A New Approach for Managing Traffic of Overlay Applications of the SmoothIT Project

Juan Fernández-Palacios[1], Maria Angeles Callejo[1], Hasan[2], Tobias Hoßfeld[3], Dirk Staehle[3], Zoran Despotovic[4], Wolfgang Kellerer[4], Konstantin Pussep[5], Ioanna Papafili[6], George D. Stamoulis[6], Burkhard Stiller[2]

[1] Telefonica Investigacion y Desarrollo, Madrid, Spain, [jpfpg¦macr]@tid.es
[2] University of Zurich, Department of Informatics IFI, [hasan¦stiller]@ifi.uzh.ch
[3] University of Würzburg, Würzburg, Germany, [hossfeld¦dstaehle]@informatik.uni-wuerzburg.de
[4] DoCoMo Euro-Labs, Munich, Germany, [despotovic¦kellerer]@docomolab-euro.com
[5] Technische Universitaet Darmstadt, Darmstadt, Germany, pussep@kom.tu-darmstadt.de
[6] Department of Informatics, Athens University of Economics and Business, Athens Greece, [papafili06¦gstamoul]@aueb.gr

Abstract. This paper outlines a new technical approach for controlling and managing network traffic of overlay applications based on incentives that lead users to configure their traffic (for their own good) in such a way that the overall situation in the network is improved. Additionally, right incentives for the network provider exist to reach a win-win situation.

Keywords: Overlay applications, ETM (Economic Traffic Management), Incentives, Locality, Peer, P2P, QoE (Quality of Experience), Underlay

1 Introduction

Overlay networks emerged in the last decade and generate now a large portion of the Internet traffic [Internet07]. Even though they were initially associated exclusively with copyright infringements and, thus, provoked considerable public odium, their good properties supersede and the overlay networking technology is here to stay. On one hand, overlay networks can provide specific applications without being supported by current underlay network, *e.g.*, multicast. On the other hand, the client-server model does not scale in different respects, and many content distributors find it more efficient to deliver their content using overlay networks. For example, free Linux distributions (*e.g.*, Red Hat's Fedora or Debian) are distributed via BitTorrent.

Furthermore, the Internet traffic stemming from overlay-based applications, *e.g.*, Peer-to-Peer (P2P) applications, increases rapidly with the increase of available bandwidth of end-nodes. Therefore, a slow paradigm shift from centrally offered services to services offered by end-nodes is in progress. For today's Telecommunication Service Providers (telcos) and Internet Service Providers (ISP) the key issue arising is how to control and manage network traffic stemming from such overlay-based applications. As the structure of overlays determines traffic flows in ISP networks, it is highly efficient for an ISP to influence the overlay configuration based on information on their structure. Thus, the management of overlays aims (a) to maximize the benefit for multiple operators/ISPs involved and (b) to increase the capability to withstand faults and balancing the network load.

Traditional traffic engineering methodologies are insufficient to satisfy simultaneously the provider's goal to maximize network usage and the user's goal to maximize his utility, since most of current overlay networks are deployed oblivious of the underlying physical network [Awduche99]. *E.g.*, traffic between two nodes in the Telefonica network in Madrid can go via Tokyo, incurring large and unnecessary costs. This is why network operators show a strong interest in influencing overlay network operations and shaping the overlay traffic as to comply with their needs.

Within that context, SmoothIT proposes a new traffic management mechanism termed Economic Traffic Management (ETM), which provides for an incentive-compatibility in those interactions foreseen between overlay applications and the underlying ISP networks in order to gain the following measurable impacts:

1. Cost saving for ISPs: lower operation costs, due to ETM-based traffic engineering, lower interconnection costs, since traffic can be kept inside an ISP's domain, and lower capacity extension cost, since capacity requirements can be forecasted with much higher accuracy.
2. Lower prices for end-users, due to competitive pricing by the ISP, which are enabled by new ETM mechanisms.
3. Better Quality-of-Service (QoS) for overlay-based applications across ISP domains, due to the usage of ETM-based traffic engineering. This leads to an improved media consumption experience for end users.

2 Technical and Economic Characteristics of Overlay Applications

This section provides an overview on a wide spectrum of overlay applications showing their different technical and economic characteristics. This is done by selecting three particular examples of overlay applications, which are highly popular among users and often used in practice: eDonkey, Skype, and Joost. In the context of the SmoothIT (***S***imple Economic ***M***anagement Approaches ***o***f ***O***verlay ***T***raffic in ***H***eterogeneous ***I***nternet ***T***opologies) project, these applications are relevant, since an optimization potential exists, with which applications are able to react to appropriate incentives for adapting themselves and the resulting overlay topology to their own and operator's needs. The selection of these three examples is done to show their differences in (a) offered service types (file-sharing, Voice-over-IP (VoIP), and Video-on-Demand (VoD)) and (b) their technical and economic characteristics. These differences lead to key effects on the overlay traffic management (cf. Section 4) and to motivate the new mechanisms for overlay traffic management.

The ***eDonkey file-sharing application*** belongs to the class of hybrid P2P architectures and comprises out of two applications: the open-source eDonkey client (or peer) and the closed-source eDonkey server. An eDonkey server operates as an index server for file locations and distributes addresses of other servers to clients. The eDonkey client stores, shares, and downloads files. A main feature of eDonkey is the multi-source download, *i.e.*, peers can issue several download requests for the same file to multiple providing clients in parallel, and providing clients can serve the requesting peer simultaneously [Hoss04].

eDonkey utilizes TCP (Transmission Control Protocol) mainly for signaling data transfer. Additional UDP (User Datagram Traffic) traffic between peers and index servers is, however, negligible at roughly 1% of the overall eDonkey traffic. In particular, the ratio of signaling traffic volume to data traffic volume is about 1:15. Although, for file-sharing, there are no strict delay and bandwidth requirements, an eDonkey client shows exhaustive bandwidth consumption via TCP and tries to fully utilize its download bandwidth. An eDonkey user is mainly interested to download files in a short time. Additionally, he might want to reduce its own upload traffic. The reduction of upload traffic is of interest, since the upload bandwidth of a user is typically much smaller compared to the available download bandwidth. While running eDonkey, users do not want that other applications, like web browsing, are negatively influenced or experience a worse quality. An ISP transporting eDonkey traffic aims at the reduction of costs for transit traffic to other ISPs without degrading (and, if possible, improving) its customers' Quality-of-Experience (QoE).

For satisfying the user's and the ISP's interest, the **optimization potential** is based on the adaptation of the overlay topology to the underlying topology. The download from close peers in the same domain reduces the download times as well as inter-domain traffic. As a necessary instrument, locality promotion has to be integrated by the ISP. SmoothIT aims at providing incentive mechanisms to enable and make this interaction between the ISP and the overlay exactly beneficial.

The ***Skype VoIP application*** defines a proprietary approach, which uses encrypted traffic and anti-debugging techniques to avoid reverse-engineering. Skype is mainly used for phone calls between Skype users, but also between Skype and PSTN users (called Skype-in and Skype-out calls, respectively). Host nodes communicate directly with each other or via super nodes to overcome Network Address Translators and firewalls. The bandwidth observed during a Skype voice call for an end-user is about 30 kbps and goes up to 120 kbps. An overview of traffic characteristics and Skype's traffic management mechanisms implemented is provided in [Hoss08a].

Voice communication has strict traffic requirements. According to the ITU-T G.114 recommendation, the end-to-end delay should be below 200 ms. If it is above 400 ms, many users are dissatisfied with the service. [Hoss07] shows that an exponential relationship between packet loss and the QoE obtained was found. This exponential interdependency clearly outlines the sensitivity of QoE to small QoS degradations. To satisfy its customers, an ISP has to minimize disturbances like loss or jitter for VoIP.

Skype is quite popular among users, since its calls are free among its users and it often offers cheap prices for national and international Skype-in and Skype-out calls. Basic requirements to use the service are quite low in terms of CPU power and network access bandwidth. Currently, Skype traffic is transported according to a "best effort" approach. However, voice services have strict QoS requirements (*i.e.*, packet loss, delay and jitter). Therefore, network operators' could improve the Skype QoE by using real-time transport connections, *e.g.*, prioritizing VoIP packets.

Optimization potential: Although Skype is quite popular among users and heavily used by a large community, the impact on ISPs is rather low. The emerging traffic amount is small compared to file-sharing or VoD applications. In addition, it is difficult to adapt the overlay topology, since mainly direct communication paths between voice partners are established. In case of relayed traffic, the selection of appropriate super nodes is possible, however, this has most likely only a minor impact on transit costs of an ISP. Thus, locality promotion is not a key issue for Skype.

Nevertheless, an ISP – beside the end-user of course – is interested in a good QoE for popular applications of its customers. The interest behind that is the selling argument to provide a good Internet access for a variety of applications. In the context of Skype, this means that the ISP could offer QoS differentiation for such real-time communication to guarantee a certain QoE to the end-user, which could be charged for by adding an extra fee to the regular flat rate.

The ***Joost VoD application*** aims at providing licensed, high-quality content using the H.264 video codec at a resolution comparable to Phase Alternating Line (PAL) (720x576). Joost is proprietary, closed-source, and provides a standalone application for watching videos and exploring video channels. For delivering their videos, Joost uses a server-assisted P2P technology, also referred to as peer-assisted server system, integrating the end-user into the data dissemination process by storing and distributing already downloaded video content. Experiments show that the local cache has a size in the order of 1 GB. The Joost network is separated into two overlays according to their functionality. These are peer management and media distribution. Joost relies on UDP for its video streaming and on TCP for its peer management.

A measurement study showed that the used downlink bandwidth of a peer is approximately at 500 kbps to 700 kbps [Hoss08b]. Since the Joost user is participating in a P2P network and provides already downloaded content to other peers in the network, an upload bandwidth of about 70 kbps to 100 kbps was observed. Around 600 UDP and 200 TCP connections were established mainly to peers in Europe and North America. About 25%-50% of the video data was identified to be downloaded from Joost servers. In contrast to eDonkey, bandwidth consumption of a Joost client is not aggressive, *i.e.*, even if more bandwidth at a peer is available, it is not utilized by Joost. Further information on Joost and its traffic characteristics are available in [Fu07, Hoss08b].

Traffic requirements of a VoD application include of course the reliable transport in-time of data packets. Depending on the local jitter buffering, delay variations can be overcome. Nevertheless, the investigation of the impact of network disturbance, either packet loss or transmission delays, on the QoE for the Joost VoD system is a topic of future work. The main incentive for users to use the Joost application is the fact that (a) the offered service is for free and (b) the application provides high quality video content. A Joost client itself has not many options to influence the behavior of the Joost application. Of course, well-known ports to the Joost content servers can be blocked. However, no explicit parameter settings are available in the user interfaces to control the Joost behavior. Operators can control the amount of traffic of P2P VoD applications and indirectly profit from them by using charging models based on the amount of the transferred data volume. This charging model is often used in wireless (*e.g.*, 3G or 2G) Internet connections.

Optimization potential: Currently, a Joost client connects to other peers and servers world-wide. The location of the peers contacted appears arbitrary and independent of the underlying topology of the ISPs involved. Optimization potentials discussed above for eDonkey and Skype could now be combined for the Joost application. As for eDonkey, the overlay topology of Joost can be adapted to the underlying topology to increase efficiency of the data dissemination and to reduce ISP's costs. Thus, incentives and mechanisms to challenge the information asymmetry between the overlay and the underlying network have to be provided such that the ISP performs location promotion and the overlay reacts adequately.

QoS differentiation provided by the ISP could be beneficial as well to fulfill traffic requirements of streaming application and to meet users' demands, *i.e.*, the QoE. Nevertheless, this depends strongly on the buffering and data exchange mechanisms implemented, whether QoS differentiation can be utilized effectively.

3 Current Situation: Key Aspects of Overlay Traffic Management

As seen, a huge portion of Internet traffic emerges from overlay applications. For example, according to a recent study about the Internet traffic distribution in Germany [Internet07], around 74% of total Internet traffic in Germany is generated by P2P applications. Overlay applications are consuming a large amount of bandwidth in operators' access, metro, and core networks. Both the amount and the distribution of overlay traffic are strongly impacting total network costs (CAPEX: Capital Expenditures and OPEX: Operational Expenditures). In particular, if an ISP customer is exchanging P2P content with a customer of a different ISP, such traffic is consuming resources in the whole network: access, aggregation, core, and interconnection (cf. Figure 1).

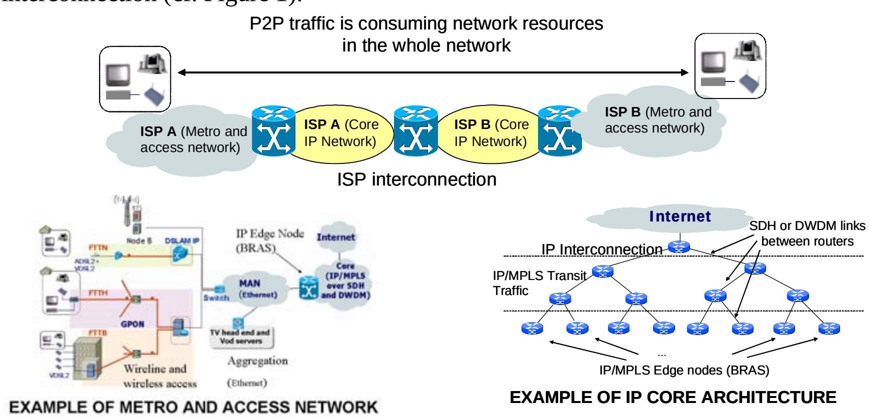

Fig. 1. Network Architecture for Overlay Traffic Transport

The end-to-end path used by an inter-carrier overlay flow begins in the access network. End-users of overlay applications (peers) can use either wire-line (*e.g.*, fiber, cable, xDSL) or wireless (*e.g.*, WIMAX, UMTS, GPRS) Internet connections. Traffic from multiple end-users is aggregated in the operator's access edge node (*e.g.*, DSLAM, GPON OLT, UMTS Node B), so that aggregated traffic flows from multiple access nodes are transported over Layer 2 networks (*e.g.*, Ethernet) toward the IP edge router (BRAS: Broadband Remote Access Server), which inspects users' packets in order to check their destination address. Afterwards, IP packets are aggregated in MPLS (Multi-protocol Label Switch) flows and sent to the destination IP node which can be either internal (*i.e.* another internal BRAS) or external (*i.e.* located in another ISP network). In the second case, traffic is sent to the IP interconnection point. ISPs' networks are interconnected as autonomous routing domains. Global routing and network reachability among these Autonomous Systems (AS) is managed by BGP (Border Gateway Protocol). This network interconnection (cf. Figure 2) is particularly important for the cost of overlay networking.

Fig. 2. ISP Interconnections (BGP)

The ISP interconnection consists of the advertisement between ISPs of routes to their customer's IP addresses. Whether such exchange of reachability information and routing of traffic could be done for free or not depends on the ISPs' business relationship. Relationships between ISPs are generally described by one of the following categories:

- Peering: Two networks exchange traffic between each other's customers freely (as long as not important traffic asymmetries are detected).
- Transit: An ISP pays to another ISP for the traffic exchange.

In the IP transit model the purchaser has to pay the difference between outbound and inbound traffic. Pricing is typically offered on a Mbps/month basis and requires the purchaser to commit to a minimum volume of bandwidth. For example a common charging model for IP transit is based on the 95th percentile method. According to this method, the difference between the average inbound and outbound traffic is measured every 5 minutes and recorded in a log file, so that at the end of the month, the top 5% of data is discarded and the following measurement becomes the billable utilization for that month.

Therefore, total networking costs depend on inter-carrier overlay traffic causing ISP transit interconnection fees. As illustrated in Fig. 3, peers connected to different BRAS consume not only access and aggregation traffic, but especially transit traffic, for which the originating ISPs have to pay for. Whereas peers connected to the same BRAS only consume access and aggregation resources of the local ISP, which can be handled without extra cost, if the network is well dimensioned.

Existing traffic management mechanisms try to throttle P2P overlay traffic to avoid high costs. However, this also impacts the profit from such traffic (if paid per packet) or the user utilizes a flat rate subscription. In order to keep interconnection fees as a main source of additional cost to a minimum, the promotion of overlay traffic locality is the key issue from a network operator's perspective. In turn, also the overlay provider will benefit from a better performance and shorter end-to-end delays.

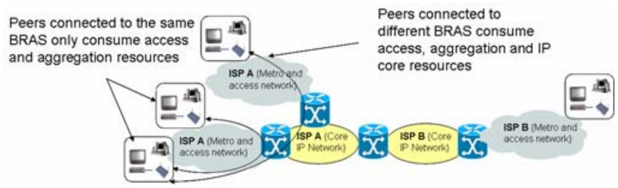

Fig. 3. ISP Interconnections (BRAS)

A second key issue with respect to overlay traffic management is QoS differentiation. Today, Internet traffic (*e.g.*, http or overlay) is transported according to a "best effort" approach only. However, some overlay applications such as IP-TV (Internet Protocol-based Television), VoD, VoIP, videoconference, or gaming present strict requirements in terms of delay and packet loss. It is difficult for the overlay provider to meet these requirements without influence on the network. Therefore, the quality of experience perceived by end-users of real-time or streaming overlay applications, such as Skype, PPlive, or Joost, could be improved significantly by introducing application-aware transport services being able to provide low packet loss rates and jitter performance. Of course, developing and implementing such application-aware traffic management mechanisms would require new network investments. However, operators could also profit indirectly by improving the QoE of overlay applications in increasing the broadband customers' fidelity and, thus, sell new broadband connectivity services to end-users.

4 New Approaches for Managing Overlay Traffic

SmoothIT proposes a new approach to manage network traffic through incentive mechanisms. This section outlines (a) the design space as a basis for developing ETM mechanisms and (b) a selected set of scenarios to show the applicability of ETM.

Design Space: A peer in an overlay performs a set of functions to achieve a user's goal of running an application. Normally, obtaining and consuming distributed shared resources, *e.g.*, multimedia contents, is a typical user's goal. 4 phases are required generally by a peer to reach this goal. In the resource finding phase, potential sources of the resource being sought are located, normally with the help of an indexing service, *e.g.*, based on Distributed Hash Tables (DHT). This is followed by the "best sources" finding phase, which aims at selecting a set of sources, from where the resource is to be obtained. The resource transfer phase enables the transfer using a specific transfer protocol. Note that during the resource transfer some applications continuously look for "best sources", *e.g.*, to maintain application-layer multicast trees for real-time video streaming. Thus, the "best sources" finding phase might overlap with the "resource transfer" phase. Finally, within the resource handling phase the resource is consumed in an application-specific way. In case of streaming resources, the resource transfer and handling phases overlap.

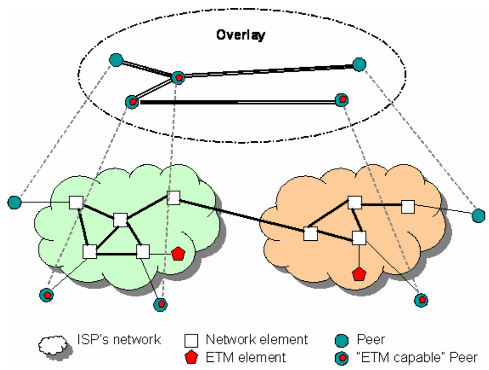

Fig. 4. Introducing ETM Elements into ISPs' Networks and Peers

Mechanisms required for these phases are well established except for the "best sources" finding phase. It is the task of ETM to aid peers in an overlay to find "best sources of resources." This is important, since it will determine benefits of both the user and ISPs. Thus, the term "best" must be viewed from both points of view and it considers cost-effectiveness next to technical performance. Without any knowledge of the underlying network topology, an application decision to choose certain sources of resources may lead to undesirable network resource consumption for the ISP.

In order to enable a peer to acquire this knowledge, SmoothIT introduces new elements in ISP networks as well as ETM-capable peers in the overlay (cf. Figure 4). These elements are responsible for evaluating utilities of each party involved, given a set of alternative sources to deliver resources requested by the application. Note that no assumption is made on possible (horizontal) relations among ETM-capable peers and among ISP's ETM elements. However, a (vertical) communication between an ETM-capable peer and an ISP's ETM element is assumed to achieve a win-win optimization solution. Table 1 shows various design parameters to be investigated in SmoothIT in order to achieve those goals stated above.

Table 1. Design Space

Dimension	Possibility
Pricing scheme for network resources	1) congestion-based, 2) demand-based, 3) performance-based, 4) flat-rate
Incentives	1) performance improvement, 2) price reduction, 3) reliability, 4) security
Information useful for overlay, which can be provided by the underlay	1) locality of nodes, 2) performance of a transport service, 3) path reliability to a node
Architecture	1) ETM elements constitute or organize themselves as a "cross ISP infrastructure" providing a special service to evaluate utility similar to indexing service provided by a DHT. 2) Each ISP operates its own ETM elements independently.
Parameter for utility evaluation	*User*: Resource quality (impact on QoE), resource charges, network performance (impact on QoE), network usage charges,

	reliability, security.
	ISP: Interconnection bandwidth usage (costs), utilization of ISP’s core networks, charge of network usage billable to customers (depending on pricing scheme or amount of usage)

Based on suitable scenarios applicable models involving these parameters are outlined in the following.

Scenario 1 — Operator’s P2P Application: An ISP can provide a P2P application on its own by offering initial resources and letting users contribute others in a P2P manner. For example, if a P2P Content Distribution Network (CDN) is assumed that uses a P2P protocol to distribute its content, an ISP could provide a search server (cf. Figure 5), where users could search for torrent files. Such a torrent would point to a tracker that runs as a service on the ISP’s own servers, monitoring and steering the download of content. The content itself would be distributed by peers, while the ISP could offer additional caches to improve the download performance.

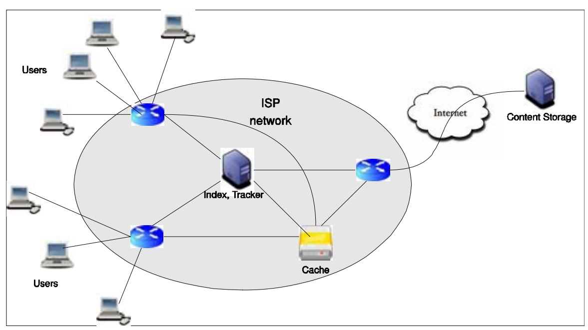

Fig. 5 Example Hybrid Application

The performance of such an application can be improved by locality-awareness, efficient bandwidth utilization and congestion avoidance. This is due to the fact that the ISP is able to monitor such information inside its network and, therefore, the ISP's tracker can steer the content distribution process. The user will receive the content with good quality and low infrastructure costs [Brecht03], thereby allowing competitive prices for the content [Bindal06]. However, note that an ISP will offer such a service only, if benefits resulting from an efficient network utilization and smaller inter-domain traffic are higher than additional infrastructure costs. The main sources of these costs are index and tracker servers as well as caches. The swapped content can be either user-generated or offered by commercial content providers, *e.g.*, movie studios. In the latter case the seed for content distribution can reside either on the ISP’s or content providers’ servers (cf. Figure 5). The first time the content is requested it is transmitted to the ISP’s network and cached by the ISP. This way, the content has to cross inter-domain borders only once. Legal issues may arise regarding what content is offered, however, since the provider offers the search service and controls the tracker, it has full control of what items can be downloaded by users.

Scenario 2 — Locality promotion by technical incentives: As explained, overlay networks, such as P2P networks, are decentralized and self-organized but they do not consider the underlying network topology when selecting peers. This independence may result in inefficient use of resources of the underlying network and in a high end-to-end delay for both overlay and underlay traffic. Therefore, one of the approaches advocated by SmoothIT is that P2P employs topology awareness to provide for good QoS at the virtual overlay network layer, as well as to impose an optimized and balanced load at the underlying network layer. This approach is in line with solutions proposed by other researches in certain cases outlined below.

For the selection of a "good" server or a close-by peer, in terms of latency, a distributed, scalable binning scheme that requires a small number of landmark machines spread across the Internet, is proposed in [Ratna02]. These landmark machines are well-known and binning is performed based on node's relative distance from them.

Because of landmarking not being self-organizing, a new model is proposed in [Le05]. The model uses Geographical Longest Prefix Matching (Geo-LPM) and RTT (Round Trip-time) to perform node clustering. A cluster is a group of nodes that share a common prefix and are close to each other. Geo-LPM with an appropriate RTT threshold ensures that a node entering the network will find other overlay nodes that belong to the same physical network. However, clusters that share a common prefix may not be physically close to each other. To this end, the above solution is combined with Geographical Partitioning (Geo-Partitioning), which creates an appropriate structure of partitions.

According to [Karag05], peer-assisted content distribution is a cost-effective and bandwidth-intensive solution for ISPs that does not require vast investments in terms of infrastructure as CDNs do. The utilization of new entrant-peers' resources increases the system's capacity at the same rate as the demand increases. Also, P2P systems such as BitTorrent are inherently more robust compared to traditional CDN systems, *e.g.*, the tit-for-tat policy makes BitTorrent ideal in terms of content availability. Although peer-assisted architectures, such as BitTorrent, provide significant benefits to end-users and content providers, they roughly double total traffic and peak load on ISP's links. As a result, ISPs often throttle BitTorrent traffic. As argued by means of experimental results by [Karag05], an ISP-friendly locality-aware peer-assisted content distribution protocol could alleviate ISPs' induced costs and prevent them from throttling P2P traffic.

Finally, Biased Neighbor Selection, an approach to enhance BitTorrent traffic locality, is examined in [Bindal06]. According to this, a peer chooses the majority of his neighbors from peers within the same ISP. Only a few neighbors are selected outside the ISP due to the lack of the content within the domain. One approach to implement this locality-awareness scheme is the modification of both tracker and client, so that the tracker selects internal peers and only a few ones outside to respond to the client's query.

The key difference of the SmoothIT approach to these topology-aware solutions is that in general these solutions are beneficial for the ISP without offering a clear advantage to the overlay and the end-nodes, *i.e.*, without offering a clear incentive to be adopted at the overlay level. In particular, SmoothIT investigates locality-aware solutions that offer incentives to the overlay while being beneficial at the same time for the ISP.

Scenario 3 — Agreements between service and network providers: In this scenario the application is controlled by a single, well-known service provider, which forms a proprietary overlay network in the backbone, *e.g.*, YouTube, Akamai, or AmazonS3. Network operators could offer economic and/or technical incentives to service providers, if they update their applications or CDN architectures in order to increase the percentage of intra-domain traffic. For example (cf. Figure 6), an ISP could offer low cost or even free network connections to the service provider Point-of-Presences (PoP) located in its own network so that user-server traffic does not consume expensive IP transit resources.

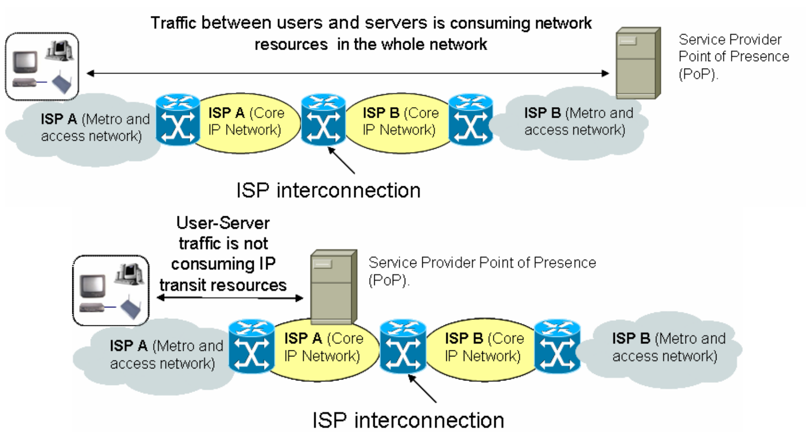

Fig. 6. Impact of server distribution in network resources consumption

5 Summary and Preliminary Conclusions

This paper summarizes the challenge of effectively ***and*** efficiently managing the traffic of overlay applications in the Internet. On one hand, overlay traffic, such as from P2P content sharing applications, constitutes a huge portion of the Internet traffic and is one of the most important contributions to total network costs (CAPEX and OPEX). On the other hand, P2P applications are also the key driver in the commercialization of broadband access. Therefore, overlay traffic management has emerged as an important task for network operators. Up to now, the overlay is operated independently of the underlying network infrastructure. This may cause threads for network operators as, for example, traffic is routed expensively across domains, but also for the overlay, since the service quality can hardly be influenced. In the approach presented, covering economic overlay traffic management aspects, a proper interrelation is addressed between the network operator and the overlay itself, where both sides benefit from. In particular incentives are identified that lead network operators and overlays to a win-win situation in respect with high-performance and cost saving traffic control. The undertaken analysis of popular P2P overlay applications (*e.g.*, eDonkey, Skype, and Joost) reveals incentives as well as optimization potential for a win-win situation. Network provider incentives are mainly motivated by avoiding high cost for ISP interconnection fees, (but also by satisfying the QoE of their users for an increased broadband customers fidelity), while the end-user is interested in a better QoE.

As key requirements for effective overlay traffic management mechanisms the following two have been identified: (a) to increase overlay traffic locality and (b) to enhance application performance. From these considerations, concrete potentials are derived, mainly for optimization including active caching or topology promotion provided by a network operator or by an adaptation to topology or packet labeling for QoS differentiation by the overlay. The newly proposed mechanisms – subsumed under the term Economic Traffic Management (ETM) – for overlay traffic management lead to the design space summarizing design dimensions and their range of possible values. The illustration of such an approach – by applying these traffic management mechanisms with three practical scenarios of overlay applications – concluded the work.

References

1. [Awduche99] D. Awduche: *MPLS and Traffic Engineering in IP Networks,* IEEE Communications Magazine, Vol. 37, No. 12, 1999.
2. [Hoss04] T. Hoßfeld, K. Leibnitz, R. Pries, K. Tutschku, P. Tran-Gia, K. Pawlikowski: *Information Diffusion in eDonkey Filesharing Networks,* ATNAC 2004, Sydney, 2004
3. [Hoss07] T. Hoßfeld, P. Tran-Gia, M. Fiedler: *Quantification of Quality of Experience for Edge-Based Applications*, 20th International Teletraffic Congress (ITC20), Ottawa, 2007.
4. [Hoss08a] T. Hoßfeld, A. Binzenhöfer: *Analysis of Skype VoIP Traffic in UMTS: End-to-End QoS and QoE Measurements,* Computer Networks, Vol. 52, No. 3, 2008.
5. [Hoss08b] T. Hoßfeld, K. Leibnitz: *A Qualitative Measurement Survey of Popular Internet-based IPTV Systems,* submitted for publication, 2008.
6. [Internet07] *Study:* http://www.ipoque.com/media/internet_studies/internet_study_2007.
7. [Karag05] T. Karagiannis, P. Rodriguez, K. Papagiannaki: *Should Internet Service Providers Fear Peer-Assisted Content Distribution?* Internet Measurement Conference, Berkely, California, U.S.A., October 2005.
8. [Le05] H. Le, D. Hoang, A. Simmonds: *A Self-Organizing Model for Topology-Aware Overlay Formation,* Seoul, 2005.
9. [Ratna02] S. Ratnasamy, M. Handley, R. Karp, S. Shenker: *Topologically-aware Overlay Construction and Server Selection,* INFOCOM 2002, New York, U.S.A., June 2002.
10. [Bindal06] R. Bindal, P. Cao, W. Chan*: Improving Traffic Locality in BitTorrent via Biased Neighbor Selection,* 26th IEEE International Conference on Distributed Computing Systems (ICDCS 2006), Lisboa, Portugal, July 2006.
11. [Fu07] X. Fu, J. Lei, L. Shi: An *Experimental Analysis of Joost Peer-to-peer VoD Service,* Technical Report IFI-TB-2007-03, University of Göttingen, Germany, October 2007.
12. [Brecht03] S. Brecht (P. Racz, B. Stiller, edts.): *Investigation of Application-level Routing,* University of Federal Armed Forces Munich, IIS, Master Thesis, 2003.