

# Interaction Patterns between P2P Content Distribution Systems and ISPs

György Dán<sup>1</sup>, Tobias Hoßfeld<sup>2</sup>, Simon Oechsner<sup>2</sup>, Piotr Cholda<sup>3</sup>, Rafal Stankiewicz<sup>3</sup>,  
Ioanna Papafili<sup>4</sup>, George D. Stamoulis<sup>4</sup>

## Abstract

Peer-to-peer (P2P) content distribution systems are a major source of traffic in the Internet, but the application layer protocols they use are mostly unaware of the underlying network in accordance with the layered structure of the Internet's protocol stack. Nevertheless, the need for improved network efficiency and the business interests of Internet service providers (ISPs) are both strong drivers towards a cross-layer approach in peer-to-peer protocol design, calling for P2P systems that would in some way interact with the ISPs. Recent research shows that the interaction, which can rely on information provided by both parties, can be mutually beneficial. In this paper first we give an overview of the kinds of information that could potentially be exchanged between the P2P systems and the ISPs, and discuss their usefulness and the ease of obtaining and exchanging them. We also present a classification of the possible approaches for interaction based on the level of involvement of the ISPs and the P2P systems, and we discuss the potential strengths and the weaknesses of these approaches.

## I. INTRODUCTION

Peer-to-peer (P2P) content distribution systems have become a dominant source of traffic in the Internet. Their primary use has been off-line content distribution, i.e., file sharing (e.g., BitTorrent, Gnutella, eDonkey), but they are increasingly used to stream live and on-demand video as well (e.g., SopCast, PPLive, Zattoo). The peers in a P2P content distribution system form an overlay network and exchange data between each other, such that the data are not delivered to the individual peers directly from the content source, but via a number of other peers. The indirect delivery of the data between peers in an overlay is called overlay routing. For overlay routing to work, the peers have to contribute with their resources, their storage capacity and their upload bandwidth. The popularity of P2P systems has increased the demand for broadband Internet access, still Internet service providers (ISPs) have started to consider P2P traffic “unwanted”, for a number of reasons.

- The popularity of P2P systems has led to increased network traffic and increased resource utilization. Higher traffic volumes require investments in the infrastructure and increase the traffic related costs.
- P2P systems usually create overlay networks unaware of the underlying physical network topology, called the underlay network. The network topology agnostic design can be motivated by the layered architecture of the TCP/IP protocol stack. But given the large amounts of data delivered via P2P systems, the network agnostic design leads to inefficient network resource

<sup>1</sup>School of Electrical Engineering, KTH, Royal Institute of Technology, Stockholm, Sweden, gyuri@ee.kth.se

<sup>2</sup>Department of Distributed Systems, Institute of Computer Science, University of Würzburg, Germany, {hossfeld|oechsner}@informatik.uni-wuerzburg.de

<sup>3</sup>Department of Telecommunications, AGH University of Science and Technology, Krakow, Poland, {piotr.cholda|rstankie}@agh.edu.pl

<sup>4</sup>Department of Informatics, Athens University of Economics and Business, Athens, Greece {iopapafi|gstamoul}@aueb.gr

usage. A one-hop overlay connection can span several intercontinental links and can traverse a number of different autonomous systems (ASs) and ISPs. Inter-AS links are often expensive, depending on the agreement between the involved ASs. Apart from being expensive, inefficient overlay routes can also contribute to network congestion.

- The properties of P2P traffic are usually transient and difficult to predict. The most popular P2P content distribution systems divide the contents into relatively small pieces, so that peers can download (and upload) different parts of the content from (and to) different peers simultaneously. The set of peers with which a peer exchanges data can change relatively fast. In a large P2P system, this may lead to traffic fluctuation on short timescales, which results in inefficient traffic management and the breakdown of network dimensioning assumptions.

A number of ISPs attempted to decrease their costs due to P2P traffic by restricting it in their networks. Some ISPs deployed traffic shaping devices to limit the sending rates of popular P2P applications. Others throttled the bandwidth of their heaviest users irrespective of the types of applications they used. Again, others injected packets to reset the TCP connections used to transfer data between the peers. These techniques rely on identifying the P2P traffic in the network, either via the ports it uses or via deep packet inspection. As an effect, to avoid identification, P2P systems started to use randomly selected ports and traffic encryption.

Such an “arms race” is, however, counterproductive for both ISPs and P2P system providers. ISPs spend resources on equipment that interferes with P2P traffic, but can lose customers if they degrade the performance of popular applications in their networks. P2P system providers spend their resources on improving the obfuscation techniques instead of on improving the systems’ performance. The alternative of the “arms race” is cooperation, which could eventually lead to benefits for both ISPs and P2P content distribution systems. The issue of how ISPs and P2P systems could cooperate, and how the cooperation would influence their respective costs and performance, has been the subject of research in the past few years, and is the topic of this paper.

The rest of the paper is organized as follows. In Section II we give a classification of the approaches that have been proposed in the literature. We describe a general framework for the interaction between overlays and underlays in Section III. We consider some legal and techno-economical challenges that the different forms of interaction might face in Section IV, and conclude the paper in Section V.

## II. CLASSIFICATION OF ISP-P2P INTERACTION

A number of approaches were proposed in the literature to decrease the tension between ISPs and P2P systems. In the following we provide a classification of the approaches among two dimensions: the involvement of the P2P system (the overlay) and the involvement of the ISPs (the underlay). Based on the level of involvement we divide the approaches in four classes.

- **Class 1: ISP’s indirect influence on the overlay.** ISPs apply traffic engineering methods, and treat aggregates of certain types of traffic preferentially in order to optimize the Quality of Service (QoS) and the resource usage in their networks. Traffic engineering may influence an overlay favorably as a by-effect of the optimization of the ISP’s operations, if the overlay’s traffic is chosen to be prioritized.
- **Class 2: ISP’s direct influence on the overlay.** The ISP performs operations that directly influence the overlay in a way that the peers are not aware of the involvement of the ISP. The goal of the ISPs is to improve the efficiency of the P2P system in terms of network resource usage (e.g., inter-AS traffic), such that their own costs decrease. The aim is a win—not-lose situation, though the P2P system performance might be slightly negatively affected eventually.

- **Class 3: Peer-to-peer system's unilateral involvement.** Contrary to the previous class, in this class it is the P2P protocol that is modified in order to optimize the overlay with respect to the underlay, but without the involvement of the ISPs. The aim is a win-not-lose situation, but the gains of the ISPs are hard to predict.
- **Class 4: Mutual direct influence.** Requires close cooperation between the P2P systems and the ISPs with the common goal of improving the performance of both. The ISP operates an infrastructure that provides information to the P2P systems, which have to be modified to make use of the information. The aim is a win-win situation.

Figure 1 shows the four classes, and the location of the proposed approaches within the classes. In the following, we give a detailed description of the existing approaches with respect to these classes.

#### A. ISP's indirect influence on the overlay

Solutions belonging to Class 1 are only applicable to IP networks that use traffic engineering to serve a limited number of QoS classes. Traffic engineering consists of the special treatment of traffic aggregates on the basis of their quality requirements, with the goal of improving the QoS and optimizing the resource usage. Aggregation needs packet inspection, as packets are classified based on their source and/or destination IP addresses, the ports as well as the protocol type. For instance, QoS can be supported by DiffServ, Multi-Topology Routing (MTR), or MultiProtocol Label Switching (MPLS).

An ISP managing an IP/DiffServ (IETF RFCs 2475, 4594, 5127) domain defines service classes and related treatment of traffic according to its policy. Packets are treated according to their classification and their marking at the domain's border. The ISP does not provide any information to the applications, neither do the applications explicitly inform the ISP about their preferred service classes.

While DiffServ operates on the level of forwarding, Multi-Topology Routing (IETF RFC 4915, 5120) performs routing on the basis of the traffic type assigned to a packet. This approach roughly consists of establishing different routing tables for various types of traffic. The goal is to treat preferentially some types of traffic, without affecting the quality of others, and additionally, to optimize the capacity usage and to perform load balancing within an ISP domain. For practical reasons, the number of classes and routing tables should not be too high. Thus, traffic classification should be rather coarse-grained, and the level of the ISP's intervention is not as high as in the case of DiffServ.

Finally, MPLS (IETF RFC 2702, 3031) assigns packet flows to different classes, and makes differentiation possible both for forwarding and for routing. Each flow has its own tunnel that uses a specific path for which additional parameters can be specified, including bandwidth and failure recovery method.

Traffic engineering does not affect the locality of P2P traffic, and consequently it does not decrease the inter-AS traffic, but it may reduce the operator's intra-domain link loads.

#### B. ISP's direct influence on the overlay

The approaches in Class 2 are based on an ISP managed proxy node that influences the operation of the P2P system, such that the peers are not aware of the involvement of the ISP and hence the P2P protocols do not have to be modified. The proxy may operate in the control plane or in the data plane of the P2P system, or even in both of them.

In the control plane, the proxy can influence the peer selection process, e.g., it can redirect a peer's requests for content to local peers that already own the content or it can modify the packets that carry information about possible neighbors. The implementation of such a redirector proxy can lead to technical issues, as it might interfere with the security features of the P2P protocol. Nevertheless, if it can be implemented, the approach could lead to a significant decrease of the traffic costs as shown in early studies on the properties of P2P traffic [1], [2].

In the data plane, the proxy can act as a transparent data cache. A transparent cache intercepts P2P traffic using deep packet inspection, and serves the local peers' requests for data, if already stored in the cache. Alternatively, the proxy can be an ISP-owned peer [3] with high upload and storage capacity. Studies based on measurements of P2P content popularity indicate that, given the practical limits on cache sizes, the hit rate and the efficiency of P2P caches is limited due to the heavy tail of the P2P content popularity distribution [4]. Depending on the ISP topology, the cache efficiency can, however, be increased by up to an order of magnitude if the caches of the neighboring ISPs are allowed to cooperate with each other [5].

ISP-managed proxies face, however, a number of technical and non-technical challenges. First, P2P traffic has to be identified in the network, which often requires deep packet inspection because of dynamic port allocation used by many P2P applications. Some P2P applications encrypt their traffic, which apart from making the identification difficult, makes ISP managed proxies infeasible. Second, the proxy has to contain P2P application specific parts, and hence it has to be updated every time a new P2P protocol appears. Third, the installation and maintenance of the proxies involves costs for the ISP, which have to be compensated by the savings in terms of traffic charges. Fourth, due to the widespread use of P2P applications for distributing copyrighted content, caching might lead to legal issues. For example, ISPs can be held legally responsible for caching copyrighted material if the caching does not comply to the Safe Harbor Provisions (§512) of the Digital Millennium Copyright Act in the U.S.A, and to §5 of the Directive 2001/29/EC in the E.U.

### *C. Peer-to-peer system's unilateral involvement*

The approaches in Class 3 do not rely on information provided by the ISP, but they are based on information obtained by the peers via some form of measurements. The measurements are used by the peers to infer their proximity according to some metric. In general, the proximity information can be used in two ways. It can be used to influence the neighbor selection process, and hence the topology of the overlay network. Alternatively, the proximity information can be used to select the peers in a given set of neighbors with which data is exchanged, and hence affecting the overlay routing but not the overlay's topology itself.

A number of metrics have been considered to measure proximity. The simplest metrics are the round-trip time (RTT) between the peers, and the number of IP hops as reported by *traceroute*. These two measures were, e.g., considered in [6] to cluster peers in BitTorrent into a hierarchical structure. Peers primarily exchange data within the same cluster, and since they are nearby with respect to the metric, the traffic over long distances decreases. The RTT and the number of hops are easy to measure, but using them for optimization does not necessarily decrease the inter-ISP traffic. For example, the round-trip time or the number of hops between two peers in a geographically spread ISP might be higher than that between two geographically nearby peers in different ISPs.

The use of the IP addresses as the basis for a distance metric is more involved. A possible way to infer locality based on the IP addresses is to use servers that map IP addresses to AS numbers. The servers could be part of the P2P system's infrastructure, or could be provided by a third party. Such solutions were considered to bias the neighbor selection [7] and to change the download policies

of the peers [8] in BitTorrent, and were shown to reduce the amount of inter-domain traffic, while the download performance was almost unaffected. Another way to infer locality based on the IP addresses is to use the DNS resolutions of CDN servers [9]: two peers are assumed to be close to each other if the name servers they use resolve CDN host names to the same IP addresses. The exact preferences of the ISPs are however hard to infer from AS numbers or DNS resolutions without support from the ISPs. For example, there are several ISPs that manage multiple ASs. They are often top-tier ISPs with complex network topologies and policies. Such ISPs might prefer traffic to stay within their networks, but might not be interested in limiting the traffic to a single AS.

Proximity can also be measured based on network coordinates, e.g., Vivaldi. Proximity aware schemes based on network coordinates have been considered for Distributed Hash Tables (DHTs) [10], [11]. Proximity routing selects from a set of possible next overlay hops the one that is shortest according to the metric used, i.e., it affects the forwarding of data. Proximity Neighbor Selection and Geographic layout create the routing tables of the peers such that they maximize proximity.

There are two major issues related to the approaches in Class 3. First, the peer selection based on these metrics does not necessarily lead to optimal choices from the underlay's perspective, and it is not clear what is the metric that would lead to the best results for the underlay. Second, it is not clear how the modified peer selection or overlay routing mechanisms will affect the performance of the systems. Most research efforts focus on answering these two questions.

#### *D. Mutual direct influence*

The approaches in Class 4 are based on collaboration between the P2P systems and the ISPs. Conceptually, the ISP deploys an entity in its network through which it provides information to the P2P system about, e.g., the network topology and the network state. The peers can use the information obtained from the entity in the control plane (e.g., to optimize the peer selection, the overlay routing) and/or in the data plane (e.g., to adapt the transmission rates between the peers). Since the information is provided by the ISP, it is more accurate than the information that the peers would obtain via reverse engineering as in Class 3. At the same time, the ISP can influence the optimization via the information it provides. The most important questions concerning collaboration are how the peers can discover such an entity, and what information the entity should provide to the peers.

The simplest implementation of an entity could provide proximity information about the peers participating in an overlay. The oracle node considered in [12] provides such a service, by ranking the potential neighbors of every peer on the basis of physical-topology proximity metrics (e.g., distance in terms of AS hops). Proximity does not have to be static, it can incorporate the actual network state, like in the case of the SmoothIT Information Service (SIS) [13]. The information provided by the entity can also include the cost of certain paths in the network as proposed for the P4P portal [14]. The costs reported to the peers are calculated such that the peers, if they make their overlay routing decisions based on them, will optimize the underlay's performance according to the criteria chosen by the ISP.

Apart from providing information that helps proximity aware overlay construction and routing, the ISPs can inform the peers about the existence of caches. If peers prioritize the cache over external peers then the transit traffic of the ISP decreases, similar to the case of transparent caches and ISP-owned peers described in Class 2. If peers prioritize the cache over local peers, then the cache can be used to decrease the congestion on the last mile up-links as well. It is, however, not clear how such priorities should be implemented and how such caches would affect the overall application performance.

In order to support a multitude of P2P systems, the discovery of the ISP provided entities and the communication with them has to rely on a standardized application layer protocol, or the entities have to be decomposed into an application specific part, and an application independent part, e.g., as described in [14]. There is ongoing work in the Application Layer Traffic Optimization (ALTO) working group of the IETF on localization services and service discovery protocols for, among others, P2P systems. Service discovery includes the discovery of ISP provided information sources, but it also concerns the discovery of other ISP supplied resources, such as caches.

This class of approaches provides the greatest potential for optimization but the success of these approaches depends on the willingness of the ISPs, the P2P system providers and the P2P users to cooperate, and on the standardization of scalable, extensible protocols for information exchange.

### III. FRAMEWORK FOR INTERACTION BETWEEN OVERLAYS AND UNDERLAYS

Among the approaches discussed in Section II the ones belonging to Class 4 have the greatest potential to improve network efficiency: both network operators and P2P applications could benefit from exchanging information with each other. This section gives an overview of the information that could be useful and could be made available by the network providers to the P2P applications and vice versa. Tables I and II give a summary of the section. In these tables we characterize the availability of the information or solution as *high*, *moderate* or *low*. This general characterization is based mainly on technical issues. Additionally, we use the term *sensitive* in some cases to indicate that the information might be confidential and not be intentionally provided from one party to the other. Figure 2 shows the types of information as a function of their availability and their sensitivity.

#### A. Information provided by the ISPs

We start with the discussion of the information about the underlay that could be beneficial for the P2P systems. A summary of the discussion is shown in Table I.

*Network topology, capabilities and state:* ISPs usually treat the information about their detailed network topology as confidential, both for commercial reasons and because the detailed information could potentially be used to attack the network infrastructure. Consequently, the information provided has to be coarse-grained but still beneficial to the P2P systems. The solutions proposed in [12], [14] were designed according to these criteria. Based on the information about network capabilities (e.g., capacity, delay, or access technology, like dial-up, DSL, FTTH, cable, wireless), a P2P application could distinguish between peers with low and high connection rates and adapt the traffic load accordingly. More sophisticated optimizations can be made by a P2P application if low-delay or high-bandwidth routes were advertised. Low delay routes can be utilized, for example, by interactive video games, VoIP and other applications requiring low latency. High bandwidth routes can be used for bulk transfers. Overlay self-organization mechanisms can also benefit from information about bandwidth, estimated delay, congested links, and the presence of bottlenecks. ISPs could provide distance metrics between peers, e.g., information about the distance to the edge of the AS (according to an IGP metric), the number of AS hops (on the basis of BGP), and whether a peer is inside or outside the AS [12]. These metrics facilitate the locality and proximity awareness of P2P applications.

*ISP policies:* P2P applications could benefit from the knowledge of how different protocols are handled by the ISP, what traffic policing rules are used in the network, and what traffic patterns are preferred. Such information, if provided by the ISP, could be used by the P2P applications to

adjust their protocols (e.g., the overlay routing) in order to meet underlay preferences. The ISPs and the P2P applications could agree on using a selected range of ports and transport protocols. This way ISPs could recognize particular applications more easily, while the applications' traffic would receive a better treatment. If the ISP supports several QoS classes, the information on traffic marking rules should be available to P2P applications. In general, if ISP policies were public, P2P users would be able to make a more conscious decision on using a particular P2P application, possibly the one that "cooperates" most with the ISP.

*Resources provided by the ISP:* An ISP can decide to provide additional resources to a P2P application, in the form of caches or ISP supplied high-capacity peers. Using these resources the ISP may influence the performance of the application and the behavior of regular peers. If the ISP provides information about the installed resources, P2P applications can take advantage of them more efficiently. The benefits are reduced traffic load and resource usage, as well as a decrease of the ISP's costs [15].

### B. Information provided by the P2P applications

In the following we discuss the information that could be provided by the P2P systems to the ISPs. A summary is shown in Table II.

*Peer location and activity:* By peer location we refer to the identity of the ISP to which a peer belongs. An ISP could make use of the information about the location of the peers participating in the distribution of a specific content, e.g., this could help nearby peers to find each other. Without support from the P2P systems it is a complex task for an ISP to find out which peers are members of a given overlay. Even though this information is sometimes available (e.g., at the trackers in BitTorrent), it can still be difficult for the ISP to obtain and to update it. Nevertheless, once the IP addresses of the participating peers are known, the peers' locations in the underlay topology can be easily obtained.

Information about the activity of the peers (e.g., upload and download bitrates) and the number of peers that belong to a particular overlay can be used to estimate the traffic demands. For some P2P systems this information is available at web sites in the form of tracker statistics (or can be obtained by directly contacting the trackers). The information about the peers' activity requires frequent updates, and the information can be potentially sensitive, since it tells about the users' content interest profiles and it might be used to disturb the operation of the overlay.

*Content popularity distribution:* Obtaining the content popularity distribution (e.g., the number of seeds and leechers for individual BitTorrent swarms in an ISP) is not straightforward, and might require the exchange of large amounts of information. Nevertheless, knowledge of the content popularity can be used for resource optimization, e.g., in caches. It is not known how well one can predict the future evolution of the content popularities, but, for example, flash-crowds (sudden increase of the peer population) are in general hard to anticipate. Information about the content popularities, and in particular, the interest profile of individual subscribers may be sensitive, as it could also be used by the underlay provider for, e.g., direct marketing.

*Application requirements:* Knowledge of a P2P application's requirements can help ISPs to provide useful QoS differentiation, and to provide incentives to the specific application by optimizing the performance parameters relevant to it. The requirements can be known to the ISP in two ways. First, via direct interaction between the applications and the ISPs: the applications use some protocol to communicate their requirements (e.g., packet loss, end-to-end delay). Second,

via employing application/traffic identification (e.g., using deep packet inspection), in which case the ISP has to have an up-to-date database of application requirements. The approach based on direct interaction is likely to be superior in terms of scalability and extensibility, if the protocol is designed appropriately.

#### IV. CHALLENGES

There are a multitude of challenges related to the interaction between ISPs and P2P systems. Some of them have a strictly technical character (e.g., method of information exchange, timescale for measurement actions or dealing with non-cooperative players), others have a legal character (e.g., network neutrality) or are more related to techno-economic issues (e.g., selection of proper incentive mechanisms). In this section, we discuss briefly the major challenges for the different classes of interaction. Figure 3 presents the major challenges for each class of interaction. Classes are placed in a coordinate system based on the level of intervention of the ISP to the overlay and vice versa. We distinguish between three different types of challenges: technical, economic and legal; they are indicated by white, light gray and dark gray boxes respectively. In the following we briefly describe the major challenges.

- *Non-cooperative behavior*: One of the two actors, the ISP or the P2P system, may try to exploit the information provided by the other actor, without providing any information in return.
- *Incentive mechanisms*: In order to establish and maintain cooperation, the ISP has to provide incentives to the P2P system, and vice versa.
- *Information exchange*: The information has to follow agreed-upon semantics, and its exchange requires agreed-upon, scalable protocols. The protocols should provide authentication, integrity, and should be robust to denial of service attacks.
- *Illegal content*: Legal issues may arise when an ISP aids the illegal distribution of copyrighted content, e.g., via certain forms of caching.
- *Cost savings*: The ISP has to deploy an infrastructure, which increases its capital and operational expenses. The cost reduction achieved by proximity aware operation must be higher than the expenses.
- *Application requirements*: The P2P application requirements have to become known to the ISP so that the QoS parameters can be adjusted accordingly, but in a manner that is general enough to be applicable for future applications.
- *Network neutrality*: Each packet in a network should be routed and forwarded impartially, not taking into account its content.
- *Network topology challenges*: Customers located nearby in the same physical network infrastructure may be provided Internet access by different ISPs. Proximity can be difficult to discover even in the case of cooperation.

Among these major challenges the last two are related to the socio-economic and legal environment, which the ISPs and the P2P application providers cannot influence much, hence we discuss them in more detail.

##### A. Network Neutrality

The idea behind network neutrality is to avoid potential discrimination of the traffic generated by some users or applications. So far, there has been a lively discussion on it in the US and in Europe, and legislation is likely to follow. There are two main variants of net neutrality. Below we discuss how three classes, Classes 1, 2 and 4, of ISP-P2P interaction are affected under these variants.



Class 3, since the solutions do not involve the ISPs, is applicable independent of the variant of net neutrality.

- **No intervention:** Operators cannot favor or punish any type of traffic. Approaches belonging to Class 1 cannot be used, as they are built on the preference for some types of flows. Approaches belonging to Class 4 can not be used if they might punish some overlays, while favoring others. Even though approaches in Class 2 have no direct influence on the quality, they might lead to an indirect improvement of the quality for some applications, which is not acceptable.
- **QoS for aggregates:** Operators can introduce some priorities, but only on the basis of the traffic type, e.g., streaming or file-sharing. However, they should not distinguish between traffic within a given traffic class, e.g., on the basis of the source and destination IP addresses. Class 1 would be an ideal approach here, as it directly conforms to the mandate of dealing with traffic aggregates. Class 2 could, however, be banned, for it is possible to prove that the operator does not give the same treatment to aggregates not related to the applications it supports. The use of approaches belonging to Class 4 is also questionable.

If neither of the above definitions of net neutrality applies then operators can employ in their domains any policy for the treatment of the packets. In this case Classes 1, 2 and 4 can be used without any legal issues.

### B. Network Topology Challenges

Historically, the Internet was considered to be a network of transit and stub ASs. While this classification still applies to some extent, the relationships between the ASs have become more complex in recent years.

On the one hand, the roles of the network owner and the service provider are separating. The reason for the separation can be economical, as in the case of *city networks*, or legal, as in the case of *bitstream access*. A *city network* is a geographically confined broadband network infrastructure that provides IP connectivity to its subscribers, but no Internet access. Internet access is provided by a number of competing ISPs via the network infrastructure. In the case of *bitstream access*, it is the wireline incumbent that operates the last mile connection, e.g., ADSL, but it makes the connection available to its competitors at some point of presence (PoP) either at the DSLAM (Digital Subscriber Line Access Multiplexer), or further. Hence, the last-mile is owned by the incumbent operator, but Internet service is provided by a competitor ISP. Both in city networks and in the case of bitstream access two customers can be far away at the network layer despite their geographic and network-wise proximity in the link layer, as data between them have to travel via their respective ISPs. Consequently, approaches in Class 3 can fail to discover proximity, but approaches in Class 1 and 4 can also be affected, if the ISPs do not know each other's topology and IP address assignment policy.

On the other hand, the ISPs that cover the same, extended geographic area often maintain *multiple peering points and bilateral peering agreements* with each other. A proximity aware peer selection scheme might choose a peer within the same ISP, even though a geographically nearby peer in the peering ISP would be optimal in terms of efficient network utilization. The existence of multiple peering points can affect the approaches in Class 3, as these do not have accurate information about the network topology.

## V. CONCLUSION

There are many alternatives to the struggle between ISPs and P2P content distribution systems. The alternatives differ in the level of involvement of the ISPs and the P2P systems, and of course,

in the achievable benefits. In this paper we have given a classification of the main approaches, and discussed their advantages and the challenges they face. We concluded that both actors hold information that is useful for the efficient operation of the other. By exchanging information both ISPs and P2P systems could optimize their operations: ISPs could improve the utilization of the network resources, while P2P systems could achieve better system performance. We gave an overview of the information that could be exchanged, how it may be used and how it could influence both the P2P applications and the underlying networks. We have discussed the challenges faced by the ISPs and the P2P systems in the implementation of the different approaches.

We conclude that mutual cooperation is the most promising approach for both ISPs and P2P content distribution systems, but its implementation faces many challenges. Much research and related standardization work will be required in order for it to prevail in the future.

#### ACKNOWLEDGEMENTS

This work was partially funded by the EU FP7 Network of Excellence Euro-NF. The authors thank Andrzej Jajszczyk, Jerzy Domzal, Robert Wojcik, and Przemyslaw Pawelczak for their comments.

#### REFERENCES

- [1] N. Leibowitz, A. Bergman, R. Ben-Shaul, and A. Shavit, "Are file swapping networks cacheable? Characterizing P2P traffic," in *Proc. of the 7th Int. WWW Caching Workshop (WCW-7)*, Aug. 2002.
- [2] K. P. Gummadi, R. J. Dunn, S. Saroiu, S. D. Gribble, H. M. Levy, and J. Zahorjan, "Measurement, modeling, and analysis of a peer-to-peer file-sharing workload," in *Proc. of the 19th ACM Symposium on Operating Systems Principles (SOSP)*, 2003, pp. 314–329.
- [3] I. Papafili, S. Soursos, and G. D. Stamoulis, "Improvement of BitTorrent performance and inter-domain traffic by inserting ISP-owned peers," in *Proc. of 6th International Workshop on Internet Charging and QoS Technologies (ICQT)*, May 2009.
- [4] M. Hefeeda and O. Saleh, "Traffic modeling and proportional partial caching for peer-to-peer systems," *IEEE/ACM Trans. on Netw.*, vol. 16, no. 6, pp. 1447–1460, 2008.
- [5] G. Dán, "Cache-to-cache: Could ISPs cooperate to decrease peer-to-peer content distribution costs?" *IEEE Trans. on Parallel and Distr. Syst.*, to appear.
- [6] J. Yu and M. Li, "CBT: A proximity-aware peer clustering system in large-scale BitTorrent-like peer-to-peer networks," *Computer Communications*, vol. 31, no. 3, pp. 591–602, 2008.
- [7] R. Bindal, P. Cao, W. Chan, J. Medval, G. Suwala, T. Bates, and A. Zhang, "Improving traffic locality in BitTorrent via biased neighbor selection," in *Proc. of 26th IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2006, pp. 66–76.
- [8] B. Liu, Y. Cui, Y. Lu, and Y. Xue, "Locality-awareness in BitTorrent-like P2P applications," *IEEE Trans. on Multimedia*, vol. 11, no. 3, pp. 316–371, 2009.
- [9] D. Choffnes and F. Bustamante, "Taming the torrent," in *Proc. of ACM SIGCOMM*, Aug. 2008.
- [10] M. Castro, P. Druschel, Y. Hu, and A. Rowstron, "Topology-aware routing in structured peer-to-peer overlay networks, Tech. Rep. MSR-TR-2002-82," 2002. [Online]. Available: [citeseer.ist.psu.edu/castro02topologyaware.html](http://citeseer.ist.psu.edu/castro02topologyaware.html)
- [11] R. Gummadi, S. Gribble, S. Ratnasamy, S. Shenker, and I. Stoica, "The impact of DHT routing geometry on resilience and proximity," in *Proc. of ACM SIGCOMM*, 2003.
- [12] V. Aggarwal, A. Feldmann, and C. Scheideler, "Can ISPs and P2P systems co-operate for improved performance?" *ACM SIGCOMM Computer Communications Review (CCR)*, vol. 37, no. 3, pp. 29–40, July 2007.
- [13] T. Hossfeld, D. Hausheer, F. Hecht, F. Lehrieder, S. Oechsner, I. Papafili, P. Racz, S. Soursos, D. Staehle, G. D. Stamoulis, P. Tran-Gia, and B. Stiller, "An economic traffic management approach to enable the triplewin for users, ISPs, and overlay providers," in *Proc. of 2nd Future of the Internet Conference*, May 2009.
- [14] H. Xie, Y. R. Yang, Y. Liu, A. Krishnamurthy, and A. Silberschatz, "P4P: Provider portal for applications," in *Proc. of ACM SIGCOMM*, Aug. 2008.
- [15] M. Merrit, D. Pasko, and L. Popkin, "Network-friendly peer-to-peer services," in *Proceedings of the IETF Workshop on P2P Infrastructures*, May 2008.

## BIOGRAPHIES

GYÖRGY DÁN (gyuri@ee.kth.se) received the M.Sc. degree in computer engineering from the Budapest University of Technology and Economics, Hungary in 1999 and the M.Sc. degree in business administration from the Corvinus University of Budapest, Hungary in 2003. He received his Ph.D. in Telecommunications in 2006 from KTH, Royal Institute of Technology, Stockholm, Sweden, where he currently works as an assistant professor. He was visiting researcher at the Swedish Institute of Computer Science in 2008. His research interests include the design and analysis of distributed and peer-to-peer systems.

TOBIAS HOßFELD (hossfeld@informatik.uni-wuerzburg.de) studied computer science and mathematics at the University of Würzburg, Germany and obtained his PhD in 2009. Currently, he is heading the research group "Future Internet Applications and Overlays" at the Chair of Distributed Systems in Würzburg. His main research interests cover virtualization, social networks, overlays and P2P systems, as well as investigations on Quality of Experience (QoE) for Internet applications like Skype or YouTube.

SIMON OECHSNER (oechsner@informatik.uni-wuerzburg.de) is a PhD student at the Chair of Distributed Systems, University of Würzburg. He received his Master's Degree in 2004 from the University of Würzburg. His research interests are the performance evaluation of search and content distribution overlays, traffic management in P2P networks, and QoE control in P2P video streaming.

PIOTR CHOLDA (piotr.cholda@agh.edu.pl) received the Ph.D. degree in telecommunications from AGH University of Science and Technology, Kraków, Poland, in 2006, which he joined the same year. His research interests focus on design of resilient multilayer networks as well as reliability modeling, including overlay networking. He is the coauthor of six refereed journal papers and two tutorials on resilient networks. Dr. Cholda is the recipient of the Communications QoS, Reliability and Performance Modeling Symposium Best Paper Award from ICC'06. Currently, he is involved in two EU projects: Euro-NF, and SmoothIT.

RAFAL STANKIEWICZ (rstankie@agh.edu.pl) received the M.Sc. and Ph.D degrees in telecommunications in 1999 and 2007, respectively, from AGH University of Science and Technology, Kraków, Poland, where he is currently employed at the Department of Telecommunication. His research interests focus on advanced methods of ensuring QoS and QoE, performance modeling, and traffic management concepts for peer-to-peer networking. He is an author of several conference and journal research papers and co-author of one book. Currently, he is involved in the EU projects SmoothIT and Euro-NF.

IOANNA PAPAFILI (iopapafi@aueb.gr) is a PhD student in Computer Science in Athens University of Economics and Business (AUEB). She has a diploma in Computer, Telecommunications and Network Engineering from University of Thessaly (UTH), Volos (2006), and a MSc in Computer Science from AUEB (2008). Her research interests include overlay traffic engineering, peer-to-peer, and socio-economics. Ms. Papafili is also member of the Network Economics and Services Laboratory of AUEB and of the Technical Chamber of Greece.

GEORGE D. STAMOULIS (gstamoul@aueb.gr) is a Professor in the Informatics Department of Athens University of Economics and Business. He received the Electrical Engineering Diploma (1987) from National Technical University of Athens, Greece, with highest honors, and the MS (1988) and PhD (1991) degrees in Electrical Engineering and Computer Science from Massachusetts Institute of Technology (MIT), USA. His main research interests are in network economics and in the application of incentives mechanisms in networks and electronic environments.

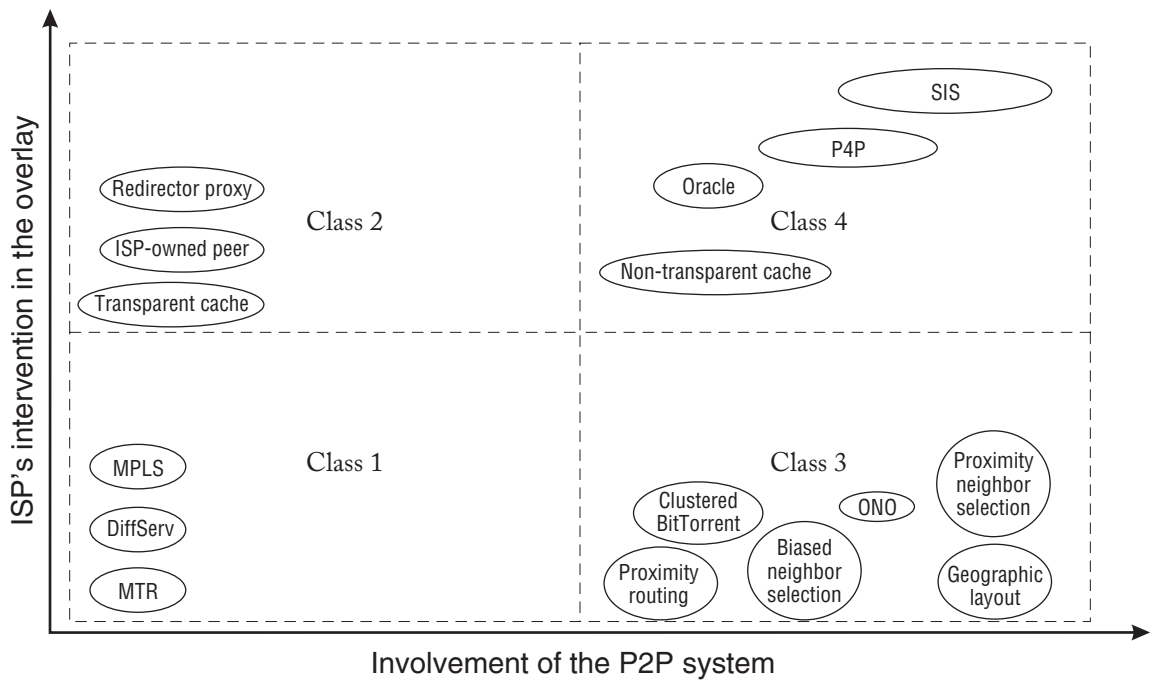


Fig. 1. Approaches with respect to the involvement of the ISP and the P2P system.




Fig. 2. Exchangeable information with respect to availability and sensitivity

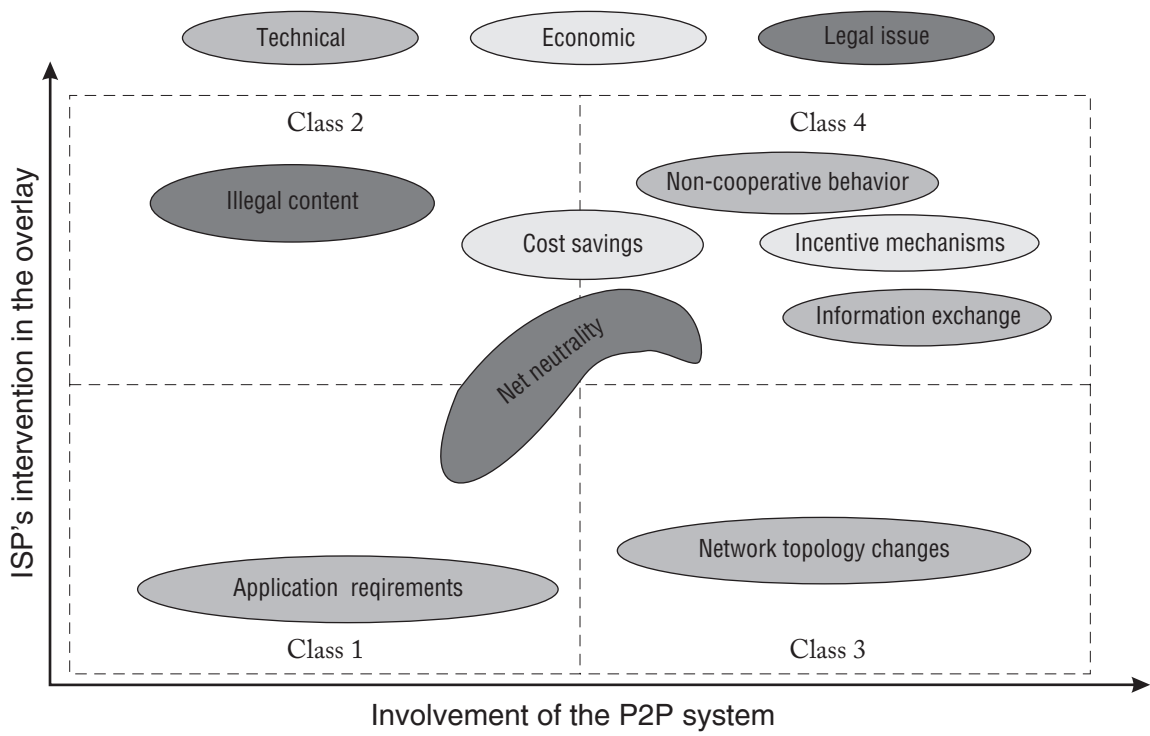


Fig. 3. Overview of challenges according to the classification of the approaches

TABLE I  
USEFUL INFORMATION FOR THE P2P SYSTEM (BASED ON [15])

Information	Use	Availability
<b>Network topology and state</b>		
Network capabilities	Distinguishing peers with low and high capacity connection	Moderate, sensitive
Optimal low-latency routes	Optimization for applications requiring low latency	Moderate, sensitive
Optimal high-throughput routes	Optimization for bulk data transfers	Moderate, sensitive
Network performance	Optimization of P2P decisions	Sensitive
Optimal low cost routes	P2P decision with respect to ISP preference	High
Network distance between peers	P2P decision based on more steady metric	Moderate/high
<b>ISP policy</b>		
Routing policies	Selection of types of packets, protocols	Moderate, sensitive
Preferences on port ranges/protocol	Selection of ports and protocols	Moderate, sensitive
Preferences on traffic profile, overlay routing, applications' activities etc.	Optimization of P2P application's behavior for better cooperation	High
Network policies	Adaptation of P2P protocols to better cooperate with ISPs	Low, sensitive
Suggestions on traffic marking	Receipt of appropriate QoS	Moderate/High
<b>Resource provided by ISP</b>		
Caching servers	More advantageous use of caching servers by P2P	High

TABLE II  
USEFUL INFORMATION FOR THE UNDERLAY

Information	Use	Availability
<b>Peer location and population</b>		
Location of peers in the overlay	Optimization of traffic management (e.g., dimensioning of links, change of routing paths) and improved caching	Moderate, possibly sensitive
Number of uploaders and downloaders per overlay	Traffic estimation in incoming and outgoing directions (on inter-domain level), in combination with locality information	High
<b>Popularity</b>		
Popularity of content	Caching policies	Low, possibly sensitive
<b>Requirements</b>		
Application requirements	QoS differentiation	Moderate