The Network Operator's Perspective on Peer-to-Peer: Business Threats or Opportunities?

Stefano Bregni*, *Senior Member, IEEE*, Giacomo Bruzzi◊, Maurizio Dècina*◊, *Fellow, IEEE*

* Politecnico di Milano, Dept. of Electronics and Information, Piazza Leonardo Da Vinci 32, 20133 Milano, ITALY
Tel.: +39-02-2399.3503 – Fax: +39-02-2399.3413 – E-mail: {bregni,decina}@elet.polimi.it
◊ ICT Consulting S.p.A., Via Vittor Pisani 22, 20124 Milano, ITALY
Tel.: +39-02-676422.36 – Fax: +39-02-676422.43 – E-mail: bruzzi@ictc.it

***Abstract* — This paper studies the impact of peer-to-peer (p2p) applications on network operators. P2p applications affect the way network resources are used, changing traffic behaviour at access, edge and peering levels. In particular, they make up overlay networks with topology decoupled from the underlying physical networks: a consequence is inefficiency and higher link loads at edge and peering levels. The growth of overlay applications is affecting the capability of operators to sustain their business: traffic grows without proportional additional revenues. Operators should not be allowed to block selectively p2p traffic or users, to not violate the "network neutrality" principle. However, operators may still use p2p traffic information without infringing privacy and neutrality principles. We highlight a possible approach for network operators to face these issues and to turn uncertainties caused by p2p applications into future opportunities.**

***Index Terms* — business, Internet, p2p, telephony, traffic.**

I. Introduction: Network Stupidity and Neutrality

The so-called "*Stupid Network*" principle was asserted in 1997 [1], arguing that the vision of an Intelligent Network was being replaced by a new perspective, in which networks would provide simple bit delivery and operators would refrain from provisioning fancy services beyond mere connectivity, thus being prevented from interfering in the virtuous relationship between service providers (offer) and users (demand). Thirteen years after its assertion, this vision is still topical.

On the one hand, telecommunications operators are more and more suffering from voice revenues decline, mainly due to the harsh competition resulting from deregulation and from new low-cost VoIP-based players (e.g., Skype). Voice services are nowadays regarded as commodities. Even mobile operators are struggling to preserve their voice margins, while revenue growth is dramatically slowing down.

On the other hand, the broadband access business provided significant incomes to network operators, compensating voice business decline, but indeed seems to represent the living manifestation of the Stupid-Network paradigm. Once the network operator has provided users with IP access to the Internet, users are allowed to go for their service providers of choice, disintermediating the network provider. Whilst this is considered beneficial for the community of users and service providers, it makes a problem for the network operator: how to replace voice revenues with new value-added services, without being bypassed at application level.

In a recent survey [2], the confidence in the capability of the telecommunications business in creating long-term sustainable growth in an all-IP world has been assessed to decrease, as market maturity increase. On the other hand, another study [3] showed that only 23% of the telecommunications services pose strict Quality-of-Service (QoS) requirements, so that the advantage of a network provider in leveraging some performance enhancement is limited to a small share of services.

As an additional issue for network operators, the so-called *"broadband incentive"* problem has become more and more evident: as the broadband access market tends to saturate, revenue growth slows down, while traffic keeps steadily on growing due to the utilization of hungrier network applications (e.g., video) [4]. Thus, the network operator should keep on investing in network capacity without additional revenue generation (due to the well-established flat-fee model). This is clearly untenable from a mere business perspective. Introducing again usage-based charging mechanisms in mature markets most probably would be not considered acceptable [5].

However, network operators should never be allowed to prevent network usage by unwelcome applications (i.e. those not generating revenue for them) thus violating the *"network neutrality*" principle [6]. Since early 2000s, this principle has assumed also legislative implications (at least in the USA), after the idea of establishing policies specifically preventing discriminatory behaviours by broadband access operators.

Since then, this principle has raised several controversies about the boundary between restrictions to discriminatory policies (e.g., against Skype) and operators' rights in differentiating quality of service (e.g., in favour of premium users). A set of behaviours, which an operator should never be allowed to adopt without violating the network-neutrality principle (a.k.a. "impermissible net bias"), is provided in [7], including for instance deliberate packet loss, artificial congestion, port blocking, affiliate favouritism violating fair trade, etc.

In summary, network operators are facing hard times. Traditional businesses (e.g., voice) are declining fast, while new ones (e.g., broadband access) are slowing down. New Internet brands make money out of overlay application networks with no involvement and revenue shares for network operators. At the same time, the network-neutrality principle inhibits adopting defensive approaches, which would prevent unprofitable usage of network resources by users and other providers. Peer-to-peer (p2p) applications are a bold example of this unintended usage of network resources.

978-1-4244-7173-7/10/$26.00 ©2010 IEEE

II. PEER-TO-PEER APPLICATIONS

P2p applications are based on distributed overlay networks among participants on the Internet, who share a portion of their resources (e.g., information, bandwidth, etc.). The term *"overlay"* denotes that the communication among users is governed by application-level algorithms that do not involve the underlying physical network, which provides simple packet routing without knowing logical associations between peers. The term *"peer"* indicates that the overlay network has no hierarchy. In a pure p2p network, there are no dedicated servers supplying information and clients demanding it: peers are both suppliers and consumers of resources.

The p2p approach has been applied to four categories of applications: file sharing, personal communication (voice and instant messages), audio and video streaming, online gaming.

One the most significant examples of *p2p file-sharing* applications is eDonkey/eMule. A recent study [8] pointed out that video is the most popular p2p file-sharing content by both volume and number of files.

P2p personal communication is represented mainly by Skype, reported to have had 8% share of international traffic in 2008 (total 416 billion minutes) [9]. According to [8], Skype's share of VoIP traffic is varying highly with the geographical region: in 2008, from more than 80% in East Europe and Middle East, to 30% in South Europe and 20% in South America.

While users of file-sharing are very tolerant to long download waiting times, in *audio and video streaming* data must be supplied in real time. In p2p streaming, as opposed to unicast from a server, every user provides its uploading capacity transmitting "chunks" of contents to other peers requesting the same content, thus overcoming network resource limitations. From the business point of view, streaming is more attractive than download&play, since it allows attaching time-variable adverts to the video content, thus providing revenues from advertising. Recent ventures like Joost (2006) and Babelgum (2007) have preferred streaming, whilst others still adopt download & play.

Finally, *Multi-player Online Gaming* (MOG) requires information exchange among all players in real time, so that the whole community share a consistent view of the game "reality". While the shared resource is storage in p2p file sharing and bandwidth in p2p video, the main advantage of adopting p2p in MOG is the opportunity to share the CPU power of peers, which process game information concurrently.

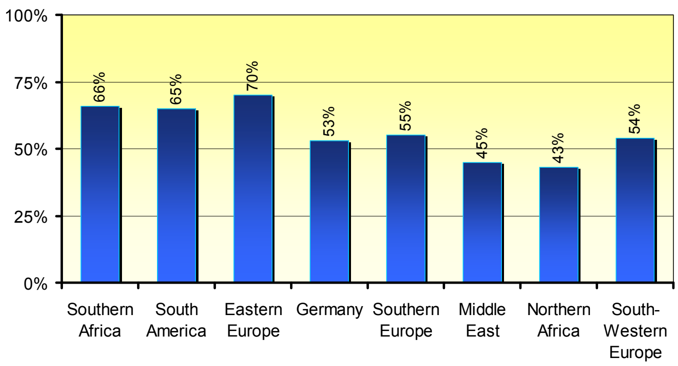

Fig. 1: Relative p2p traffic volume measured in various regions in 2008 [8].

III. IMPACT OF PEER-TO-PEER APPLICATIONS ON THE USAGE OF NETWORK RESOURCES

A closer look on p2p traffic characteristics allows identifying some good reasons for network operators to be concerned about p2p overlays. P2p applications are reshaping the way network resources are used.

A. Impact of P2P on Network Traffic Behaviour

First, p2p applications may take a large share of the overall Internet traffic, although statistical data vary highly depending on where and when they are measured. For example, Fig. 1 compares the relative p2p traffic volumes measured in various geographical regions in 2008 [8] (unfortunately, important markets have been ignored in this study). Moreover, the white paper [10] reported the p2p traffic to account for 44% of residential broadband traffic in North America in 2007. Finally, the study [11] has assessed that about 51% of the 2007 global Internet consumer traffic was taken by p2p, 22% by Internet video to TV and PC, 21% by web, email and other data, 4% by online gaming, 2% by VoIP and video communications.

Second, network operators have no knowledge about the logical associations among peers in a p2p overlay network. Peers may be millions, belong to different ISP networks and be distributed over different countries or continents. The geographical distribution of active peers in the overlay network even changes over time. The decoupling between physical networks and overlay p2p topologies causes a large share of IP traffic to behave unpredictably for the network operator.

Moreover, p2p applications do not take only a large part of the network capacity, but they also change the traffic spatial patterns, as illustrated in Fig. 2, and the proportion between required upstream and downstream capacities.

In traditional client-server applications (e.g., web browsing), upstream traffic is typically much less than downstream. In p2p networks, instead, traffic is mostly balanced or the opposite even holds (see e.g. Fig. 3 [10]). To this regard, it is interesting to notice that the average downstream throughput, available to a single peer in a p2p network, is limited by the average maximum upstream capacity provided by the physical network to all other peers belonging to the p2p network.

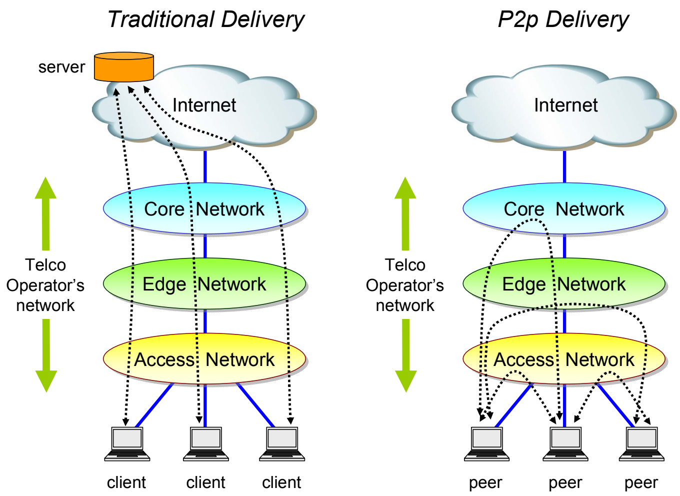

Fig. 2: Change of spatial traffic patterns due to p2p applications.

At network *access*, in particular, p2p applications taking higher upstream capacity is not an issue, as long as the upstream capacity is underutilized (e.g., in a consumer access line). If, on the contrary, the upstream access capacity is shared among many users (e.g., in a corporate or campus access to the Internet), then the p2p traffic generated by some users may subtract valuable transmission capacity to institutional applications.

B. P2P for Streaming?

At the *edge* of the IP network, as illustrated in Fig. 4, operators usually concentrate downstream capacity (i.e., downstream capacity at network edge is lower than the aggregate downstream capacity of all access lines), based on the assumption that not all customers require to receive traffic with the access line speed simultaneously (statistical gain).

On the other hand, more advanced applications like TV multicast streaming over IP (IPTV) require higher downstream capacity at edge level, since many users can be watching TV at the same time, not to mention that the benefit from multicasting at edge level is limited by the fact that users can be watching different channels. Therefore, because of IPTV the edge is dimensioned with lower statistical gain, to provide at least the same video quality as in legacy broadcasting.

The adoption of p2p video streaming might help to reduce investments for downstream capacity at edge level, because each user could provide video content to neighbouring users to avoid replicated traffic flows to be sent over the core network. However, this result would be achieved only *if the operator could shape the p2p overlay network topology to reflect, somehow, the physical network topology*. The goal is to avoid that the streaming content is provided by remote users from networks of other operators, as it happens most commonly,

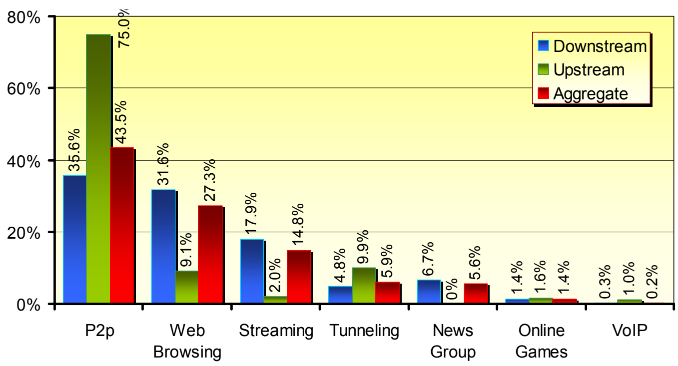

Fig. 3: Comparison between downstream and upstream traffic shares of different traffic types (sample: a number of leading North-American service providers, residential broadband traffic, 2008) (data from [10]).

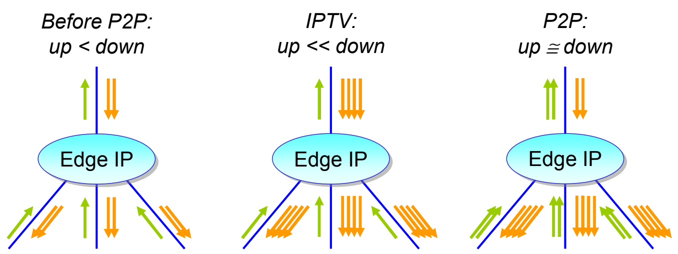

Fig. 4: Evolution of traffic balance at the IP network edge.

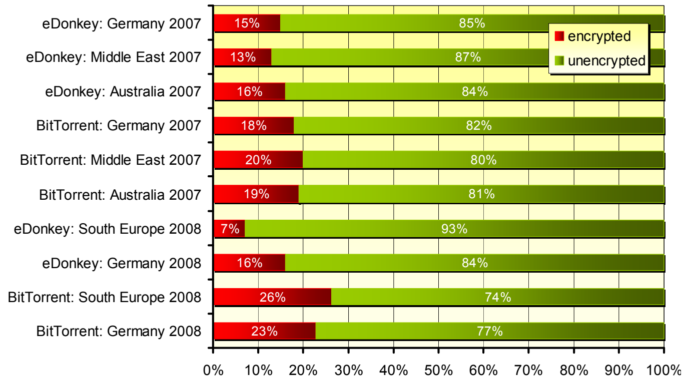

Fig. 5: Proportion of encrypted and unencrypted BitTorrent and eDonkey traffic measured in various regions in 2007 and 2008 (data from [8]).

because the topologies of the overlay and physical networks are fully decoupled. Having p2p traffic routed independently from the physical topology brings inefficiency of network engineering and higher link loads at edge and peering levels.

In particular, this made the traffic between ISPs at peering points to grow uncontrollably in the last few years. For example, Vodafone Iceland found in 2007 that p2p was accounting for 60% of incoming traffic and 80% of outgoing traffic along its undersea Internet links. As soon as they began throttling the total capacity available to p2p users, the network performance improved substantially [12].

IV. Defensive Approaches to Prevent Unintended Usage of Network Resources by P2P Applications

Based on these considerations, it is evident that network operators are very much interested to block or discourage p2p traffic in their network, seen as a threat for their business.

A. Identifying and Blocking P2P Traffic and Users

Several methods for identifying p2p flows in aggregate traffic have been proposed so far, but this is still an open issue and a topic of research (anyway, beyond the scope of this paper).

Traffic inspection methods are computationally intensive and require customization for every p2p application. In addition, they may be defeated by traffic encryption or obfuscation, although this may be ineffective against modern traffic identification systems. However, a recent study [8] has estimated that the proportion of encrypted p2p traffic is small and did not increase significantly from 2007 to 2008 (see Fig. 5).

More intrusive approaches of p2p traffic identification are based on "spy" peer clients participating to the p2p network and recovering information about other peers, network topology and shared contents. This enables the network operator not only to identify p2p traffic, but also to identify customers, usage patterns and contents, thus raising major privacy issues.

Once p2p traffic or users have been identified, however, operators *should not be allowed to block them sharply and violate the "network neutrality" principle* [6]. To this regard, we like to cite the press statement of the USA FCC Chairman [13]: "If we aren't going to stop a company that is looking inside its subscribers' communications (…), blocking that communication when it uses a particular application regardless of whether there is congestion on the network, hiding what it is doing by making consumers think the problem is their own, and lying about it to the public, what would we stop?".

B. Discouraging P2P Applications

Network operators may still use p2p traffic information, formally without infringing privacy and network neutrality principles, by *discouraging* customers in various ways from using p2p applications. Several followed this approach.

For example, Bell Canada, against alternative ISPs protesting because it throttled P2P traffic, pointed out that it spent nearly $500 million in 2008 and $110 million in 2007, just to deal with network congestion [14]. In USA, AT&T has recently announced to FCC that it does not allow mobile subscribers to use P2P applications, as they have detrimental effect of network capacity. The company issues a written warning to any customers found using p2p applications and, if the breach continues, it terminates the user's contract [15].

It is possible to adopt various countermeasures to discourage customers from using p2p applications without formally violating network neutrality principles, namely:

- traffic management policies shaping p2p traffic (i.e., giving p2p a lower priority at pre-defined times or when resource contention occurs);
- capacity planning policies restraining upstream capacity upgrades as demand grows;
- commercial policies discouraging the usage of upstream capacity, for example re-establishing usage-based charging mechanisms, or asking extra fees for allowing p2p traffic, or capping p2p traffic (i.e., setting a threshold on the volume of p2p traffic exchanged say per month).

Thus, the network operator, without explicitly blocking p2p traffic, can make things somehow more difficult for p2p than for other applications. In other words, p2p traffic can be legitimately allowed to traverse the network, *as long as it can be restrained to cause minimum disturbance to more privileged applications*.

Since the network operator is not participating to the business behind the p2p application, it should be granted the right to differentiate traffic treatment in favour of more remunerative services. In UK, for example, some ISPs threatened to start throttling iPlayer traffic unless the BBC offered them cash to fund network upgrades. In April 2008, BBC executives threatened to shame publicly ISPs who were doing so and even received support from Ofcom [16].

V. Turning P2P Applications to a Future Opportunity for Network Operators

Once the operator has detected p2p traffic over its network, it can use this information in an alternative way, trying to make the best business out of it.

A. Endorsing the P2P Paradigm

The p2p paradigm can be endorsed, while driving the way network resources are used. For a network operator, adopting a p2p architecture would mean to regain control, somehow, of the way its network resources are used by customers.

New services based on p2p may be also launched. A bold example is represented by p2p video streaming services (e.g., AOL In2TV, launched in 2006 and based on the Kontiki Peer-Assisted Content Delivery Technology platform).

The user's p2p application may retrieve contents from both other peers and network operator's servers distributed in edge network nodes. The key idea is that content should be retrieved from the closest point where it is available (a.k.a. the *"geolocalization of content"* concept).

This way, network resource usage is optimized at all network layers (edge, backbone, interconnection), allowing the ISP to minimize traffic exchange with other operators (i.e., to reduce interconnection costs). Moreover, this p2p service can be offered also to users of other ISPs, who will then exchange content with users of the ISP providing the p2p service, bringing new p2p traffic at interconnection points.

Therefore, a network operator provided with attractive contents, by endorsing the p2p paradigm, can leverage the intrinsic scalability of p2p networks (better performance with less investments), optimize the use of its own network resources (operational cost reduction), launch new value-added services (new revenues from selling content and advertising) and participate to the new arena of application-layer services, avoiding to be confined to the "bit-pipe" provider role. All this would be achieved by embracing the p2p paradigm in a legal manner (i.e., without copyright infringements or privacy violation) and in full compliance with network neutrality principles.

B. Proactive Provider Participation for P2P (P4P)

What outlined above depicts more or less the idea behind the DCIA (Distributed Computing Industry Association) initiative called Proactive Provider Participation for P2P (P4P) [17], defined as "A set of business practices and integrated network topology awareness models designed to optimize ISP network resources and enable P2P based content payload acceleration".

The mission statement of the DCIA P4P Working Group [18] depicts a scenario in which researchers, ISPs and p2p software distributors work together to envisage new mechanisms to maximise content distribution speed, while jointly optimising network resources. The underlying principle is that ISPs themselves supply p2p applications, providing contents on the shortest routes between peers by means of network policies based on, say, the relative distance between peers and the network usage. This way, network operators are back again in condition to engineer their networks and optimize resource usage, still letting p2p applications to find the shortest path to content. A core group of 19 member companies (including large operators AT&T, Telefonica, Telecom Italia and Verizon), with 18 observing members, sponsors this P4P initiative.

In March 2008, the Yale University, Pando Networks (a p4p software provider participating to the DCIA core group), Telefonica International Wholesale Services and Verizon carried out p4p field trials on Telefonica and Verizon networks, providing p4p software with network map information [19][20]. Telefonica also declared (April 25, 2008) to look forward to expanding their tests, to involve larger operations in Europe and elsewhere as the next phase of their participation in the P4P Working Group.

These trials demonstrated benefits for both p2p users and ISPs [19][20]. P2p FTTH clients experienced an average improvement of download rates above +200% from other FTTH clients and about +23% from all clients. From the ISP's viewpoint, tests demonstrated a dramatic drop in the number of average hop count (from 5.5 to 0.89), an average peering link load decrease about -42% (outbound) and -35% (inbound) and an internal backbone load decrease about -71%.

This DCIA P4P initiative is the first practical attempt to match the needs of p2p users and the operational efficiency requirements of network operators. It is difficult to judge whether the p4p technology is the best solution. However, it is clear the need for allowing network operators to tie the behaviour of p2p applications to network awareness, so that the user experience can be improved while maximising network operational efficiency.

VI. Conclusions

Nowadays, network operators face an awkward situation. Used to the traditional business of voice services, in which revenues are proportional to network resource usage, they are now facing an exponential growth of traffic over their IP networks with declining revenue growth for broadband services.

Today, p2p applications take a large share of the overall Internet traffic. The decoupling between physical and overlay topologies causes this traffic to behave unpredictably for the network operator, hampering network planning strategies. Moreover, the network operator does not participate to the business behind the p2p application, with limited margins and revenues from traffic delivery.

At the same time, the network-neutrality principle inhibits operators adopting hard defensive approaches, which would prevent unprofitable usage of network resources by users and other providers. Some network operators, without explicitly blocking p2p traffic, attempt to discourage customers from using p2p applications, still without formally violating network neutrality principles.

Alternatively, the network operator might endorse the p2p paradigm, to improve efficiency in traffic engineering and even launch new p2p services. In order to stay in business, without relying on external financial support like public financing, network operators should climb the value chain and provide more than mere connectivity, harnessing p2p traffic and subduing the underlying business model to their advantage while putting network resources back under control.

This paper presented an example of this strategy: providing content distribution services based on p2p architectures (e.g., p2p video streaming services), both for recovering network efficiency and for launching new value-added services.

For the first time, the DCIA P4P initiative has put forward a practical technical solution to this aim, based on making p2p applications aware of both network topology and network resources availability. This way, the p2p overlay and network topologies are coupled, matching the efficiency objectives of network operators with the quality-of-experience requirements of p2p users.

Whether this strategy will be successful or not largely depends on the capability of network operators to embrace a commercial p2p-based service proposition, together with a suitable technical solution allowing an increased operational efficiency. Doing so, a network operator will be still able to play a relevant role, beyond pure bit delivery, in the business models of future applications like p2p IPTV.

References

[1] D. S. Isenberg, "The Rise of the Stupid Network", *Computer Telephony*, Aug. 1997, pp. 16-26.

[2] Simon Torrance Limited for Telco, "Telco 2.0™ Market Study", 2007.

[3] G. Salisbury, "Trends in Telecommunications Services - Does the Traditional Operator Have the Advantage?", Detecon Managem. Rep., 2006.

[4] "The Broadband Incentive Problem", Broadband Working Group, MIT Communications Futures Program. Available: http://cfp.mit.edu/docs/incentive-wp-sept2005.pdf.

[5] A. M. Odlyzko, "Paris Metro Pricing for the Internet", *Proc. ACM Conf. on Electronic Commerce (EC'99)*, 1999.

[6] T. Wu, "Network Neutrality, Broadband Discrimination", *Journal of Telecommunications and High Technology Law*, Vol. 2, p. 141, 2003.

[7] Rob Frieden, "Network Neutrality or Bias? Handicapping the Odds for a Tiered and Branded Internet", Pennsylvania. State Univ., Sept. 2006.

[8] H. Schulze, K. Mochalski, "Internet Study 2008/2009 - The Impact of P2P File Sharing, Voice over IP, Instant Messaging, One-Click Hosting and Media Streaming on the Internet" and "Internet Study 2007", IPOQUE reports, 2008-09.

[9] Telegeography, "Telegeography Report 2008", Primetrica Inc, 2008.

[10] Sandvine, "2008 Analysis of Traffic Demographics in North-American Broadband Networks", White Paper, 2008.

[11] Cisco, "Cisco Visual Networking Index – Forecast and Methodology, 2008–2013", White Paper, 2009.

[12] Nate Anderson, "ISPs to BBC: We will throttle iPlayer unless you pay up", *Ars Technica*, 2007. Available: http://arstechnica.com/old/content/ /2007/08/isps-to-bbc-we-throttle-iplayer-unless-you-pay-up.ars.

[13] K. J. Martin, USA FCC Chairman, Press Statement on *"Formal Complaint of Free Press and Public Knowledge Against Comcast Corporation for Secretly Degrading Peer-to-Peer Applications; Broadband Industry Practices, WC Docket No. 07-52"*, 1 Aug 2008. Available: http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-284286A2.pdf.

[14] Nate Anderson, "Bell Canada: congestion numbers look low, but actually aren't", *Ars Technica*, 2008. Available: http://arstechnica.com/news.ars/post/20080625-bell-canada-congestion-numbers-look-low-but-actually-arent.html.

[15] Telegeography, "AT&T tells FCC no P2P", 2008. Available: http://www.telegeography.com/cu/article.php?article_id=24376.

[16] Nate Anderson, "Who should pay? BBC, UK ISPs argue over iPlayer traffic", *Ars Technica*, 2008. Available: http://arstechnica.com/news.ars/post/20080425-who-should-pay-bbc-uk-isps-argue-over-iplayer-traffic.html

[17] Open P4P Working Group, "P4P Home Page", March 2008. Available: http://www.openp4p.net.

[18] Distributed Computing Industry Association, "DCIA P4P Working Group Mission Statement". Available: http://www.dcia.info/documents/P4PWG_Mission_Statement.pdf.

[19] Open P4P Working Group, "P4P Field Tests", March 2008. Available: http://www.openp4p.net/front/fieldtests.

[20] Pando Networks, *"Results of Second Round of P4P Field Trials with Pando Networks Yield Greater Improvement in Performance of Broadband Networks"*, Nov. 2008. Available: http://www.pandonetworks.com/node/166.