A Survey of P2P Traffic Management Approaches: Best Practices and Future Directions

R. Dunaytsev, D. Moltchanov, Y. Koucheryavy, O. Strandberg, and H. Flinck

Abstract—**Over the last decade, we have witnessed the tremendous growth of peer-to-peer (P2P) file sharing traffic. Built as overlays on top of the existing Internet infrastructure, P2P applications have little or no knowledge of the underlying network topology, generating huge amounts of "unwanted" inter-domain traffic. Bandwidth-hungry P2P applications can easily overload inter-domain links, disrupting the performance of other applications using the same network resources. This forces Internet service providers (ISPs) either to continuously invest in infrastructure upgrades in order to support the quality of service (QoS) expected by customers or use special techniques when handling P2P traffic. In this paper, we discuss the best practices and approaches developed so far to deal with P2P file sharing traffic, identifying those that may provide long-term benefits for both ISPs and users.**

Index Terms—**File sharing, peer-to-peer, quality of service, traffic management.**

I. Introduction

For a long time, the reference model of data exchange in the Internet was the client/server model, resulting in the so-called "downstream paradigm", where the vast majority of data is being sent to the user with low traffic load in the opposite direction. As a consequence, many communication technologies and networks have been designed and deployed keeping this asymmetry in mind. The best-known examples are the Asymmetrical Digital Subscriber Line (ADSL) and Data Over Cable Service Interface Specification (DOCSIS) technologies. For instance, ADSL2 and ADSL2+ provide a downstream rate of up to 24 Mbps and an upstream rate of up to 1 Mbps [1]. In fact, bandwidth asymmetry with high data rates in the downstream direction together with low-rate upstream links fits well the environment with dedicated servers and conventional users. Everything changed with the arrival of Napster and other peer-to-peer (P2P) file sharing systems in the late 1990's and early 2000's. In such systems, all participants (although to different extents) act as both content providers and content requestors, thus transmitting and receiving approximately equal amounts of data. Hence, uplink and downlink data flows tend to be symmetric [2]. Since then, P2P networks have experienced tremendous growth, and for several years P2P file sharing traffic used to be the dominant type of traffic in the Internet [3] [4] [5]. The situation has changed dramatically in recent years with the increasing deployment of multimedia applications and services such as Flash video, IPTV, online games, etc. In 2010, global Internet video traffic has surpassed global P2P traffic [6]. According to recent studies [6] [7] [8], P2P traffic is growing in volume, but declining as a percentage of overall Internet traffic. However, the prevalence of real-time entertainment traffic (Flash, YouTube, Netflix, Hulu, etc.) with a decrease in the fraction of P2P file sharing traffic is usually the result of cheap and fast Internet access and is more typical for mature broadband markets, while many emerging broadband markets are still in a phase in which P2P file sharing accounts for a large (or even a dominant) portion of global traffic [8]. In any case, P2P file sharing is still fairly popular among users and continues to be one of the biggest consumers of network resources. For instance, it is expected that P2P file sharing traffic will reach 8 exabytes per month by 2015, at a compound annual growth rate (CAGR) of 15% from 2010 to 2015 [6].

From the early days of P2P file sharing systems, P2P traffic and its impact on the performance of other applications running on the same network has attracted the attention of the academic community and Internet service providers (ISPs), and continues to be a hot research topic (e.g., see [9] and references therein). What makes this type of traffic so special? Let us consider the most distinctive features of P2P file sharing from the ISP's point of view.

1) P2P file sharing applications are bandwidth-hungry: Network applications (and hence traffic sources) can be classified into 2 basic types: constant bit rate (CBR) and variable bit rate (VBR). CBR applications generate data traffic at a constant rate and require a certain bandwidth allocation in order to operate successfully and support the desired quality of service (QoS). At the same time, allocating bandwidth above the requirement does not improve the user satisfaction. VBR applications generate data traffic at a variable rate and are typically designed to quickly discover and utilize the available bandwidth. In general, the more the bandwidth, the better the user-perceived QoS. However, in order to provide scalability to a large number of clients accessing the system simultaneously, both CBR and VBR servers usually limit the maximum data rate per user. As a result, this effectively places an upper bound on the amount of data that can be transmitted per unit time and thus the bandwidth used by an individual user.

As reported in [7] [8] [10], BitTorrent is the most popular P2P file sharing system today. In contrast to client/server systems, BitTorrent is more robust and scalable: as more users interested in downloading the same content join an overlay network, called a torrent or a swarm, the download rate that is achieved by all the peers increases [11]. With BitTorrent, when multiple users are downloading the same

Manuscript received May 15, 2012.

Roman Dunaytsev is with the Space Internetworking Center (SPICE), Electrical and Computer Engineering Department, Democritus University of Thrace, Greece (corresponding author's phone: +30-6995119041; fax: +30-2541079554; e-mail: roman.dunaytsev@spice-center.org).

Dmitri Moltchanov and Yevgeni Koucheryavy are with the Department of Communications Engineering, Tampere University of Technology, Finland (e-mail: moltchan@cs.tut.fi, yk@cs.tut.fi).

Ove Strandberg and Hannu Flinck are with Nokia Siemens Networks, Espoo, Finland (e-mail: ove.strandberg@nsn.com, hannu.flinck@nsn.com).

file at the same time, they upload pieces of the file to each other. Instead of relying on a single server, this mechanism distributes the cost of storing and sharing large amounts of data across peers and allows combining upload capabilities of multiple peers into the compound download rate observed by a user [12]. As a rule, if there are enough active peers, the maximum achievable throughput of a user is mostly limited by either congestion somewhere in the ISP's network (if any) or the last-mile bandwidth of the user. In well-provisioned networks, this results in higher data rates (e.g., up to 80 Mbps over a 100 Mbps Ethernet link) and potentially larger amounts of data that can be sent per unit time, when compared to traditional client/server applications and services such as the World Wide Web and IPTV. In addition, since P2P file sharing requires very little human intervention once it is initiated, some users tend to run P2P programs 24/7, generating huge amounts of traffic.

High-speed P2P traffic interferes with other traffic on the same network, degrading the performance of delay-sensitive applications such as multimedia streaming, online games, and VoIP. Poor application performance during congestion causes low customer satisfaction and aggravates subscriber churn, leading to a decline in service revenues. In turn, this forces ISPs either to continuously invest in infrastructure upgrades in order to support the QoS expected by customers or use special policies when handling P2P traffic.

2) P2P file sharing applications are topology-unaware: P2P file sharing applications establish overlays on top of the Internet infrastructure with little or no knowledge of the underlying network topology (see Fig. 1). In P2P networks, data are often available in many equivalent replicas on different hosts. However, the lack of topology information leads P2P applications to make random choice of peers from a set of candidates. For example, it is common for a peer that wants to download some data (music, movies, software, etc.) to choose sources randomly, possibly picking one located on the other side of the world (see Fig. 2). Such random selection ignores many peers that are topologically closer and therefore could provide better performance in terms of throughput and latency. Moreover, this leads to inefficient utilization of network resources, significant amounts of costly inter-ISP traffic, and congested inter-domain links (both incoming and outgoing). Thus, "better-than-random" peer selection would be beneficial for both users and ISPs.

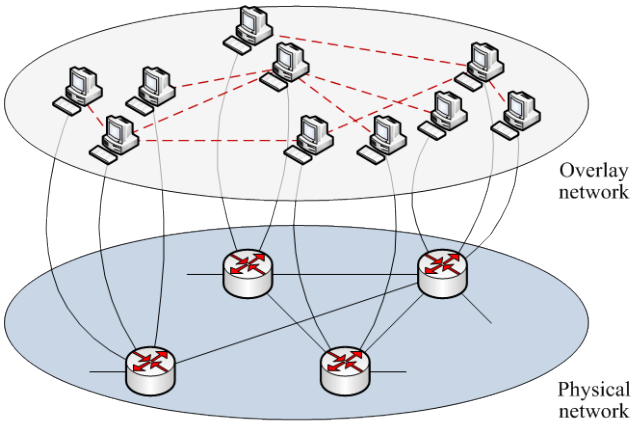

Fig. 1. An overlay network is a virtual network of end systems and logical links built on top of an existing physical network. The overlay is a logical view that might not directly mirror the physical network topology.

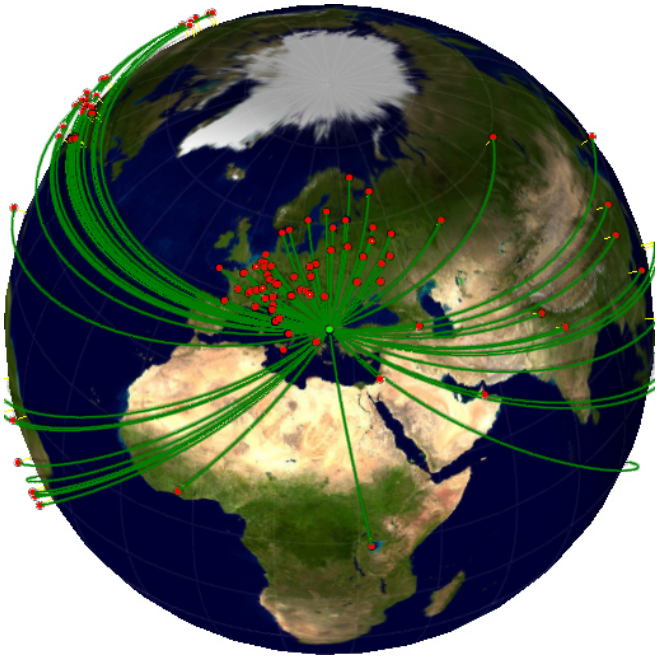

Fig. 2. A user in Greece downloads the latest version of Ubuntu. Most of the peers are from other countries and even the other side of the globe. Does it mean that no one is seeding this distro in Greece? Of course not! It is just a consequence of the random peer selection.

3) P2P file sharing brings no additional profit to ISPs, only additional costs and legal headaches: Nowadays, some ISPs tend to charge content providers and content delivery networks (CDNs) additional fees for either carrying high-bandwidth content over their networks or providing QoS guarantees for premium traffic. Meanwhile, end users are typically billed for Internet access based on flat-rate pricing, so ISPs do not generate any additional revenue from delivering P2P traffic to/from their customers.

High-speed P2P data transfers may also increase the cost of traffic exchange with other ISPs. The vast majority of ISPs (except a small group of Tier 1 ISPs that rely completely on settlement-free peering) are charged for traffic transit according to the 95th percentile billing model, which works as follows. Transit providers poll customer interface ports at regular intervals (typically, every 5 minutes) throughout the billing cycle. Each sample contains the number of bytes transmitted to and received from the customer. These samples are converted into data rates measured in Mbps, and then sorted from largest to smallest. The transit provider discards the top 5% of the lines from the list, and bills the customer for the next line which is called the 95th percentile. Thus, customer ISPs can transmit data as fast as possible for 36 hours per month free of charge. With P2P file sharing, users can contribute a lot to the ISP's transit cost by remaining active most of the time and transferring data at high rates. Moreover, peers are randomly distributed in the Internet and operate on an irregular basis, which makes predicting traffic patterns and optimizing transit arrangements very difficult, if not impossible.

P2P file sharing applications not only consume network bandwidth and increase transit costs but also encourage illegal distribution of copyrighted material among users. According to [13], 63.7% of all BitTorrent traffic is estimated to be non-pornographic copyrighted content shared illegitimately (the copyright status of pornography is more difficult to identify but the majority is believed to be

copyrighted and shared illegitimately too). For other P2P file sharing applications, such as eDonkey and Gnutella, about 86.4% traffic is estimated to be non-pornographic and infringing. The findings of a study of P2P exchanges of copyrighted material on a college campus show that at least 51% of students living on campus engaged in P2P file sharing, and at least 42% attempted to transfer copyrighted material [14]. As a result, ISPs receive numerous cease-and-desist letters from lawyers about copyright infringements (e.g., see [15]).

However, it should be emphasized that, despite all these issues, P2P file sharing has been a key driver for broadband adoption in the early 2000's, and ISPs do not want to lose customers who wish to use these applications today. Besides the above-mentioned problems, ISPs have to cope with many other challenges such as:

--*Tremendous growth of Internet traffic:* According to [6], global Internet traffic is expected to quadruple in the next few years, reaching nearly 1 zettabyte per year by the end of 2015. Although such growth can be partially compensated for by recent advances in transmission technologies (most notably in optical communications for wide and metropolitan area networks), it may still pose certain problems for small and medium-sized ISPs, both technically and financially.

--*Rich mix of applications with different QoS requirements in terms of delay, delay jitter, and bandwidth:* In such an environment, providing different QoS levels to different applications while maintaining efficient use of expensive network resources becomes an issue of special importance for all ISPs.

--*Highly competitive market of Internet access services:* In order to be successful in the long run, ISPs need to maximize their profits and minimize costs. Moreover, to avoid subscriber churn, ISPs must provide their customers with consistent QoS that is at least perceived to be no worse than that provided by competitors. Among other ways, these goals can be achieved by managing traffic on a per-application and per-subscriber basis, replacing flat-rate pricing by multi-tiered one with several levels of QoS based on different subscription rates to meet the requirements of applications and users, deferring investments in infrastructure upgrades, and reducing inter-ISP transit expenses.

A number of approaches have been developed over the years to deal with P2P file sharing traffic and to mitigate its impact on the network and application performance or, in other words, ISPs and their customers (see [16] [17] [18] [19] and references therein). These approaches can be broadly classified into the following categories: acquiring more bandwidth (also known as overprovisioning), blocking P2P traffic, implementing bandwidth caps (widely referred to as traffic quotas), bandwidth management, and localizing P2P traffic. The reminder of the paper provides an overview of these approaches, including their strengths and weaknesses. In Section 2, we consider the conventional strategies like bandwidth overprovisioning, blocking P2P traffic, implementing bandwidth caps, and bandwidth management. Section 3 presents an overview of recently proposed P2P-friendly solutions, with the main focus on P2P caching and biased choice of peers. Conclusions are drawn in Section 4. The paper consolidates the main findings and conclusions arising from recent studies to provide guidelines on the most effective strategies in P2P traffic management.

II. Conventional Approaches

A. Acquiring More Bandwidth

The Internet carries many types of traffic, each of which has different characteristics and requirements. Many file transfer applications, including P2P file sharing, require that some quantity of data be transmitted in a reasonable amount of time but can tolerate variable delays and packet losses, which frequently occur during periods of network congestion. Other applications, like VoIP and multimedia streaming, require that the one-way delay between endpoints should not exceed a certain threshold and have unacceptable performance when long queuing delays occur. If enough bandwidth is available, the best-effort service, where all packets and flows are treated in the same way, meets all these requirements. If bandwidth is insufficient, real-time traffic suffers from congestion. Acquiring more bandwidth, or overprovisioning, is the most straightforward way to alleviate congestion in the network and address the QoS issue. As bandwidth became cheaper and communication technologies evolved, overprovisioning of network resources, in the hope that the full capacity will not be reached (at least for a while), has become a common approach among ISPs. Very often, overprovisioning comes from the fact that the link capacity must be selected from a set of discrete values (e.g., 155 Mbps or 622 Mbps), which inevitably leads to acquiring more bandwidth than currently needed. As compared to the other approach, known as service differentiation, where different packets and flows are treated differently in order to provide QoS guarantees and give preferential treatment to some types of traffic, overprovisioning has a number of benefits. Firstly, it is more difficult to control a network that does not have enough bandwidth than a well-provisioned one. This is caused by additional complexity required to provide service differentiation in the network and additional costs associated with deploying and managing QoS solutions. Secondly, overprovisioning not only allows to meet the current needs, but also leaves enough room for future traffic growth.

For many years, overprovisioning is the de facto standard for capacity engineering and network planning. For example, extensive measurements of Sprint's backbone network between September 2000 and June 2001 show that most links in the core network are not highly loaded: 69% of links never experience 30% load even once during their lifetime [20]. The results of a private large-scale survey conducted in 2007 also indicate that overprovisioning is a common choice of ISPs, while the overprovisioning ratio varies widely, depending on the underlying network topology and technology, the number of users and anticipated variation in traffic loads, etc. [21]. According to [22], the rule-of-thumb for backbone links is to (a) upgrade when the link utilization reaches about 40% and (b) ensure that the maximum utilization of the link does not exceed 75% under failure scenarios. Router vendors also insist that link utilization levels should not exceed about 80% because otherwise routers can slow down or even crash. As a result, backbone networks contain large amounts of bandwidth that is not currently being used [23]. For instance, when telecommunications companies run fiber-optic cable, they usually run 2 or 3 times the amount of fiber that they actually need [24]. These spare strands of fiber are often referred to as "dark fiber". Hence, fiber-optic networks have excellent

expansion capabilities assuming some dark fiber is placed in the cable alongside of the operational strands. For example, as it follows from [25], most National Research and Education Networks (NRENs) in Europe have access to dark fiber and can increase capacity easily and economically whenever required. Moreover, 85% of NRENs state that they either prefer to overprovision their networks or see no need for QoS engineering [26]. The report also claims that the lower the congestion index score on the backbone of a NREN, the more likely this NREN is to be adopting overprovisioning.

As noted in [27], ISPs do their best to build networks such that their backbone and inter-domain links are uncongested. This is because the access network infrastructure usually represents the greatest component of the total network cost, so it only takes a relatively small fraction of the revenue from each customer to overprovision the core network by adding more bandwidth. In the late 1990's and early 2000's, most capital investment was devoted to upgrading and deploying long-distance backbone links. More recently, investment has been tied more to access networks, including investment in upgrading copper networks, cable television networks, and new fiber-optic access networks (also known as FTTx, Fiber To The x, where "x" can be Node, Cabinet, Building, Home, depending on the degree of optical fiber penetration) [28]. Facilities-based competition has driven ISPs to upgrade or enhance their infrastructure for faster speeds, better QoS, and provide larger amounts of bandwidth to their customers. Fiber, cable, and DSL deployments grew at a CAGR of 25%, 9%, and 7% between 2007 and 2009, respectively (see Table 4.2 in [28]).

Today, broadband network coverage continues to improve. For example, most of the 34 countries, members of the Organization for Economic Cooperation and Development (OECD), report nearly full DSL network coverage (see Table 4.14 in [28]). An extensive survey, covering 686 offerings of fixed broadband services from 101 ISPs across all OECD countries revealed that the average advertised speed increased by more than 20%, from up to 30.5 Mbps in October 2009 to up to 37.5 Mbps in September 2010 [28]. Of course, actual speeds are often lower than advertised ones. This is also confirmed by the results of a nationwide performance study of fixed broadband services in the USA [29]. The study examined service offerings from 13 of the largest broadband ISPs using automated measurements of the broadband performance delivered to the homes of thousands of volunteer broadband subscribers during March 2011. One of the main findings is that the majority of ISPs deliver speeds that are generally 80% or better than their advertised rates. The study also provides interesting insights into performance variation by access technology. On average, during peak periods, DSL-based services delivered 82% of advertised download speeds, cable-based services delivered 93%, and FTTH-based services delivered 114%. In the opposite direction, DSL meets 93%, cable meets 108%, and FTTH meets 112% of advertised upload speeds. During peak periods with many consumers online, speeds decreased from 24-hour average speeds by less than 10% for all access technologies (however, results may differ significantly among different ISPs). This suggests that the majority of broadband access networks are massively provisioned with bandwidth (the networks are not congested even during peak hours), whereas the difference between actual and advertised speeds is mostly a DSL-related issue. Indeed, with DSL, the download and upload speeds deteriorate with distance, so the speeds advertised for DSL-based services are dependent on the distance between the switch and the user.

Thus, we conclude that overprovisioning plays a key role in providing QoS over the Internet. As noted in [30], careful network design and bandwidth overprovisioning not only make a network more resilient, but also prevent many problems from occurring and hence eliminate the need for complex mechanisms designed to solve those problems. In [31], by taking Sprint's backbone network as an example, the authors argue that satisfying end-to-end delay requirements as low as 3 ms requires only 15% extra bandwidth above the average data rate of the traffic.

Taking into account the increasing amount of video content coming online [6], bandwidth overprovisioning becomes more important now than ever. However, adding extra bandwidth cannot solve the P2P problem alone, since uncontrolled P2P traffic tends to expand and fill all the available bandwidth, thus requiring frequent and costly infrastructure upgrades. For this reason, as emphasized in [16], overprovisioning as a part of a development strategy is a necessary step, whereas acquiring additional bandwidth to address just the P2P problem is a dead end.

B. Blocking P2P Traffic

Blocking all (or almost all) P2P traffic is another way to eliminate the performance degradation associated with P2P file sharing. As a rule, this is achieved by blocking ports commonly used by popular P2P applications [32]. While this approach allows to substantially reduce bandwidth consumption and to avoid legal headaches for ISPs caused by illegal distribution of copyrighted content via their networks, it has a number of shortcomings. Firstly, it is not so easy to block P2P traffic these days as popular P2P programs, such as uTorrent and BitComet, enable users to select a desired port or randomize port each start. Secondly, there is a certain trend to use P2P systems, especially BitTorrent, for delivering totally legal content such as Linux distros or copyright-free music, movies, games, and software updates [33]. Thirdly, P2P file sharing has become a driver of broadband adoption, so blocking P2P downloading and/or uploading activities can easily lead to decreased customer satisfaction and loyalty and, ultimately, reduced revenue and market share. Hence, this approach is mainly suitable for campus and corporate networks, rather than commercial ISPs with residential customers.

Today, in many colleges and universities, installing and using software for P2P file sharing on computers connected to the campus network is forbidden (e.g., see [34]). However, there is a way to restrict illegal distribution of copyrighted material and to avoid wasting valuable bandwidth, but still allow users to access legal P2P content. This technique is known as "tracker whitelisting" [15]. In BitTorrent, file sharing between peers is assisted by central servers called trackers. Therefore, by permitting access to trusted trackers with legal content and denying access to others, network administrators can mitigate many of the threats posed by P2P file sharing. What is more, users can also participate in creating and updating these whitelists.

C. Implementing Bandwidth Caps

Implementing bandwidth caps is a means to discourage users from consuming excessive amounts of bandwidth. Typically, if a user exceeds the bandwidth cap, the ISP restricts his/her connection speed for a while (e.g., up to the end of the accounting period). The ISP may also offer to purchase additional bandwidth for an extra fee and hence to recover some costs caused by heavy users. Within North America specifically, 1% of the subscriber base generate about 50% of upstream and 25% of downstream traffic [7]. Unfortunately, this approach allows to achieve bandwidth savings in the long run (in terms of traffic per month) but offers very little on a short time scale (in terms of traffic per second/minute/etc.). Therefore, implementing bandwidth caps cannot effectively prevent and manage congestion in the network during peak hours [35]. Additionally, this approach can limit only the total bandwidth consumption but lacks the granularity to deal with P2P file sharing traffic as a separate phenomenon [16]. The situation can be slightly improved by using more sophisticated schemes such as imposing caps on specific applications, time of the day, day of the week, throughput in terms of Mbps, etc. In [35], it is proposed to use sliding windows for measuring the bandwidth usage instead of calendar-dependent traffic quotas. That is, the sliding window count adds the traffic from the latest historic period and subtracts the traffic from the oldest one. But, as pointed out in [35] [36], a common drawback of all these caps is that many users do not understand them well and tend to blame the ISP rather than themselves when their service is degraded or interrupted. Eventually, customer confusion and frustration aggravate subscriber churn, especially if other ISPs do not implement such caps.

Recently, bandwidth caps have become less frequent in fixed broadband networks, while continue to be quite common in wireless networks, where the use of smartphones and tablets is starting to challenge the overall network capacity (e.g., see [6] [37]). According to a pricing survey conducted in September 2010 [28], only 29% out of 686 fixed-broadband offerings had monthly caps on the amount of traffic which users can download or upload, compared to 36% in September 2008. In particular, the fraction of offerings with bandwidth caps decreased from around 40% to 32% for DSL-based services and from approximately 31% to 20% for cable-based ones. At the same time, the fraction of fiber-based offerings with caps increased from 8% to 26%. However, most of OECD countries (20 out of 34) had no caps at all among their broadband offerings. Moreover, bandwidth caps tend to rise with time, while monthly subscription prices remain the same or even decline slightly. For example, Shaw Communications, a Canadian ISP, has increased the monthly caps of its broadband offerings as follows: from 75 GB for $47/month (for December 2010) to 125 GB for $45/month (for April 2012), from 125 GB for $57/month to 200 GB for $55/month, etc. At the end, it is worth mentioning that broadband prices have been continuously falling over the last decade across all countries, while connection speeds have been getting faster.

It is interesting to notice here that imposing limitations on the amount of data a user can generate during a billing cycle contradicts the current trend to use the Internet as the common carrier for all kinds of services, especially video [6]. Today, communication services are frequently sold as bundles, consisting of broadband Internet, telephone, and television services, often referred to as “triple-play”. As reported in [8], the majority of real-time entertainment traffic on North America’s fixed access networks is destined not for laptop or desktop computers, but for game consoles, TVs, tablets, etc. A typical situation, described in [8] as “multiple screens drive multiple streams”, arises when several persons living in a household enjoy a number of TV programs simultaneously. In addition, when the same video is available in multiple bitrates, users tend to select a high-definition format with the highest bitrate possible. This results in a huge amount of traffic per bundle with HDTV. Then, if bandwidth caps are set too high, only minimal (if any) bandwidth savings can be achieved in the network. If bandwidth caps are set too small, there is a risk of inconsistency between the cap and the average traffic load generated per month, resulting in annoying service interruptions. This may also result in extra charges on the subscriber’s bill. Last but not least, with game consoles and TVs, tracking the total bandwidth usage in order to avoid exceeding traffic quotas is a nontrivial task. These issues can be partly addressed by alerting users when they reach certain threshold values (e.g., 25/50/75% of the cap) and not to impose extra charges until they have exceeded the caps several times.

Due to long-term contracts (e.g., for 6 months or more), bundled services are a promising strategy for ISPs to increase customer loyalty and reduce churn, which is a major issue during economic downturns. However, forcing customers to “watch the meter” instead of watching TV may seriously affect the practical usability and user-friendliness of these services, thus slowing down their adoption process. Let us illustrate this with an example. Suppose there are 2 individuals (e.g., husband and wife) in a household. As it follows from [38], watching TV is the most popular leisure activity in the USA and occupies, on average, 2.73 hours per day. According to measurement results of IPTV traffic collected from an operative network [39], video stream bitrates range from 2.5 Mbps to 4 Mbps. Since bitrates tend to increase rather than to decrease with time, let us assume that the bitrate of video streams in our example is 4 Mbps. Thus, watching 2 TV programs simultaneously will result in 2 concurrent data streams, each with a rate of 4 Mbps, and about 290 GB of traffic per month. Note that this value does not include traffic caused by other applications such P2P file sharing and online games. At the same time, average bandwidth caps in OECD countries go from a few GB to several hundred GB per month (see Fig. 7.30 in [28]). For instance, Shaw Communications offers triple-play, including HDTV, with the following caps (for April 2012): 125 GB for $124.90/month, 200 GB for $149.90/month, and 400 GB for $199.90/month. It is easy to see that 2 out of these 3 bundles have caps that are much smaller than the monthly traffic volume in our example. As a result, users cannot totally rely on such bundles as a substitute for conventional TV, which prevents the adoption of these services and reduces the revenue the ISPs could extract from them.

Similarly to [40], we conclude that bandwidth caps is a crude tool when it comes to targeting potentially disruptive applications and heavy users. These caps (if applied) should be based on better understanding of real usage patterns in order to avoid punishing the wrong users.

D. Bandwidth Management

Compared to the previous approach, bandwidth management is a more flexible strategy, where P2P traffic can either be dropped or marked as low-priority before being admitted into the network. As a result, certain applications and/or users get preferential access to the bandwidth, while others get less or none at all. In general, this is accomplished in 3 steps: Deep Packet Inspection (DPI), priority assignment, and differential treatment of packets and flows.

DPI is a relatively new technology in the field of traffic management and QoS provisioning, but has been around for a long time in firewalls, spam filtering, content caching, intrusion detection and prevention [41]. Today, DPI becomes a mandatory element of service differentiation in the Internet since inspecting just packet headers and port numbers does not provide a reliable application classification anymore. This is because many P2P systems have introduced encryption and obfuscation in order to prevent ISPs from identifying and restricting P2P traffic. DPI systems analyze traffic in real time and use a number of techniques to classify flows according to the communicating application. These techniques include scanning for specific strings in the packet header and payload, behavioral analysis, statistical analysis, etc. With DPI, there is no need to analyze all packets in a flow. Instead, DPI systems only scan for certain patterns in the first few packets of each flow: 1-3 packets for unencrypted protocols and about 3-20 packets for encrypted ones [42]. Indeed, this introduces certain processing delays for all types of traffic and makes DPI-based bandwidth management sensitive to the accuracy of the DPI systems in use. On the other hand, this approach avoids the drawbacks of blocking all P2P traffic. As a consequence, it is widely used by ISPs to prioritize certain applications, while throttling others, especially P2P file sharing [43]. Plus, it allows to address the problem caused by flat-rate pricing and heavy users. According to recent measurements [7] [44], a small fraction of users generates the majority of traffic, causing slow speeds and performance problems for other users. That is, these heavy users pay the same fee but use more resources than the rest of the customers. Bandwidth management is very beneficial in this case as it can be used not only at the application level but at the subscriber level as well to assure fair access to the network and equal bandwidth distribution.

Once packets and flows are classified at the edge of the network, the ISP uses a priority allocation scheme to specify how the traffic should be queued, shaped, and policed. Priority assignment can be static or dynamic. Static priority assignment implies that applications or subscribers belonging to a certain class have a certain level of priority that does not change with time or with the traffic load. Dynamic priority assignment means that the priority of a certain class changes with time or based on the traffic load. Then, specific forwarding treatment on nodes along the path is applied, providing the packet with appropriate guarantees on performance metrics such as delay, delay jitter, and bandwidth. With respect to P2P file sharing, the following "smart" policies can be used to deal with P2P traffic: deprioritizing P2P during congestion periods; throttling upstream P2P traffic (file uploads) while not limiting downstream P2P traffic (file downloads); limiting P2P during certain periods of the day or week (e.g., business vs. evening hours, weekdays vs. weekends); limiting P2P traffic traversing expensive inter-ISP links, etc. [45].

It is interesting to note here the following. In the literature, it is widely accepted that P2P file sharing traffic, crossing domain borders, typically results in heavy expenses associated with a transit fee paid by lower-tier ISPs to transit providers (e.g., see [46] [47] [48] [49]). However, the results of recent studies in this area contradict this common belief that P2P file sharing always incurs intolerable costs to ISPs for its inter-domain traffic. In [50], the author describes how UK ISPs, like Plusnet, are charged for the broadband services that they provide. The calculation shows that transit and peering costs (i.e., the cost of transferring data between the ISP's network and the rest of the Internet) make up less than 14% of the total bandwidth costs. Obviously, the traffic-related costs imposed by P2P file sharing are even less. According to [51], for fixed broadband networks the costs of carrying traffic are a small percentage of the total connectivity revenue and, despite traffic growth, this percentage is expected to stay constant or decline. This suggests that the primary objective of P2P traffic management (at least today) is not to achieve cost savings due to reduced inter-domain traffic charges for ISPs, but to improve the performance of real-time applications and user-perceived QoS.

Blocking and throttling P2P traffic, employed by ISPs to manage bandwidth on their networks, has led to heated debates about "network neutrality" that refers to efforts to keep the Internet open, accessible and neutral to all users, applications, and providers (e.g., see the network neutrality timeline in [52]). The problem is that most ISPs are not very open about their traffic management policies, because they fear losing customers to their competitors. To empower the public with useful information about broadband performance and advance network research, a new initiative, called M-Lab (Measurement Lab) and led by researchers in partnership with companies and other institutions, has been started [53]. The tools running on M-Lab servers can help users test their connections and make informed choices about purchasing and using broadband services. In particular, Glasnost, a network monitoring test, allows to detect whether a local ISP is performing application-specific bandwidth management such as BitTorrent throttling [54]. The results of these tests, covering the period from January 2009 to the present day, are publicly available on [55]. The collected data indicate that DPI-based bandwidth management seems to be declining, with a peak in 2009 at 24% and about 15% these days.

P2P users may try to bypass ISP's bandwidth throttling with Virtual Private Networking (VPN) or SSH tunneling, forcing encryption (many P2P programs have this feature built-in), or getting a dedicated high-speed server, used exclusively for P2P file sharing and known as a "seedbox". But a more network-friendly approach for P2P file sharing without overloading the network and adversely affecting the performance of delay-sensitive applications is to make P2P systems more responsive to congestion conditions in the network. That is, they should be able to quickly adjust their transmission rates in response to information they receive describing the status of the network. It is also a good way to sidestep network neutrality issues and save ISPs money.

To facilitate this approach, the uTorrent Transport Protocol or uTP (also referred to as the micro-Transport Protocol, μTP) has been added to uTorrent, the most popular BitTorrent client, and enabled by default in 2009 [56]. Today, many other BitTorrent clients (Vuze, Transmission, KTorrent, qBittorrent, etc.) also implement uTP. In 2008, a new IETF Working Group, called LEDBAT (Low Extra Delay Background Transport), has been formed aiming to standardize a novel congestion control mechanism that should use the available bandwidth efficiently, maintain low delay, and be no more aggressive than TCP congestion control [57]. Originally, BitTorrent used to run on top of TCP. TCP congestion control is a loss-based mechanism in the sense that it uses either packet losses or excessively delayed packets to trigger congestion-alleviating actions [58]. TCP connections seek to utilize any available capacity, including bottleneck buffers. This results in long queuing delays and delay jitter, causing poor performance of real-time applications that share the same network path. LEDBAT is an experimental congestion control mechanism that attempts to utilize the available bandwidth on an end-to-end path while limiting the consequent increase in queuing delay on that path [59]. A LEDBAT sender uses one-way delay measurements to estimate the amount of queuing on the end-to-end path and controls the transmission rate based on this estimate: once it detects a growing one-way delay, it infers that some router's queue is building up and reacts by reducing the amount of data injected into the network. Thus, it reacts earlier than TCP, which instead has to wait for a packet loss event to detect congestion. This minimizes interference with competing flows and improves the overall performance of the network. Being more network-friendly than TCP, uTP and LEDPAT may be able to make BitTorrent throttling unnecessary. In addition, uTP congestion control layered on top of UDP means the end of the TCP RST packet attacks some ISPs have used to throttle P2P traffic [60].

As noted in [36], for any traffic management policy to be successful, it should be:

--*narrowly tailored*, with bandwidth constraints aimed essentially at times of actual congestion;

--*proportional*, in order to ensure that the policy applied to applications and users is proportional to the level of impact they are having on the network;

--*reasonable*, in order to prevent needless discrimination of lawful content, applications, and user activities;

--*transparent* by making the information about the policy publicly available;

--*auditable* by disclosing information about the rationale behind the policy and the performance improvement, if any.

However, even in this case, there is a risk that bandwidth throttling may alienate P2P users, resulting in cancelled subscriptions, lost revenue, and negative company image. Hence, more P2P-friendly solutions are needed.

III. P2P-friendly Approaches

A. Rationale and Research Work

The problems of P2P file sharing, such as inefficient utilization of network resources and significant amounts of inter-domain traffic, can be addressed by implementing "better-than-random" peer selection and traffic localization. Numerous measurements carried out in P2P overlay networks demonstrate that a large fraction of P2P traffic crosses inter-ISP links. For example, a study of eDonkey file sharing revealed that 99.5% of P2P traffic traversed national or international networks [61]. It also showed that about 40% of this traffic could be localized if appropriate mechanisms were integrated in the P2P protocol. Comcast reported that approximately 34% of BitTorrent traffic was localized in their field trials [62]. In [63], the authors demonstrated that the share of intra-domain traffic in P2P systems can be increased from 10% to 80%. In [64], the authors crawled 214,443 torrents representing 6,113,224 unique peers spread among 9,605 autonomous systems (ASs). They showed that whereas all these peers generated 11.6 petabytes of inter-ISP traffic, a locality policy could have reduced this traffic by up to 40%. However, there is somehow contradictory evidence from [65]. The paper presents a comprehensive study of BitTorrent, using data from a representative set of 500,000 users sampled over a 2-year period, located in 169 countries and 3,150 networks. Surprisingly, the results indicate that BitTorrent traffic exhibits significant geographic and topological locality: 32% of traffic stayed in the country of origin, and 49% of traffic was intra-domain or crossed an inter-ISP link only once. As noted in [65], the observed geographic locality could be explained by certain trends in user content interests (e.g., based on language) and activity patterns (i.e., most users on a continent tend to use BitTorrent at the same time), whereas the locality across networks could be affected by ISP-imposed throttling of inter-domain P2P traffic.

Although traffic localization techniques were originally proposed for P2P file sharing, they can also be beneficial for P2P streaming. However, it is important to note that the effectiveness of these techniques is somewhat limited. For example, language barriers still exist between countries and the media content that is popular in, say, Finland, may not receive much attention elsewhere, naturally leading to country-wide localization. While this observation is valid for a significant part of P2P content, there is enough media content that is popular around the world irrespective of the language of its audio tracks. For instance, this is true for major sporting events such as the FIFA World Cup or Olympic Games. But even when certain content is popular within a single country, it may still lead to large amounts of inter-ISP traffic, especially when small and medium-sized ISPs are concerned.

Traffic localization can significantly impact the QoS perceived by P2P users. On the one hand, it may ensure a higher throughput and lower latency since traffic between peers passes through fewer hops and travels shorter distances. For example, the authors in [66] noticed that the average download time in a locality-aware P2P system may be decreased by up to 20%. Through extensive experiments, the authors in [63] demonstrated consistent improvements in the end-user experience with a decrease in content download times by up to 30%. The increase in the average download rate by up to 80% was also observed in [62]. On the other hand, some traffic localization techniques may lead to clustering of peers [67]. In this case, even inherently robust P2P systems become vulnerable to service interruptions as a result of loss of connectivity between clusters. For instance, the field trials performed by Comcast demonstrated that a

high degree of traffic localization adversely affects the user-perceived QoS in terms of content download times [62]. This could be a result of a variety of factors, including the non-uniform distribution of seeders over swarms and bandwidth among P2P users, where most of the capacity of BitTorrent overlays comes from a small group of broadband peers. Clustering peers on the basis of geographic locality also tends to cluster them by bandwidth, which leads to some improvement in download performance for broadband clusters with many active seeders but causes performance degradation for other clusters [68]. Hence, P2P traffic localization involves a tradeoff between reducing inter-ISP traffic and maintaining the QoS for P2P users.

The most critical part of traffic localization is how to enable locality-awareness in P2P systems. In [69], the authors provided a good overview of the research work in this field. They classified the proposed solutions as end-system mechanisms for topology estimation and ISP-provided topological information. In this paper, we use another classification. We distinguish between 2 types of traffic localization approaches: caching of P2P traffic and biased choice of peers. It should be emphasized that although most of the solutions highlighted below were originally developed for P2P file sharing, they do not heavily depend on the type of transferred content and can be applied to P2P streaming as well [70].

One possible way to localize traffic in P2P systems is to use network caching (also known as in-network storage). In [71], the authors proposed to use the existing Web cache infrastructure already deployed by ISPs to cache P2P traffic. The idea is that a newly connected peer should first determine whether the requested content is available in its ISP cache. This approach tries to benefit from geographical correlation of the content distributed in P2P systems and can be extremely useful for streaming content. However, caching of P2P traffic faces some challenges. In particular, taking into account a large number of P2P systems available today, we need to find a way to differentiate between them. One possible solution is to enumerate all P2P systems and handle them in an ad hoc way by designing and maintaining a separate cache for each system [71]. However, this does not seem to be an easy way due to a number of reasons. First of all, P2P applications may use dynamic ports, randomizing port each start. In this case, it is impossible to identify the type of a P2P system using a single "lightweight" solution. Secondly, taking into account that the number of P2P systems constantly grows, it would be very difficult to deploy as well as update these individual caches. However, recent studies of P2P traffic demonstrate that not all file sharing systems are equally popular [7] [8] [10]. Besides, it is reasonable to expect that this will hold for streaming systems too. Then, as a possible solution to this problem, ISPs could support only those systems that generate the largest shares of traffic [71]. Finally, we note that network caching is equal to introducing a hierarchy into the P2P overlay by implementing caching devices as "superpeers". Such superpeers are always on, providing services to local peers. Indeed, appropriate modifications need to be made to the P2P protocol in order to make regular peers aware of these caching devices.

Biased choice of peers is one of the most promising approaches to enable traffic localization in tracker-based P2P systems [72]. To illustrate the concept, let us consider the file sharing service in BitTorrent. In this system, a new node that wants to join a swarm created for certain content needs to contact the tracker first. The information about the tracker is contained in a special file that is downloaded separately. The tracker maintains information about all peers in this particular overlay network. Upon receiving a request from a new peer, the tracker randomly selects several peers and replies with a list containing their IP addresses. The peer contacts those from the list and initiates P2P file sharing. According to biased choice of peers, the peer of interest contacts only those peers that are geographically close to its location. Alternatively, the tracker itself can be configured to perform this biased peer selection. There are numerous details in this mechanism that affect the performance of the P2P system [73]. For example, the list of peers sent by the tracker should also contain some peers that are not in the neighborhood of the requesting peer. Otherwise, the overlay network will be too clustered. To simplify localization, a new peer can also provide some additional information about its location (e.g., the country of residence).

In [74], the authors carried out a pioneering work comparing different approaches to quantify the impact of peer-assisted file delivery on the end-user experience and resource consumption. They firstly assumed an ideal caching scheme and demonstrated that up to 80% of traffic crossing inter-ISP links can be localized. The authors further noticed that P2P caching solutions are not so easy to implement and proposed to use biased choice of peers, differentiating between peers based on IP address prefixes. They first tried to use a domain matching scheme based on the information extracted from DNS servers. Their results indicate that localization solutions based on domain matching are not as effective as caching of P2P traffic. However, they still reported up to 50% reduction in inter-ISP traffic compared to what a traditional distribution mechanism would produce. Then they proceed with a prefix matching scheme and found that prefixes like /8 result in coarse localization with peers spread out over multiple ISPs. On the other hand, prefixes like /24 result in very small groups of peers. The number of peers in these groups may not be sufficient to provide satisfactory performance of the P2P application. The obtained results indicate that the best grouping scheme is /13. The authors also suggested to use a hierarchical matching scheme, where users are first matched by the /24 prefix, then /16, /14, and finally /13. Although the performance of these algorithms was found to be worse than that of P2P caching, they do not require any assistance from ISPs and can be easily embedded into BitTorrent software.

Successful implementation of many localization algorithms is conditioned on the ability of nodes to measure the distance between each other in the Internet. Basically, there are 2 ways to do that. According to the first approach, they have to use a certain Internet coordinate system (ICS). An ICS provides mapping between IP addresses of hosts in the network and their geographical locations. Nowadays, such systems are mainly used to provide localized advertisements in the Internet. The error of localization using these systems can be as large as 500-700 km [75]. Although this is clearly unacceptable, the best modern techniques can identify a node within 30 km of its actual geographical position [76]. It is important to note that most of ICSs provide somewhat inaccurate results due to non-perfect correlation

between geographical and network distances. Another way to measure the distance between nodes is to allow them to carry out active network measurements themselves. However, it may significantly increase prefetching delays as these measurements are often time consuming. As a measure of distance between 2 nodes, a number of network-related metrics can be used: the number of hops, the round-trip time (RTT), etc. Most approaches proposed to date fall in this category (see [69] and references therein).

In [73], the authors defined the so-called "oracle" service that could be maintained by ISPs. This service provides the topology information to a P2P system such that nearby peers can be identified and chosen by a new node in the overlay. The rationale behind this proposal is that ISPs have complete information about the topology of their own network. As expected, the performance of this solution was shown to be much better compared to other localization techniques. In [77], the authors extended their oracle service by proposing an Internet-wide global coordinate system based on cooperation between ISPs. Elaborating this approach further, they suggested to consider not only the geographical location of a peer but its connection speed as well. Making this information available to new peers that join the overlay will enable a proper choice of peers in the system to be made. Note that an ICS can also be supplemented with bandwidth estimation measurements performed by hosts in the network as was proposed in [78]. It would enable them to make the right choice of peers in terms of higher throughput and low latency.

An interested reader is encouraged to refer to [18] [19] [79] [80] for further details about P2P traffic localization. In the rest of this section, we outline implementation and standardization efforts in P2P caching and biased peer selection.

B. *Caching of P2P Traffic*

Network caching became popular in the mid-1990's as a way to speed up Web transfers over low-speed links. With the massive deployment of broadband technologies, the value of caching of small Web objects decreased. Later, the growing popularity of P2P file sharing and Internet video has attracted much attention to caching again. It was found that P2P and video content responds well to caching, because it has high reuse patterns. In addition, since efficient network caching requires DPI support, caching and DPI are often used in combination.

It should be emphasized that today P2P file sharing traffic is growing in volume but declining as a percentage of global traffic, while Internet video is quickly becoming dominant. Therefore, modern caching systems should be able to deal with both P2P file sharing and Internet video. In practice, many products available on the market today do meet this requirement. Examples include OverCache (Oversi Networks Ltd.), UltraBand (PeerApp Ltd.), CacheFlow (Blue Coat Systems Inc.), iCache (Huawei Technologies Co. Ltd.), etc. Unfortunately, there is a common belief that once an ISP gets into the business of caching of P2P traffic, then it runs a risk of losing its legal immunity from piracy charges. In this context, a widely cited document, Online Copyright Infringement Liability Limitation Act of the U.S. Code, title 17, chapter 5, §512 (a) [81], states that:

"A service provider shall not be liable for monetary relief … for infringement of copyright by reason of the provider's transmitting, routing, or providing connections for material through a system or network controlled or operated by or for the service provider, or by reason of the intermediate and transient storage of that material in the course of such transmitting, routing, or providing connections, if … (4) no copy of the material made by the service provider in the course of such intermediate or transient storage is maintained on the system or network in a manner ordinarily accessible to anyone other than anticipated recipients, and no such copy is maintained on the system or network in a manner ordinarily accessible to such anticipated recipients for a longer period than is reasonably necessary for the transmission, routing, or provision of connections."

As a result, to minimize the risk of legal prosecution, many vendors focus their activities in countries where bandwidth is quite expensive and the return of investments (ROI) is very fast (i.e., mostly in developing countries). For instance, for February 2010, Oversi's customers were mainly located in Eastern and Central Europe (Montenegro, Serbia), Asia (Macao, the Philippines, Singapore, Taiwan, Thailand), and Latin America (Brazil, Chile, Columbia, Mexico, Puerto Rico) [82]. However, in the same act, §512 (b) clearly allows ISPs to use caching in order to speed up content delivery:

"A service provider shall not be liable for monetary relief … for infringement of copyright by reason of the intermediate and temporary storage of material on a system or network controlled or operated by or for the service provider."

The most important requirement here (Condition E in [81]) is that when a user makes some content available online without authorization of the copyright owner of that material, the service provider should respond quickly by removing or disabling access to the material upon receipt of a notification about copyright violation. Besides, ISPs could try to keep their legal immunity from piracy charges by avoiding caching complete files. Instead, they could cache everything except several chunks of a file, making it unusable as is. Anyway, as big players (like Cisco Systems, Huawei Technologies, Juniper Networks, and others) enter the game, we can expect worldwide deployment of caching systems in the Internet.

Recently, the IETF started working toward standardization in this area through the DECADE (Decoupled Application Data Enroute) Working Group [83]. The ultimate goals of the project are as follows:

1) Improve the QoS and reduce resource consumption;
2) Alleviate congestion in the downstream direction;
3) Develop a standard protocol for various P2P and other applications to access in-network storage.

The DECADE architecture consists of 2 planes: control and data [84]. The control plane focuses on application-specific functions, whereas the data plane provides in-network storage and data transport functions. This decoupling of application control and signaling from data access and transport allows to reduce the complexity of in-network storage services. However, there are still a number of open issues that need to be addressed in order to guarantee the efficiency of DECADE and its wide acceptance by both ISPs and users:

--*Late market arrival:* Since DECADE is still in a stage of standardization, while commercial products are already

available, additional efforts should be made to speed up the standardization activities. Close collaboration with companies that are already on the market would be beneficial for the DECADE project.

--*The need for cooperation with the P2P community*: DECADE requires a certain level of cooperation between ISPs and the P2P development community. This is hard to achieve in practice as P2P applications are often developed by loosely organized individuals.

--*Legal issues:* The ability of ISPs to trace and filter out copyrighted content can repel users and the P2P development community. Moreover, in theory, DECADE caches can be used by content owners for spoofing and flooding P2P networks with junk files in order to frustrate P2P users looking for illegally distributed content.

However, the results of the field trials conducted within China Telecom are quite promising [85]. After deploying a DECADE system, consisting of 16 caching devices, each with 1.8-terabyte hard drives, traffic in the network has decreased a lot, resulting in bandwidth savings of up to 55%.

Finally, it is worthwhile to note that some P2P file sharing systems have built-in caching capabilities and follow the main design principle of network caching: "bring popular content closer to consumers". The best known example is Freenet, a distributed P2P file sharing and data storage system, supporting Internet-wide information storage with anonymous information publication and retrieval. In Freenet, once inserted into the network, files migrate closer to where they are most frequently requested (see [86] for details).

C. Biased Choice of Peers

Since P2P applications generate a substantial amount of Internet traffic, locality-awareness of P2P systems has gained much attention of both researchers and practitioners. As a result, a number of techniques have been proposed to address the overlay/underlay topological and routing mismatch. Unfortunately, most of them are still far from maturity for commercial use. To the best of our knowledge, Oversi's NetEnhancer, developed in collaboration with BitTorrent Inc., is the only commercial product available today that allows ISPs to optimize P2P traffic flows across their networks. It supports both commercial and non-commercial P2P applications and can be also integrated with network caching solutions for further traffic management and network performance optimization.

In order to aid in "better-than-random" peer selection while taking into account the underlying network topology, the IETF established the ALTO (Application-Layer Traffic Optimization) Working Group [87]. The ultimate goals of this project are mostly the same as those of the DECADE project, except that traffic localization allows to alleviate congestion both in downlink and in uplink [88]. The ALTO Working Group does not specify or mandate a certain architecture, since there are various architectural options for how the ALTO service could be implemented. However, it does itemize several key components, which should be elaborated and standardized, namely:

1) ALTO server, which provides guidance to applications that have to select one or several hosts from a set of candidates. This guidance should be based on parameters that affect the performance and efficiency of data transmission between the hosts;
2) ALTO protocol, which is used for sending queries and responses between ALTO clients and ALTO servers;
3) Discovery mechanism, which is used by ALTO clients in order to find out where to send ALTO queries.

Similarly to DECADE, ALTO has a number of open issues and security concerns:

--*Information disclosure:* For efficient traffic localization, ISPs should communicate some information about the network topology and resources to external entities. This poses a security threat because such information can reveal some business-related aspects, including the configuration of the network and the load it carries.

--*Clustering of peers and swarm weakening:* Excessive traffic localization via biased choice of peers can cause swarm weakening and thus performance degradation for P2P applications. In order to provide the best possible trade-off between the level of traffic localization and the QoS, additional field studies are needed.

--*Intentional concealment of content and tracking user activities:* Potentially, ALTO servers can misguide ALTO clients on purpose, in order to frustrate P2P users looking for illegally distributed content. In addition, the ability of ISPs to trace and filter out copyrighted content can repel users and ruin ALTO's public image. While this may seem as a perfect solution for copyright owners to retain their rights, it can easily decrease the popularity of ALTO-enabled systems leading to migration of users to other systems and making all efforts obsolete.

Nevertheless, the results of the field trials conducted within China Telecom demonstrate that ALTO is very effective in reducing inter-ISP traffic [85]. For instance, after deploying the ALTO service, the inter-province traffic has been reduced from 75% to 23%. This study also revealed that when P2P traffic is localized, the average download rate decreases. Therefore, ALTO should be used in combination with some network caching technology, like DECADE, or other performance enhancement mechanisms.

IV. Conclusion

P2P file sharing and real-time multimedia (including P2P streaming) are the target applications for traffic management in the nearest future. Hence, for any traffic management solution to be successful, it should be able deal with both P2P file sharing and multimedia streaming.

Overprovisioning and DPI-based bandwidth management are considered the best conventional strategies to deal with P2P traffic. Notice that overprovisioning is nowadays used by most ISPs and is expected to be quite efficient as the volume of global traffic rapidly grows.

Blocking P2P traffic is an approach designed to fully avoid the problems and costs associated with P2P file sharing. However, since P2P file sharing remains fairly popular today, there is a strong probability that commercial customers will not stand for it. Thus, this approach is suitable for campus and corporate networks only.

While bandwidth caps are often the most annoying thing one can get with an Internet connection, traffic quotas fail to address such issues as network congestion and performance degradation of real-time applications caused by bandwidth-hungry P2P traffic.

TABLE I

Approach	Improve performance of real-time applications	Alleviate congestion in the downstream direction	Alleviate congestion in the upstream direction	Reduce transit costs caused by P2P file sharing traffic	P2P-friendly
Overprovisioning	Yes	Yes	Yes	No	Yes
Blocking P2P traffic	Yes	Yes	Yes	Yes	No
Bandwidth caps	No	No	No	Yes	No
Bandwidth management	Yes	Yes	Yes	Yes	No
Caching of P2P traffic	Yes	Yes	No	Yes	Yes
Biased choice of peers	Yes	Yes	Yes	Yes	Implementation-defined

Network caching and biased choice of peers are less expensive, yet more flexible, solutions for P2P traffic management compared to straightforward overprovisioning of network resources. Although the IETF standardization process is still underway, products featuring P2P caching are starting to appear on the market. Field trials and commercial deployments demonstrate good results in terms of reducing inter-ISP traffic and alleviating network congestion. Consequently, ISPs are able to deliver more traffic, more services, and get more revenues over existing infrastructures. The major shortcoming of P2P caching is that it requires additional investment on the caching infrastructure and may cause legal issues in the future. It is worth mentioning here that the success of DECADE and ALTO will ultimately depend on end users, who will evaluate these services based on the observed performance.

In conclusion, we note that the most promising approaches to P2P traffic management and optimization, namely overprovisioning, bandwidth management, network caching and biased choice of peers, are best used in combination. Table 1 summarizes the above discussion about their major benefits and limitations.

ACKNOWLEDGMENT

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013_FP7-REGPOT-2010-1, SP4 Capacities, Coordination and Support Actions) under grant agreement n° 264226 (project title: Space Internetworking Center-SPICE). This paper reflects only the authors' views and the Community is not liable for any use that may be made of the information contained therein.

REFERENCES

[1] *Asymmetric Digital Subscriber Line 2 Transceivers (ADSL2) – Extended Bandwidth ADSL2 (ADSL2plus)*, ITU-T Recommendation G.992.5, 2009.
[2] *Entering the Zettabyte Era*, Cisco Systems Inc., White Paper, June 2011.
[3] *Global IP Traffic Forecast and Methodology, 2006–2011,* Cisco Systems Inc., White Paper, January 2008.
[4] H. Schulze, K. Mochalski, "Internet Study 2007," ipoque GmbH, White Paper, 2007.
[5] H. Schulze, K. Mochalski, "Internet Study 2008/2009," ipoque GmbH, White Paper, 2009.
[6] *Cisco Visual Networking Index: Forecast and Methodology, 2010–2015*, Cisco Systems Inc., White Paper, June 2011.
[7] *Global Internet Phenomena Report: Spring 2011,* Sandvine Inc., Technical Report, 2011.
[8] *Global Internet Phenomena Report: Fall 2011*, Sandvine Inc., Technical Report, 2011.
[9] *Report from the IETF Workshop on Peer-to-Peer (P2P) Infrastructure*, IETF RFC 5594, July 2009.
[10] Internet observatory, http://www.internetobservatory.net
[11] R. Xia, J. Muppala, "A survey of BitTorrent performance," *IEEE Communications Surveys and Tutorials*, vol. 12, no. 2, pp. 140–158, April 2010.
[12] B. Cohen, "Incentives build robustness in BitTorrent," in *Proc. P2PECON*, June 2003.
[13] *An Estimate of Infringing Use of the Internet*, Envisional Ltd., Technical Report, January 2011.
[14] A. Mateus, J. Peha, "Dimensions of P2P and digital piracy in a university campus," in *Proc. TPRC*, September 2008.
[15] *Case Study: RWTH Aachen, Germany. Legal File Sharing with BitTorrent Whitelisting*, ipoque GmbH, White Paper, 2011.
[16] *Meeting the Challenge of Today's Evasive P2P Traffic*, Sandvine Inc., White Paper, September 2004.
[17] *Strategies for Managing the P2P Phenomenon*, Ixia, White Paper, November 2007.
[18] V. Gurbani, V. Hilt, I. Rimac, M. Tomsu, E. Marocco, "A survey of research on the application-layer traffic optimization problem and the need for layer cooperation," *IEEE Communications Magazine*, vol. 47, no. 8, pp. 107–112, August 2009.
[19] J. Wang, C. Wang, J. Yang, C. An, "A study on key strategies in P2P file sharing systems and ISPs' P2P traffic management," *Peer-to-Peer Networking and Applications*, vol. 4, no. 4, pp. 410–419, December 2011.
[20] S. Iyer, S. Bhattacharyya, N. Taft, C. Diot, "An approach to alleviate link overload as observed on an IP backbone," in *Proc. IEEE INFOCOM*, March 2003, pp. 406–416.
[21] *ISP Traffic Management Technologies: The State of the Art*, Heavy Reading, Technical Report, January 2009.
[22] T. Telkamp, "Traffic characteristics and network planning," presented at the NANOG 26 Meeting, Eugene, OR, October 27–29, 2002.
[23] A. Nucci, N. Taft, P. Thiran, H. Zang, C. Diot, "Increasing the link utilization in IP over WDM networks using availability as QoS," *Journal of Photonic Network Communication*, vol. 9, no. 1. pp. 55–75, January 2005.
[24] A. Oliviero, B. Woodward, *Cabling: The Complete Guide to Copper and Fiber-Optic Networking*, Wiley Publishing Inc., 2009, p. 17.
[25] *TERENA Compendium of National Research and Education Networks in Europe, 2011 Edition*, TERENA, Technical Report, 2011.
[26] *TERENA Compendium of National Research and Education Networks in Europe, 2009 Edition*, TERENA, Technical Report, 2010.
[27] M. Mathis, "Reflections on the TCP macroscopic model," *ACM SIGCOMM Computer Communication Review*, vol. 39, no. 1, pp. 47–49, January 2009.
[28] *OECD Communications Outlook 2011*, OECD Publishing, 2011, ch. 4.
[29] *Measuring Broadband America. A Report on Consumer Wireline Broadband Performance in the U.S.*, FCC's Office of Engineering and Technology and Consumer and Governmental Affairs Bureau, 2011.
[30] X. Xiao, T. Telkamp, V. Fineberg, C. Chen, L. Ni, "A practical approach for providing QoS in the Internet backbone," *IEEE Communications Magazine*, vol. 40, no. 12, pp. 56–62, December 2002.
[31] C. Fraleigh, F. Tobagi, C. Diot, "Provisioning IP backbone networks to support latency sensitive traffic," in *Proc. IEEE INFOCOM*, March 2003, pp. 375–385.
[32] M. Perenyi, T. Dang, A. Gefferth, S. Monlnar, "Identification and analysis of peer-to-peer traffic," *Journal of Communications*, vol. 1, no. 7, pp. 36–46, November/December 2006.
[33] P. Rodriguez, S. Tan, C. Gkantsidis, "On the feasibility of commercial, legal P2P content distribution," *ACM SIGCOMM Computer Communication Review*, vol. 36, no. 1, pp. 75–78, January 2006.
[34] *University Rules and Regulations. Application Instructions*, Tampere University of Technology, 2012, [Online]. Available: http://www.tut.fi/idcprod/groups/public_news/@l102/@web/@p/documents/liit/p012430.pdf
[35] *Quota Management Does Not Solve Congestion – Traffic Management Does*, Sandvine Inc., White Paper, August 2010.

[36] *The Evolution of Network Traffic Optimization: Providing Each User Their Fair Share*, Sandvine Inc., White Paper, April 2011.
[37] *Mobile Broadband Capacity Constraints and the Need for Optimization*, Rysavy Research LLC, White Paper, February 2010.
[38] *American Time Use Survey – 2010 Results*, Bureau of Labor Statistics, U.S. Department of Labor, 2011.
[39] M. Mellia, M. Meo, "Measurement of IPTV traffic from an operative network," *European Transactions on Telecommunications*, vol. 21, no. 4, pp. 324–336, June 2010.
[40] *Do Data Caps Punish the Wrong Users? A Bandwidth Usage Reality Check*, Diffraction Analysis, Technical Report, November 2011.
[41] M. Mueller, "DPI technology from the standpoint of Internet governance studies: an introduction," School of Information Studies, Syracuse University, Technical Report, October 2011.
[42] K. Mochalski, H. Schulze, "Deep Packet Inspection: Technology, Applications & Net Neutrality," ipoque GmbH, White Paper, 2009.
[43] Bad ISPs, http://wiki.vuze.com/w/Bad_ISPs
[44] *Cisco Visual Networking Index: Usage*, Cisco Systems Inc., White Paper, October 2010.
[45] *Managing Peer-To-Peer Traffic with Cisco Service Control Technology*, Cisco Systems Inc., White Paper, February 2005.
[46] M. Lin, J. Lui, D.-M. Chiu, "Design and analysis of ISP-friendly file distribution protocols," in *Proc. 46th Annual Allerton Conference on Communication, Control, and Computing*, September 2008, pp. 952–959.
[47] *E-Cohorting – An Opportunity for Telecom Operators Content Delivery Networks*, Tech Mahindra Ltd., White Paper, June 2010.
[48] J. Dai, B. Li, F. Liu, B. Li, H. Jin, "On the efficiency of collaborative caching in ISP-aware P2P networks," in *Proc. IEEE INFOCOM*, April 2011, pp. 1224–1232.
[49] E. Agiatzidou, G. Stamoulis, "Collaboration between ISPs for efficient overlay traffic management," in *Proc. 10th IFIP Networking*, May 2011, pp. 109–120.
[50] D. Tomlinson. (2008, February). How UK ISPs are charged for broadband – the cost of IPStream. [Online]. Available: http://community.plus.net/blog/2008/02/28/how-uk-isps-are-charged-for-broadband-the-cost-of-ipstream/
[51] B. Williamson, D. Black, T. Punton, "The open Internet: a platform for growth," Plum Consulting, Technical Report, October 2011.
[52] J. Fitzsimmons. (2012, April). Network neutrality timeline. [Online]. Available: http://josephfitzsimmons.com/archives/network-neutrality-timeline
[53] M-Lab, http://www.measurementlab.net/
[54] Glasnost, http://broadband.mpi-sws.org/transparency/bttest.php
[55] MLab data, http://dpi.ischool.syr.edu/MLab-Data.html
[56] *uTorrent Transport Protocol*, BEP 29, 2009. [Online]. Available: http://www.bittorrent.org/beps/bep_0029.html
[57] LEDBAT status pages, http://tools.ietf.org/wg/ledbat/
[58] *TCP Congestion Control*, IETF RFC 5681, September 2009.
[59] *Low Extra Delay Background Transport (LEDBAT)*, IETF Internet-Draft, October 2011.
[60] N. Weaver, R. Sommer, V. Paxson, "Detecting forged TCP Reset packets," in *Proc. NDSS*, February 2009.
[61] L. Plissonneau, J. Costeux, P. Brown, "Detailed analysis of eDonkey transfers on ADSL," in *Proc. 2nd EuroNGI Conference on Next Generation Internet Design and Engineering*, April 2006, pp. 256–262.
[62] *Comcast's ISP Experiences in a Proactive Network Provider Participation for P2P (P4P) Technical Trial*, IETF RFC 5632, September 2009.
[63] V. Aggarwal, O. Akonjang, A. Feldmann, "Improving user and ISP experience through ISP-aided P2P locality," in *Proc. IEEE INFOCOM*, April 2008, pp. 1–6.
[64] S. Blond, A. Legout, W. Dabbous, "Pushing BitTorrent locality to the limit," *Computer Networks*, vol. 55, no. 3 pp. 541–557, February 2011.
[65] J. Otto, M. Sanchez, D. Choffnes, F. Bustamante, G. Siganos, "On blind mice and the elephant: understanding the network impact of a large distributed system," in *Proc. ACM SIGCOMM*, August 2011, pp. 110–121.
[66] H. Xie, R. Yang, A. Krishnamurthy, Y.-G. Liu, A. Silberschatz, "P4P: provider portal for applications," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4, pp. 351–362, December 2008.
[67] *Improving Peer Selection in Peer-to-peer Applications: Myths vs. Reality*, IETF Internet-Draft, September 2010.
[68] M. Piatek, H. Madhyastha, J. John, A. Krishnamurthy, T. Anderson, "Pitfalls for ISP-friendly P2P design," in *Proc. 8th ACM HotNets*, October 2009, pp. 1–6.
[69] V. Gurbani, V. Hilt, I. Rimac, M. Tomsu, E. Marocco, "A survey of research on the application-layer traffic optimization problem and the need for layer cooperation," *IEEE Communications Magazine*, vol. 47, no. 8, pp. 107–112, August 2009.
[70] D. Moltchanov, "Service quality in P2P streaming systems," *Computer Science Review*, vol. 5 no. 4, pp. 319–340, November 2011.
[71] G. Shen, Y. Wang, B. Xiong, Y. Zhao, Z. Zhang, "HPTP: relieving the tension between ISPs and P2P," in *Proc. 6th IPTPS*, February 2007.
[72] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, A. Zhang, "Improving traffic locality in BitTorrent via biased neighbor selection," in *Proc. 26th IEEE ICDCS*, July 2007, pp. 66–72.
[73] V. Aggarwal, A. Feldmann, C. Scheideler, "Can ISPs and P2P users cooperate for improved performance?" *ACM SIGCOMM Computer Communication Review*, vol. 37, no. 3, pp. 31–40, July 2007.
[74] T. Karagiannis, P. Rodriguez, K. Papagiannaki, "Should Internet service providers fear peer-assisted content distribution?" in *Proc. 5th ACM SIGCOMM IMC*, October 2005, pp. 6–18.
[75] V. Padmanabhan, L. Subramanian, "An investigation of geographic mapping techniques for internet hosts," *ACM SIGCOMM Computer Communication Review*, vol. 31, no. 4, pp. 27–35, October 2001.
[76] B. Wong, I. Stoyanov, E. Sirer, "Octant: a comprehensive framework for the geolocalization of Internet hosts," in *Proc. 4th USENIX NSDI*, April 2007.
[77] V. Aggarwal, A. Feldmann, R. Karrer, "An Internet coordinate system to enable collaboration between ISPs and P2P systems," in *Proc. 11th ICIN*, October 2007, pp. 33–39.
[78] V. Aggarwal, O. Akonjang, A. Feldmann, "Improving user and ISP experience through ISP-aided P2P locality," in *Proc. 11th IEEE GI*, April 2008.
[79] G. Dan, T. Hossfeld, S. Oechsner, P. Cholda, R. Stankiewicz, I. Papafili, G. Stamoulis, "Interaction patterns between P2P content distribution systems and ISPs," *IEEE Communications Magazine*, vol. 49, no. 5, pp. 222–230, May 2011.
[80] M. Hefeeda, C.-H. Hsu, K. Mokhtarian, "Design and evaluation of a proxy cache for peer-to-peer traffic," *IEEE Transactions on Computers*, vol. 60, no. 7, pp. 964–977, July 2011.
[81] *Copyright Law of the United States of America and Related Laws Contained in Title 17 of the United States Code*, Circular 92. [Online]. Available: http://www.copyright.gov/title17/92chap5.html
[82] G. Peleg. (2010, February). Delivering quality to the Internet video experience. [Online]. Available: http://www.ciscoexpo.ru/club/sites/default/files/seminar_attachments/02_delivering_quality_video.pdf
[83] DECADE status pages, http://tools.ietf.org/wg/decade/
[84] *DECADE Architecture*, IETF Internet-Draft, March 2012.
[85] *ALTO and DECADE Service Trial within China Telecom*, IETF Internet-Draft, March 2012.
[86] I. Clarke, T. Hong, S. Miller, O. Sandberg, B. Wiley, "Protecting freedom of information online with Freenet," *IEEE Internet Computing*, vol. 6, no. 1, pp. 40–49, January/February 2002.
[87] ALTO status pages, http://tools.ietf.org/wg/alto/
[88] *Application-Layer Traffic Optimization (ALTO) Problem Statement*, IETF RFC 5693, October 2009.

Roman Dunaytsev is a Senior Research Associate in the Space Internetworking Center at Democritus University of Thrace, Greece. He received his M.Sc. and Cand.Sc. degrees from Saint-Petersburg State University of Telecommunications, Russia, in 1999 and 2005, correspondingly, and a Ph.D. degree from Tampere University of Technology, Finland, in 2010. His current research interests include space internetworking and delay-tolerant networking, P2P traffic management and optimization.

Dmitri Moltchanov is a Senior Research Scientist in the Department of Communications Engineering at Tampere University of Technology, Finland. He received his M.Sc. and Cand.Sc. degrees from Saint-Petersburg State University of Telecommunications, Russia, in 2000 and 2002, respectively, and a Ph.D. degree from Tampere University of Technology, Finland, in 2006. His research interests include performance evaluation and optimization issues of wired and wireless IP networks, ad hoc and sensor networks and P2P networks. Dmitri Moltchanov serves as TPC member in a number of international conferences. He authored more than 40 publications.

Yevgeni Koucheryavy is a Full Professor in the Department of Communications Engineering at Tampere University of Technology, Finland. He received his M.Sc. degree (1997) from Saint-Petersburg State University of Telecommunications, Russia, and Ph.D. degree (2004) from Tampere University of Technology (TUT), Finland. Yevgeni has been teaching for different target groups, students and professionals in Austria, Brazil, China, Czech Republic, Ireland, Finland, Russia, Sweden and Spain.

He acted as Ph.D. evaluation committee member or examiner in several countries. He has been working in a number of R&D projects within different frameworks, e.g. FP7, and companies. During 2010-2012, from sources external to TUT, Yevgeni managed to attract over 2 million euros as research funding. He is an invited expert in ITU-T and Skolkovo Foundation (Russia) and acts as external reviewer for state funding agencies of several European countries. Yevgeni has authored or co-authored over 100 papers in the field of advanced wired and wireless networking and communications. Yevgeni serves on TPC and Steering Committee of a number of international conferences; he also serves as editor for several international journals. His current research interests include various aspects in heterogeneous wireless communications and systems, network and services performance evaluation, Internet of Things and standardization. Yevgeni is an IEEE Senior Member.

Ove Strandberg is a Senior Specialist at Nokia Siemens Networks, Espoo, Finland. He received the M.Sc. degree in electrical engineering from Helsinki University of Technology, Finland, in 1992. He has been working for the Nokia Research Center for over 22 years, working in areas of transmission technology to IP technology. His current interests are radio access features and, especially, QoS issues in IP networks.

Hannu Flinck is a Senior Specialist at Nokia Siemens Networks, Espoo, Finland. He received his M.Sc. degree (1986) and Lic.Tech. degree (1993) in Computer Science and Communication Systems from Helsinki University of Technology, Finland. He joined Nokia Telecommunications in 1987 where he acted as Group Manager, Section Manager, and Deputy Head of the ATM Switching Department. During 1995-1996 he participated in the TINA consortium acting as Visiting Researcher in Bellcore, NJ, USA. After this he joined the Nokia Research Center where he conducted research on IPv6, IP mobility and its applications. In 2007, he joined Nokia Siemens Networks and continues working on Internet technologies.