argument mentioned by Giovannetti and Ricuccia[33] (proliferation of European IXPs) refers to the development in Europe of several Internet Exchange Points – IXPs. These are physical network infrastructure operated by single entities with the purpose of facilitating the exchange of Internet traffic between ISPs through second-level peering agreements, thereby reducing dependency on their respective upstream transit providers. Any ISP that is connected to that IXP can exchange traffic with any of the other ISPs connected to the IXP, using a single physical connection to the IXP, thus overcoming the scalability problem of individual

---

[27] E. Giovannetti and C.A. Ricuccia, "*Estimatine Market Power in the Internet Backbone Using Band-X Data*", Working Paper No. 332, University of Cambridge, 2003.

[28] Catching is the storage of already accessed data; multi-homing is an alternative IBP which to recur in case of transmission hiccups; mirroring is the geographical or backbone multiplication of a web site content.

[29] E.g.: Telia, France Telecom, Deutsche Telekom, Telecom Italia, KPN.

[30] OECD, "*Internet Traffic Exchange and the Development of End-to-End International Telecommunication Competition*", DSTI/ICCO/TISP(2001)5/FINAL, Working Party on Telecommunication and Information Service Policies, 2002.

[31] Ibid. note 26, above.

[32] See the Band-X's web site: www.band-x.com.

[33] Ibid. note 26, above.

interconnections. Also, by enabling traffic to take a more direct route between many ISP networks (there are fewer hops between networks), an IXP can improve the efficiency and the fault-tolerance of the ISPs' provisions.[34] Most European IXPs are non-commercial co-operatives funded by membership fees paid by the connected ISPs. The actual saving will depend on the cost of membership (the range of membership fees is quite wide) and the amount of traffic that can be exchanged in relation to the ISPs total traffic.[35] The growing number of IXPs allows traffic to be exchanged on a regional basis rather than traversing transcontinental backbone networks.[36]

The last two points (cultural and linguistic differentiation and new technologies and practices: such as catching, multi-homing and mirroring) were keenly grasped by the EC Commission, but it was argued that at that time their effectiveness in easing the dependency of large European ISPs from U.S. IBPs was overweighed by the persisting U.S. centric feature of the Internet.

The likely consequences of the first three centrifugal forces for the competitive assessment of the "backbone market" are far-reaching as regard the market definition both in its product and geographical dimension. The product boundaries of the "backbone market" may be enlarging: this happens once second-level ISPs achieve IBP-status (the vertex of the hierarchy is flatting), trading floors ingenerate a process of commoditization and the diffusion of second-level public peering platforms, IXPs, eases the dependency of ISPs for transit provision. The geographical extension of the "backbone market" may no longer be global but increasingly regional: the combined effect of these centrifugal forces are contributing to the development of a denser matrix of interconnected networks, which, in turn, allows and implies a greater proximity in the provision of universal connectivity.[37]

In so far as this trend will consolidate, the argument for the creation of a dominant U.S. IBP able to act independently worldwide, regardless of the reactions of buyers and competitors, may lose its significance. This is not to say that the "backbone market" shall no longer be monitored, since the foreseeable development trajectories cast new concerns over the next-to-come Internet.

VII. THE NEXT COMPETITIVE CONCERNS

The economic framework underpinning the model of selective degradation outlined at Section 5.2. prescribes that a dominant IBP aims at vertically differentiating its offering as a higher quality IP connectivity relative to the one of the targeted rival. Given the IP/TCP perfect interoperability, this is the only kind of product differentiation available for the provision of basic IP connectivity.

Given the ongoing commoditization of the provision of universal IP connectivity, as outlined in the previous Section, IBPs might look for other strategies in order to differentiate their offering and, therefore, to price at a premium. This sort of differentiation strategies might pursue an horizontal differentiation through the provision of enhanced Internet services delivered on a proprietary QoS protocol platform. In accordance with the principles of strategic interaction, actual IBPs would strive to follow suit. This could imply a market outcome whereby

---

[34] J.S. Marcus, "*Global traffic Exchange among Internet Service Providers(ISPs)*", Paper presented at the OECD conference on Internet Traffic Exchange, Berlin, June 7, 2001.

[35] See the European Internet Exchange Association web site, www.euro-ix.net.

[36] See note 33, above.

[37] See E.J. Malecky, "*The Economic Geography of the Internet's* Infrastructure", Economic Geography, vol. 78 No. 4(2002) pp.399-424, E.J. Malecky, "*Fiber Tracks: Explaining Investments in Fiber Optic Backbones*", Entrepreneurship and Regional Development, vol. 16 No. 1(2004) pp. 21-39 and J. Rutherford, A. Gillespie and R. Richardson, "*The Territoriality of Pan-Europeann Telecommunications Backbone Networks*", GaWC Research Bulletin No. 136A(2004), Loughborough University. For an assessment of the geographical boundaries for the provision of universal connectivity in the UK, see OFTEL, "*Effective Competition Review of Internet Connectivity*", vol. 0.6 No. 23(2001).

future voluntarily interoperability between private networks is in jeopardy, as IBPs could leverage strategically the fast-packet service they run onto their network facilities as a means to differentiate their market position.

Indeed, fast-packet services are enhanced Internet services transit-customers must pay for and, thus, these cloud technologies are a foreseeable main source of revenue for tier-1 ISPs, once interactive applications spread across final users. The fact that, so far, no standard protocol spontaneously emerged shows that the need for interoperability on different cloud technologies among IBP peering partners is not an argument in private peering negotiations, and, thus, the "best effort" *status quo* is far from being substituted by some form of deeper cooperation among private peering partners.

A. THE "BALKANIZATION" OF THE INTERNET

The concept of relative independence of a dominant IBP and the incentive to differentiate its offering against the countervailing incentive to cooperate in the provision of seamless universal connectivity (pursuing positive network externalities), invokes a widely discussed issue about a concerning likely development of the next-to-come Internet: "the balkanization of the Internet".[38]

In order to provide QoS for enhanced Internet services, such as voice over IP, video conferencing and Internet banking, the reliability of Internet connections is very important. Therefore, IBPs would have to agree on a standard transmission protocol (fast-packet service), which runs just below the IP/TCP layer, and guarantees interoperability across competing private networks. Currently, no agreed protocol exists for such cloud technologies. This could entail an Internet truly universal only for basic services, such as WWW and e-mail, but subdivided among networks for the provision of enhanced Internet services (balkanization).

From a private perspective, the decision to interconnect for the provision of QoS services appear to be relatively similar to the one IBPs have so far made when deciding whether to peer with one another by balancing out the costs and benefits of interconnecting. The benefits stem from the positive network externalities from interconnection that attract new consumers and encourage consumption by the existing ones; the costs come from competitive network externalities: a backbone decision to interconnect with another backbone makes the other backbone more attractive to customers.[39]

There is, however, a difference between the current interconnection arrangement and new ones for the exchange of QoS traffic. Universal connectivity achieved through peering agreements is a legacy of the cooperative spirit characterizing the Internet in its early days that made basic services, such as e-mail and Web access, universally available. In this context, no IBP could differentiate itself based on the unique provision of these services. On the contrary, in the commercial spirit that pervades the Internet today, IBPs view the new services that rely on QoS as a means of differencing themselves from their competitors, and, if the strategy proves to be successful, of charging a premium to its own transit customers.

B. A MORE EXTREME SCENARIO

The "balkanization of the Internet" describes a static configuration of the backbone market, whereby differentiated QoS protocol platforms compete simultaneously. The same analysis conducted in a dynamic fashion could lead to a more extreme outcome, whereby a dominant IBP would have an incentive to develop a proprietary standard for such fast-packet services that

---

[38] See note 10, above, M. Kende and D.C. Sicker, "*Real-time Services and the Fragmentation of the Internet"*, paper presented at the 28th Telecommunication Policy research Conference, September 23-25, 2000, Alexandria (Virginia) and R. Frieden, "*Does Hierarchical Internet Necessitate Multilateral Intervention?*", paper presented at the 28th Telecommunication Policy research Conference, September 23-25, 2000, Alexandria (Virginia).
[39] See note 10, above.

would be offered only on its network, thus internalizing network economies by refusing to concede interoperability to other networks. Given the large consumer base relative to competitors, the QoS protocol platform of the dominant IBP might become the *de facto* industry-wide standard, thereby allowing the dominant IBP to tip the market and exclude other competing technologies, or, at least, to impose and control the industry technological development.

The economic foundation of this scenario is grounded on the analysis of allocation under increasing returns pioneered by Arthur.[40] Therein, a model is presented where users chose between technologies competing for a market of potential adopters and where each technology improves as it gains in adoption. The model prescribes that if one alternative technology gains an early lead in adoption it may eventually "corner the market" of potential adopters, with the other competing technologies becoming locked out. This framework fits well with high technology markets that call for compatibility with industry standard, and where increasing returns are due to network effects that benefit users each time a new user joins the market. Thus, the economics of networks applied to Internet-based new markets suggests a tendency toward winner-take-all outcomes. In the case of the emerging market for the provision of QoS services a dominant IBP could leverage its existing large customer base in order to gain a substantive first mover advantage and thereby monopolize the market.

C. IMPLICATIONS FOR COMPETITION POLICY

The two possible scenarios outlined above point out the risk that the Internet will develop in a suboptimal manner. Lack of interoperability or (maybe worse) monopoly power may reduce consumer and social welfare. On one hand this risk seems to call for some public intervention. On the other hand the spectacular expansion of the Internet in very few years, absent any form of regulation, indicates that market forces may prove much more reliable in fostering the adoption of new technologies and the surge of economic benefits. Moreover, uncertainty about the pros and cons of competing technologies and the fast pace of their development should induce public administrative authorities to refrain from meddling with technical and complex issues. Even the most benevolent regulator runs the risk of getting it completely wrong.

In this complex state of affairs competition policy may play an important role. Even if it is not immune from mistakes, its *ex post* character makes it more apt to correct market avenues leaning toward anticompetitive settings without interfering with initial technological choices. This is even more true if one believes (as we do) that the analysis of the backbone market contained in this paper proves that consolidation in this market should not raise competitive concerns and therefore should not warrant the strict application of the merger regulation that the EC Commission reserved to previous cases. Competition law provisions regarding agreements and dominance will probably constitute the main tools to preserve competition. With respect to horizontal agreements, the EC Commission and National Competition Authorities should have a permissive attitude towards those agreements concerning technical standards and interoperability. Their pro-competitive effects are likely to be sufficiently high to outweigh the risk of a collusive outcome. Moreover they seem to be the best way out the unpleasant choice between balkanization and monopoly.

The main reason for believing that firms will not opt for this cooperative solution, given the presence of strong network externalities, is that each of them (or most of them) believes to be able to win the war for standard and gain monopoly rent afterwards. In the recent *Microsoft* case, the EC Commission has made clear that it is not going to tolerate such strategy as it is willing to demand interoperability to the dominant operator that should emerge from a war for standard. If firms anticipate such attitude by an antitrust authority, the incentives to keep developing their cloud technologies in isolation should disappear. If network effects are significant, then a cooperative solution is likely to prevail.

[40] W.B. Arthur, "*Competing Technologies, Increasing Returns, and Lock-in by Historical Events*", Economic Journal, vol. 99(1989) pp.116-131.

Vertical relationships deserve further considerations. Both vertical mergers and vertical contractual arrangements could alter the prospects faced by Internet operators and modify their attitude toward cooperative solutions. Indeed, a way to win the war for standard is to secure complementary products such as premium content or exclusive access to advanced services (e-health, e-government, etc.). These links have two opposite effects. On one side they can help mitigating double marginalization problems, if the complementary product market is not competitive. On the other side they may stabilize an inefficient equilibrium with multiple standards and a balkanized Internet as the benefit stemming from the complementarities between the Internet and the "proprietary" advanced services could offset the loss due the unexploited network externalities. One way to solve this problem might be to treat these complementary products as part of the technical standards and to require their sharing if a concrete risk of balkanization materializes.

VIII. CONCLUSIONS

As regards the competitive assessment of the backbone market, the main concern raised by the EC Commission was that of unilateral effects by a dominant US IBP. Undoubtedly an industry-wide regulatory intervention setting mandatory rules for interconnection agreements between IBPs would not be the proper policy, as the current evolution of the upstream Internet shows that the market auto-regulates as there are sufficient market forces countervailing the strategic moves of dominant players.

Regarding competition policy interventions during the late '90s, the same cannot be said, as there is no counterfactual evidence that the market outcome, absent the EC Commission interventions in the two merger cases, would have been virtuous as well. It is arguable that, at the time the EC Commission decided, there was a wide consensus about the U.S. centric feature of the upstream Internet structure, which is the pivotal prerequisite of the EC Commission assessment.

The issue about QoS is not self contained in the "backbone market". The provision of enhanced Internet services is carried out by the contribution of operators acting at several layers of the Internet value chain: i) telecommunication carriers; ii) content providers; iii) downstream Internet access providers; iv) broadcasting service providers (on various technological platforms: cable, satellite, mobile and fixed telephony, terrestrial digital TV); v) client-software providers. Thus, it is far from clear which Internet operator will preside the key strategic layer and, thereby, will affect the delivery process throughout the Internet value chain. Even more difficult is to forecast that IBPs will be the ones to succeed in this task on their own forces, leveraging their proprietary fast-packet platforms. What seems more arguable is that the QoS issue will spur competing processes of vertical integration between complementary economic agents and, therefore, the market outcome will depend on the competitive structures and dynamics at different layers of the value chain. Thus, it appears that competition law enforcers will have to carry out a careful assessment of the dynamics throughout the Internet value chain both horizontally and vertically, instead of a focused assessment on a layer which appears to be the strategic bottleneck of the delivering process.