Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies

European Seventh Framework Project FP7-2008-ICT-216259-STREP

Deliverable D1.1
Requirements and Application Classes and Traffic Characteristics (Initial Version)

The SmoothIT Consortium

University of Zürich, UZH, Switzerland
DOCOMO Communications Laboratories Europe GmbH, DOCOMO, Germany
Technische Universität Darmstadt, TUD, Germany
Athens University of Economics and Business - Research Center, AUEB-RC, Greece
PrimeTel Limited, PrimeTel, Cyprus
Akademia Gorniczo-Hutnicza im. Stanislawa Staszica W Krakowie, AGH, Poland
Intracom S.A. Telecom Solutions, ICOM, Greece
Julius-Maximilians Universität Würzburg, UniWue, Germany
Telefónica Investigación y Desarollo, TID, Spain

For more information on this document or the SmoothIT project, please contact:

Prof. Dr. Burkhard Stiller
Universität Zürich, CSG@IFI
Binzmühlestrasse 14
CH—8050 Zürich
Switzerland

Phone: +41 44 635 4355
Fax: +41 44 635 6809
E-mail: info-smoothit@smoothit.org

Document Control

Title: Requirements and Application Classes and Traffic Characteristics (Initial Version)

Type: Public

Editor(s): Tobias Hoßfeld, Dirk Staehle

E-mail: hossfeld@informatik.uni-wuerzburg.de, staehle@informatik.uni-wuerzburg.de

Author(s): Hasan, Fabio Hecht, Peter Racz, Burkhard Stiller, Zoran Despotovic, Wolfgang Kellerer, Maximilian Michel, Konstantin Pussep, Sergios Soursos, Haris Neophytou, Jan Derkacz, Jerzy Domzal, Rafal Stankiewicz, Krzysztof Wajda, Robert Wojcik, Michalis Makidis, Tobias Hoßfeld, Simon Oechsner, Dirk Staehle, Maria Angeles Callejo Rodriguez, Juan Fernandez-Palacios

Doc ID: D1.1-v3.6.doc

AMENDMENT HISTORY

Version	Date	Author	Description/Comments
V0.1	January 23, 2008	Tobias Hoßfeld	First version, proposing table of contents
V0.2	January 30, 2008	Tobias Hoßfeld	feedback from mailing list to list of overlay applications and examples
V0.4	April 14, 2008	Tobias Hoßfeld	including the ToC as fixed during Athens meeting
v1.0	April 17, 2008	Tobias Hoßfeld	merging bullet lists and input received from all partners
V1.1	April 21, 2008	Juan Fdez-Palacios	TID's contribution to section 5
V1.2	May 10, 2008	Sergios Soursos	AUEB's contribution to section 5
V1.3	May 16, 2008	Jan Derkacz, Jerzy Domzal, Robert Wojcik	AGH's contribution to section 5
V1.4	May 16, 2008	Haris Neophytou	PrimeTel's contribution to section 5
V1.5	May 20, 2008	Peter Racz, Hasan, Fabio Hecht	UZH's contribution to section 5
V1.6	May 22, 2008	Sergios Soursos	Merging and editing of section 5
V1.7	May 24, 2008	Sergios Soursos	Re-structuring of section 5 after UZH's proposal and final editing
v2.0	May 25, 2008	Tobias Hoßfeld, Maria Angeles Callejo Rodriguez, Hasan, …	integrating Section 5 and 6
V2.10	June 08, 2008	Tobias Hoßfeld, Konstantin Pussep, Michalis Makidis, Sergios Soursos, Zoran Despotovic, Maria Angeles Callejo Rodriguez	inputs integrated for Sections 3, 4, 7
V3,2	June 21, 2008	Tobias Hoßfeld, Wolfgang Kellerer, Sergios Soursos, Peter Racz, Michalis Makidis, Jan Derkacz, Haris Neophytou	update of sections; this version is send for final review to Wolfgang Kellerer and Maria Angeles Callejo Rodriguez
V3.6	June 28, 2008	W. Kellerer, Maria Angeles Callejo Rodriguez, Peter Racz, T. Hoßfeld	review and comments of entire deliverable, update of Section 6.2, addressing review comments
V3.7	July 24, 2008	Burkhard Stiller	Final checks, TriplWin included, submission prepared, ok
V4.0	March 31, 2009	Tobias Hoßfeld, Wolfgang Kellerer, Burkhard Stiller	Review Comments embedded

Table of Contents

1 Executive Summary

The objective of the SmoothIT project is to define, develop and test Economic Traffic Management (ETM) mechanisms to optimize the traffic impact of overlay applications on ISP and telecommunication operator networks based on a cooperation of network operators, overlay providers and application users. This deliverable summarizes the initial investigations in respect with overlay applications, their classification, their characteristics, and their significance for ETM and explains initial ETM approaches. A set of requirements for the system design and development in SmoothIT are also derived from this study and reported in the document.

This deliverable contains the results of the Tasks T1.1 “Overlay Application Classification”, T1.2 “Economic Discussion of Overlay Applications”, and T1.4 “Determination of Overlay Applications Utilized in Trials” running from Jan 08 – Jun 08. In addition, the milestone MS1.1 “Application and Traffic Descriptions for Internal Trial” is documented in this deliverable.

The major objectives of D1.1 are 1) to give an overview on overlay applications describing their features and traffic characteristics, 2) to provide an overview on technical, economic, and functional aspects of overlay applications within SmoothIT, 3) to identify the high-level requirements for the SmoothIT architecture based on the technical and economic discussions, 4) to classify the overlay applications for SmoothIT, and finally 5) to select the overlay application to be implemented for the internal trial. These objectives are reflected in the structure of D1.1.

The main outcome of this deliverable is:

- A set of **18 functional and non-functional requirements for the SmoothIT system** design (Section 6)
- **Selection of P2P-based streaming as the reference overlay application class for SmoothIT** (Section 7)
- Selection of **a concrete P2P-based Video on Demand overlay application implementation** for its usage in the internal trial (Section 7)

The overlay application selection process is based on an investigation of seven overlay application types (Section 2) and selected overlay application examples (Section 3) and their initial discussion with respect to their relevance for SmoothIT, i.e., the potential to apply Economic Traffic Management most efficiently. This initial discussion leads to the specification of nine classification criteria (Section 4) indicating the overlay application relevance for SmoothIT that are finally used for the application selection (Section 7). In addition to the application classification criteria discussion, economic and regulatory mechanisms are investigated (Section 5) before concrete requirements for the SmoothIT architecture can be derived (Section 6). Section 7 finally presents the final application classification resulting in the prioritization of the Peer-to-Peer (P2P) based streaming application.

This selected overlay application class forms the basis for all theoretical investigations in work package WP2. It has also been selected as the target overlay application class for our development in respect with the internal trial. The determined overlay application for the internal trial impacts the engineering and test-bed integration in WP3, as well as the trial integration and assessment in WP4. In the follow-up deliverable D1.2 “Commercial Application Classes and Traffic Characteristics”, (due M19) we will further focus on the selected overlay application class. We evaluate traffic analysis methods and apply them to the selected application class. In addition, we will utilize existing ISP data and data from traffic measurements to develop a suitable traffic analysis model for overlay applications.

In the following, we summarize the outcome and the achieved goals of this deliverable D1.1 in more detail in respect with the particular sections.

Section 2 ***"Overlay Application Service Types"*** summarizes the variety of existing overlay application service types in the Internet. They are distinguished according to the offered service types which are file sharing, video-on-demand (VoD), live TV streaming, voice-over-IP (VoIP), gaming, content distribution networks (CDN), and virtual private networks (VPN). For these service types, various applications exist which may significantly differ in their actual implementation and performance. However, for all applications of such a service type, there are common Quality of Service (QoS) requirements and traffic characteristics, as well as Quality of Experience (QoE) aspects. Based on these characteristics, we show the optimization potential to achieve a win/win/win situation (TripleWin) between network operators, overlay providers and application users. Hence, we determine the relevance of this overlay application service for SmoothIT. As a result, we consider file sharing, P2P VoD and P2P live TV as being interesting for the SmoothIT project, while the overlay application service types P2P VoIP, P2P gaming, CDNs and VPNs are less interesting due to lower optimization potential or traffic impact. This reasoning is a first step for the selection of the main application to be focused on in SmoothIT and, in particular, for the internal trial.

In Section 3 ***"Discussion of Overlay Application Examples"***, the most prominent examples for the service types described before are presented. Significant characteristics of individual applications are highlighted regarding their technical, QoE related, and economic characteristics. Relevant differences among the example overlay applications are discussed revealing their impact on a particular SmoothIT solution. A main focus is put on the optimization potential for an overlay application. We have marked, if it differs from general observations found with the respective overlay application service type in Section 2. As a result of the analysis, file sharing applications are of interest for SmoothIT mostly due to the high optimization potential. In particular, open-source variants for the popular eDonkey and BitTorrent protocol are available to be used in internal and external trials, including the option to use Chordella for lookup. P2P streaming is regarded as a future-proof application class. Usage and popularity of P2P live TV and video-on-demand applications is expected to increase in the next years. Moreover, a performance benefit for a user in terms of visible performance improvement has a higher potential when using a SmoothIT mechanisms for ETM. Thus, the P2P video streaming class offers even more optimization potential than file sharing, as topology promotion as well as QoS provisioning may be utilized. However, most popular applications are proprietary implementations what makes it particularly hard to influence them to achieve the TripleWin situation. In this context, Vuze or Tribler (used in the P2P Next project, see Section 7.1) have been identified as promising P2P streaming applications, since they are open-source.

Section 4 ***"Classification Criteria"*** explains the classification criteria that are used in SmoothIT in detail. With respect to the application selection process in SmoothIT, we distinguish between major and supplementary classification criteria. Additionally, we discuss the impact of the technical environments besides classical ISPs operating fixed networks that have to be taken into account when classifying overlay applications. Heterogeneity of network systems (including wireless networks) and user mobility have an impact on all evaluated overlay applications. As they are not addressed in most of the currently available applications, their impact is considered as a general requirement for the project work. The major classification criteria fall into three different categories: technical criteria, optimization potential and non-technical criteria. The technical criteria comprise traffic intensity and characteristics, source code availability, and traffic recognition and emulation. The optimization potential addresses end-user controllability, which might be utilized in an automatic way by the SmoothIT solution, the utilization of QoS provisioning and the utilization of locality information. In the category of non-technical criteria, the popularity of the overlay applications, the emerging ISP costs, the legality of content, as well as the acceptance for additional charging for an improved user experience are taken into account. These criteria are used in Section 7 for the classification of the overlay application service types, described in Section 2, and the overlay application examples, described in Section 3.

In order to derive the requirements, not only overlay application characteristics are analyzed, but also the potential incentives for the different stakeholders and their effects are investigated from an economic and regulatory viewpoint in Section 5 ***"Economic and Regulatory View"***. From the analysis, we see that incentives provided to one stakeholder may introduce negative effects to another one. For example, the performance improvements that an overlay provider may want to introduce may come in direct conflict with the economic incentives for the operator (ISP), since such improvements may change the traffic patterns, affecting the interconnection agreements and charges for the specific ISP. Incentives for end users include performance improvement, peer availability and peer reputation. For overlay providers incentives are performance improvement and user loyalty toward the overlay application. Operators could benefit from traffic optimization and related general performance improvement and loyalty of users selecting and staying with the ISP. Monetary benefits may also be applicable strongly depending on the application. Based on an investigation of overlay related costs, possible ETM mechanisms such as interconnection agreements and locality promotion were discussed initially. A more detailed description and evaluation can be found in deliverable D2.1.

The overlay application and incentive discussion lead to the following high-level functional requirements for the SmoothIT architecture design manifesting around the lack of information exchange (information asymmetry) between overlay providers and network operators. This is described in Section 6 ***"Functional Requirements"***. This section identifies the key information that should be shared or exchanged between overlay and underlay, in order to address the current information asymmetry. Furthermore, it also discusses possible ways to share this information. An information service (referred to as SmoothIT Information Service, SIS) should be provided by network operators to optimize overlay traffic taking into account the underlying network. Distributed SISs should interact. They should provide an open, reliable, scalable service that can be differentiated in free and premium services provided anonymously or customer aware.

Section 7 ***"Overlay Application Classification for SmoothIT"*** evaluates the overlay applications and applies the introduced classification criteria to the individual overlay applications. The evaluation is based on technical and non-technical criteria in order to judge their relevance for the SmoothIT objectives. The introduced evaluation criteria include traffic intensity, traffic recognition, optimization potential, popularity, legal content and charging possibilities. In addition, overlay applications are characterized by their design parameters such as overlay algorithms, overlay topology, or QoS requirements, which are implicitly contained in the optimization potential discussion and have more impact on WP2 algorithm development. As a result of this evaluation, P2P-based video streaming and P2P (BitTorrent-style) based file sharing applications have been identified as most relevant and important for SmoothIT. This is motivated by their high traffic impact, their popularity and their optimization potentials such as locality promotion, as well as the application service classes seem to be future-proof. In particular, P2P streaming was prioritized as its popularity is expected to increase overtaking file sharing and it shows a higher optimization potential, i.e., sensitivity in respect with QoS and QoE. This is also reflected in the trial application selection where P2P VoD was selected for the internal trial. Here especially the availability of open source software was taken into account additionally. The SmoothIT consortium decided in unison to devise not one but two action plans concerning the decision on the selection of the application to be utilized for the internal trial. Plan A is driven by a strong desire of the SmoothIT project to engage in a bilateral beneficial agreement with the P2P NEXT project and the availability of the source code of the BitTorrent based application. Plan B suggests selecting Vuze (formerly called Azureus) as the overlay application for the internal trial which is based on the popular Azureus Java BitTorrent client implementation with optional integrated adaptations for video streaming.

The achieved goals of this deliverable are summarized in Section 8.

This deliverable documents all initial project investigations in respect with high-level requirements and application selections. These findings will be updated in M1.2 "Application and Traffic Descriptions for External Trial" in D1.2. In addition to D1.1, we will provide later on the Deliverable D1.1-Annex "SmoothIT Related Work" which contains a comprehensive discussion of SmoothIT related work and is a living document. The content of this document D1.1-Annex can be used for various forthcoming deliverable and publications of SmoothIT.

2 Overlay Application Service Types

The goal of this section is to describe and summarize the variety of existing overlay application service types in the Internet. They are distinguished according to the offered service types which are file sharing, video-on-demand (VoD), live TV streaming, voice-over-IP (VoIP), gaming, content distribution networks (CDN), and virtual private networks (VPN). For these service types, various applications exist which may significantly differ in their actual implementation and performance. However, for all applications of the same service type, there are major common quality of service (QoS) requirements and traffic characteristics, as well as quality of experience (QoE) aspects. These common points and requirements from application's point of view will be discussed and outlined in this section. As a result, we show which of the application service types are skipped for further investigations since no optimization potential is foreseen or since the impact on global network traffic is not relevant and, therefore, they are not interesting for SmoothIT. This reasoning is a first step for the selection of the actual application used in SmoothIT further on and, in particular, for the internal trial.

Before that, we introduce and **define notions and terms** recurring in this deliverable.

- A ***node*** is an entity within a telecommunication network. If not mentioned explicitly, nodes are physical entities which are connected and form the network.
- ***Underlay* or underlying network** is used for the native networks which might span several metro and access networks, as well as core IP networks of several Internet Service Providers (ISPs) or Telecommunication Service Providers (telcos).
- An ***overlay*** or an overlay network is a flexible, logical network which is built on top of another physical network, i.e. underlay. In the context of SmoothIT, we define that flexibility is a mandatory feature of an overlay. Flexibility means that the overlay and the connections among the overlay nodes can be adapted, i.e., individual routes among overlay nodes are based on local decisions of the overlay nodes. To define this precisely, the logical nodes in an overlay network are referred to as *overlay nodes*, which must not necessarily correspond to physical nodes.
- In a ***peer-to-peer (P2P) network***, the nodes of this network called *peers* share resources, e.g., bandwidth or memory storage, in order to provide or support a certain service, like file sharing. Typically, the peers form an overlay for communicating with each other.
- The ***application*** refers to a software program designed for end-users. An ***overlay application*** is an application which utilizes an overlay for communication. It has to be noted that two or even more different overlays may be used to provide some particular functionality. The basic functions are a) resource mediation and b) resource access control. On one hand, r*esource mediation* is the capability to locate resources. For example, a distributed hash table may be used for lookup of files in a file sharing system resulting in a structured search overlay. On the other hand, r*esource access control* specifies how resources are exchanged or shared among the users of this application. For example, in order to exchange video data in a live TV system, an application-layer multi-cast tree might be established as communication overlay for distributing the real-time video contents.

In the context of SmoothIT we distinguish between relevant and interesting applications. An application is called a *relevant application* if it is used in practice and causes problems or significant costs in the network of an ISP or telco. Thus, a relevant application is of importance for network providers. An *interesting application* means additionally that there is optimization potential with respect to SmoothIT, i.e., the application might react on appropriate incentives to adapt itself and the overlay topology (to its own and the operator's needs). Thus, as outcome of this deliverable we want to highlight interesting overlay applications and reason which ones are most promising for SmoothIT.

2.1 File Sharing

File sharing overlay applications are very popular among users and widely used in the Internet, especially to share files of large size, like copyright-protected movies, music albums, or software releases. Therefore, file sharing applications generate a huge amount of traffic volume in the Internet. Due to the scientific interest in P2P file sharing systems, additional several academic implementations exist. However, there are only a few main file sharing systems interesting for SmoothIT, like BitTorrent or eDonkey (which will be derived in Section 3). Nevertheless, there exist a lot of different applications for the same system, *e.g.,* BitTorrent, which speak the basic protocol, *e.g.*, BitTorrent protocol. For the most popular file sharing applications, open source variants exist which can be modified according to the user's own needs. The rich set of different file sharing systems and various available applications software allows the user to decide and select a particular platform out of these. In addition, the user can control and set some parameters of the application to its own benefit via a graphical user interface or a parameter configuration file. Among the typical parameters are for example the maximum numbers of parallel upload and download connections, the maximum shared upload capacity, or the port numbers. Often, proxies can be announced for redirecting the communication paths.

The mechanisms to control and manage content distribution in P2P networks can be distinguished in two major categories: a) resource mediation mechanisms, which are functions for searching and locating resources and b) resource access control mechanisms, i.e. function for exchanging files or parts of it. As a result, P2P applications for file sharing may form two logical overlays, a) search overlay and b) file distribution overlay.

There are several approaches focusing on *resource mediation mechanisms* including resource discovery mechanisms. They vary from centralized concepts such as index servers or trackers, as in eDonkey or BitTorrent, to highly decentralized approaches such as flooding protocols, as in the original Gnutella network, or distributed hash tables, as used in the Chord protocol. Especially, the DHTs and hierarchical derivates have gained a lot of scientific interest addressing refinements to cope with reliability and efficiency, *e.g.*, in cellular environments. Special architectural entities like crawlers are used to locate files and sources of files on behalf of other users to improve the performance. This is especially important in mobile environments with scarce and expensive resources of users.

The *resource access control mechanisms* determine the coordination and cooperation among peers which means to permit, prioritize, and schedule the access to shared resources. In this context, incentive mechanisms are implemented to promote cooperative behavior. This means they try to make peers participate in the network and share their resources. Examples are credit point systems as used in eDonkey or tit-for-tat strategies like in BitTorrent. An efficient and robust way of cooperative content delivery is the multiple source download (MSD), which means that the recipient peer orders and downloads the desired data from many providing peers instead from a single one. The efficiency of MSD was demonstrated by the success of the P2P files sharing platforms eDonkey and

BitTorrent, which are able to overcome flash crowd by their scalability features. Gnutella and Wuala also implement MSD for an efficient content distribution, cf. Section 3.1. It has to be noted that Wuala is a file storing system instead of a file sharing system, like eDonkey, BitTorrent or Gnutella. The main difference between storing and sharing is that a user of a file storage system wants to access his/her own files, while a file sharing user aims at downloading new files injected by other users. However, the basic mechanisms to coordinate the resources among users in both systems are the same.

P2P applications for file sharing and storing implement resource mediation as well as resource access control mechanisms. Pure P2P architectures are implementing both mechanisms in a fully decentralized manner, while *hybrid P2P systems* utilize central entities that collect mediation data. In *hierarchical P2P systems,* special entities like super nodes accomplish more tasks than regular peers, which often result in hierarchical overlay structures instead of unstructured meshed networks. This can happen in the resource access (file distribution) overlay, as well as in the resource mediation (search) overlay. However, in efficient file sharing overlay networks the major traffic volume is generated by the exchange of content data. Therefore, the adaptation of the file distribution overlay is of major interest in the context of SmoothIT. Typically, in file sharing networks, the peers form a meshed overlay due to the MSD mechanisms and the related incentive mechanisms. However, the search overlay or the central entities providing lookup functionality deliver the sources of files to requesting peers and hence determine implicitly the topology of the data exchange overlay.

The traffic characteristics of different file sharing applications are mostly the same. The signaling traffic as well as the data traffic is transported via TCP. Most applications are bandwidth demanding and try to fully utilize the available upload and download bandwidth of a particular user. However, the traffic is elastic which means that the data transfer is tolerant to throughput variations. Current file sharing applications are not topology or location aware and establish many parallel connections to available sources of files. Many upload and download connections are established at the same time.

From the user's point of view, the quality of experience is mainly indicated by the average throughput. However, since the downloading user is also participating in the network by providing resources to other users, the QoE might also be influenced by the amount of uploaded data or the currently consumed upload bandwidth. In the case of a large upload bandwidth, the user might be disturbed in using different applications like web browsing or VoIP. Nevertheless, in general, QoE is less critical for file sharing applications, because users might also tolerate longer waiting times for the download, *e.g.*, when starting downloads of large contents over night.

For the usage of file sharing applications, no special requirements have to be satisfied. However, the performance of the various systems differs significantly and depends on particular features. In eDonkey for example, a user behind a NAT or a firewall gets a worse performance due to a so-called 'low-id', as other peers cannot initiate a connection to this peer. In BitTorrent, the perceived download throughput of a user depends heavily on its upload bandwidth due to the tit-for-tat principle.

For sure, file sharing traffic is a relevant application and causes a large amount of traffic and costs for ISPs and telcos. For SmoothIT, file sharing is also interesting, as there is a lot of optimization potential. The download from close peers in the same domain could reduce the download times as well as inter-domain traffic. As a necessary instrument, locality promotion has to be integrated by the ISP. This can be enabled by cooperation strategies (also called sharing strategies), which achieve cooperative behavior among the peers in the overlay and minimizes ISP's costs or maximizes ISP's revenue while decreasing the download times. If peers are given incentives to enhance and utilize locality awareness, then this could be considered as an economic traffic management (ETM) mechanism.

Furthermore, it is of interest to provide robust incentive frameworks which are resilient to attacks by malicious and greedy peers. For example, poisoning and pollution are caused by malicious peers in a file sharing system. This means that fake or corrupted contents are offered by the malicious peers which disturb the entire file dissemination process. As a result, more traffic is transferred before the user gets the desired content correctly or the user aborts the download leading to a bad experience with this file sharing system.

2.2 P2P Video-on-Demand

There are already a lot of P2P-based video-on-demand applications available in the Internet, such as Joost, Vuze, SopCast, or PPLive. However, the popular applications offering popular contents are mainly proprietary and the source code of the client's application is not open. Similar to file sharing applications, the scientific community has a strong interest in P2P VoD and several open-source applications are available. But in that case, the offered content is not interesting for the majority of the Internet users and thus the open-source P2P VoD applications do not generate a large amount of data traffic. Nevertheless, the popularity of the proprietary systems shows that P2P VoD applications are interesting for SmoothIT, see the results in Section 7.1. In the near future, it is expected that their popularity and thus the traffic volume is further increasing.

In contrast to file sharing applications, P2P VoD requires more sophisticated resource mediation and resource exchange mechanisms in order to cope with the stronger traffic requirements. The average throughput in such a video streaming system has to be higher than the actual video bit rate in order to provide a good QoE. Packet loss, reordering of packets and large jitter in the network may lead to a strong QoE degradation. This means the QoE sensitivity is high.

At each user, a small *jitter buffer* or *playout buffer* is required to overcome packet delay variations and guarantee a smooth playout of the video stream. Additionally, a large *video storage buffer* at the peers is required in order to make all video contents available in the network, as a user decides on contents on demand. Although server farms are often used to support the storage of contents (*e.g.*, Joost video servers transport a large amount of video data to the users currently), the server systems are not scalable with the number of requesting users and the amount of stored contents. Thus, the video storage buffer has to be large enough to provide a reliable VoD system.

The topology of the data exchange overlays of P2P VoD systems is determined by the applied multi-source download mechanisms. Similar to file sharing, the peers in such systems download parts of the video contents in parallel from different peers and video content servers, while simultaneously uploading to requesting users. The resulting topologies range from meshed overlays to structured, tree-based overlays. However, the tree-based overlays are only useful in case of flash crowds, when many peers want to see the same content within a short period of time. This might happen for example when new video contents are inserted into the system. Especially in the case of flash crowds, the scalability features of P2P systems, either mesh-based overlays or tree-based, are utilized to cope with the heavily increased resource demands. For the prevailing proprietary, popular P2P VoD applications it is assumed that they use mesh-based overlays, while the underlying MSD mechanisms are BitTorrent-like, but adapted to video streaming requirements. For example, [DLH+05,VIF06,SP07,LCJ+07] propose enhancements of BitTorrent for supporting streaming, however, their applicability and performance in real environments has to be investigated. Instead of downloading always the least-shared chunks of a file, an individual peer has to consider the deadlines for the chunks. A chunk received too late is neglected on application later and is perceived like packet loss, thus significantly decreasing the QoE. Nevertheless, a 'too late' chunk can be further disseminated to other users.

The traffic characteristics of typical P2P-based VoD applications show that these applications are less aggressive and demanding than file sharing applications. It can be seen that the available bandwidth is neither in uplink nor in downlink direction fully utilized. At a peer, the download bandwidth often appears constant on a larger time-scale. However, short bandwidth peaks arise often due to signaling. Current implementations are neither topology nor location-aware, however, due to license agreements the traffic for particular video contents is restricted to specific countries [HL08]. Typically, the video contents are transported via UDP, while the signaling traffic goes via TCP.

Similar to file sharing, locality promotion can be utilized to improve the download of the video data. In addition, QoS provisioning might be useful to handle large jitters or packet loss on links. However, it has to be noted that for VoD applications the delay variations are not crucial for the QoE, as long as the jitter buffer overcomes them, thus, QoS provisioning might be overkill. Nevertheless, P2P VoD is interesting for SmoothIT due to its optimization potential.

2.3 P2P live TV

P2P-based live TV (television) streaming applications are comparable to P2P VoD in many aspects including the current and expected popularity of these applications among users, the amount of traffic in the Internet, the availability of source code, as well as the traffic characteristics. However, P2P live TV systems differ significantly in the requirements of the application itself compared to P2P VoD. Additionally to the requirements in terms of average throughput larger than the actual video rate and small jitter delays, P2P live TV systems only allow a maximum delay in the system. As a result, these systems require only a low storage demand, as only video contents for this maximum tolerable delay have to be made available in the system. Therefore, the basic resource access control and resource mediation mechanism differ significantly from P2P VoD or file sharing systems. Different mechanisms are required to fulfill the even stronger QoE and QoS requirements than for VoD. These additional aspects will strongly affect the SmoothIT solutions.

While in literature, application layer multicast solutions are proposed, see survey on multicast communication [PCE+07] and references therein, it is not clear whether the prevailing systems in the Internet use such tree-based overlays or if the peers form a meshed overlay, as in typical file sharing overlays. In order to construct and maintain the structure of a tree, a lot of signaling overhead may be generated. The reason for this is that leaving and failing peers cause a restructuring of the tree. Thus, it is important to predict and monitor the availability of peers and resources to deal with churn. Although simpler mesh-based algorithms implementing a push- or pull-based multi-source download mechanism might be less efficient in theory, these practical issues make them be used in current P2P live TV software. Another hint for this is the fact, that several P2P live TV applications offer also VoD. But, again the popular systems are closed-source and measurement studies are required to derive the resulting topologies which will be addressed in Task 1.3 in SmoothIT.

As stated above, the traffic characteristics of P2P live TV systems is comparable to P2P VoD in terms of upload and download bandwidth, as well as in the number of upload and download connections of a peer. Thus, a peer will not completely utilize the entire download bandwidth, but only as much as is required for smooth video playout. However, in contrast to P2P VoD, more flash crowd effects are expected due to special events, like the European Football Championship, or common viewing times, *e.g.*, 8 o'clock in the evening.

The strict QoS requirements in terms of end-to-end delay, jitter, loss, and throughput shows the optimization potential for P2P live TV. Beside the utilization of topology promotion to select near peers in the same ISP's network with smaller e2e delays, QoS provisioning might be beneficial to fulfill the service requirements. A combination of both means that the locality-awareness of the overlay will make the content keep locally within the ISP's network, such that QoS provisioning of the connections within the ISP's network results in QoS provisioned e2e path. Thus, a mutual benefit may be achieved by taking into account topology awareness and QoS provisioning.

2.4 P2P Voice-over-IP

Although P2P VoIP seems to be quite interesting because of the popularity of Skype, a pure P2P VoIP application is of minor interest for SmoothIT for two reasons. First, the totally exchanged amount of traffic due to VoIP is not the major source of costs for ISPs and telcos. Second, the voice call is usually directly established between the caller and the callee of an application. Only in case of NATs/firewalls or bad e2e connections, the call is relayed via a third peer, which is in the case of Skype a super peer. The selection of this super peer could be done with respect to the locations of the two end voice users, such that the costs are minimized for the ISP and a good voice quality is available on the e2e link for the end users. However, the optimization potential for this kind of applications is rather low, compared to file sharing, VoD or live TV. Due to the strict QoS requirements in terms of e2e delay, jitter, and packet loss of voice calls, QoS provisioning is useful for the end user. However, this is mainly an ISP internal approach and thus not interesting for SmoothIT.

The resource mediation mechanisms for lookup of users and super peers might be realized with Distributed Hash Tables (DHT). In this case, topology-awareness can be taken into account to form the structure of the DHT appropriately. Again, this won't have a large impact on the costs of an ISP.

The required bandwidth for a VoIP call is also low compared to real-time or non-real-time video-streaming. The required signaling traffic for peers is marginal, only super peers show a slightly higher throughput which is described in Section 3 for Skype in more detail. It has to be noted that the traffic characteristics and requirements will change for voice or video conferencing. However, for SmoothIT the optimization potential and reachable gain is still lower than for the other applications.

While QoS provisioning as ISP internal solution is useful for the QoE of P2P VoIP, other optimization potential like locality promotion is not a key issue. Therefore, P2P VoIP applications will be skipped in the application selection process in SmoothIT.

2.5 P2P Gaming

For P2P gaming, there are currently no implementations available which are used in the Internet. Therefore, this kind of overlay applications is skipped in SmoothIT. However, it has to be noted that in near future an increased popularity of P2P gaming applications is expected which might mainly depend on a successful (likely commercial) release of a P2P gaming software.

The traffic characteristics and requirements of the gaming application strongly depend on the supported gaming type. For example, an ego shooter game like Unreal Tournament (which is not based on P2P) is comparable to P2P VoIP conferencing. The users playing against each other are located in the same arena map, thus, the information has to be delivered to all players in the map. However, these are typically only a few ones. For this kind of game, QoS requirements in terms of delay and especially of jitter are crucial. Topology information cannot be utilized. This is different for massive multi-player online games like World of Warcraft (also not P2P-based). Here, the player are acting in a larger world, however, the current users only see a small part of this. Thus, a P2P-based approach allows storing, computing and updating parts of the world in the background, which are delivered when the user enters this part of the map. For this kind of games, bursty traffic patterns might appear according to the virtual movements of the different players. Depending on the actual game context, this application also requires strict QoS requirements, *e.g.*, when players are fighting, while during the rest of the game small jitter and delays are tolerable.

2.6 Content Distribution Networks

In order to clearly distinguish CDNs from file sharing networks, we provide the following definition for SmoothIT. CDNs are commercial applications which appear as C/S to the end-user. Server farms are used as basic technology. Thus, there is no significant upload traffic of users in CDNs. However, CDNs build overlays which generate large amounts of traffic, as there are many users and often large contents are downloaded via CDNs.

For sure, CDNs can be optimized by utilizing locality information, *e.g.*, Akamai also does, but the CDNs likely do not offer much 'freedom' for the SmoothIT purposes. This means the topology cannot be influence by the hosting ISPs. In fact, CDN nodes are placed on the ISP/IBP premises with their agreement. ISP and overlay provider (like YouTube, Akamai, AmazonS3) will likely agree on such bilateral contract to regulate the interrelationships. Therefore, we conclude that although CDNs are relevant overlay applications, they are not interesting for the SmoothIT project.

2.7 Virtual Private Networks (VPN)

VPN are becoming an attractive and easy to use solution to allow the end users to use some applications as they were in the same private network using the Internet, skipping all the problems associated to NATs, Firewalls. One of the most common usages of VPNs in the residential environment is gaming scenario, where several players distributed in the Internet can play against each other as they were in the same LAN. Even though this scenario is not quite representative in terms of amount of traffic for an ISP, it is very interesting due to the users' requirements to have Local Area Network conditions, in terms of delay or jitter. Therefore, this scenario could be a representative scenario in terms of providing (and maybe charging) advanced connectivity capabilities to the end users. This example is later used in Section 3 for discussion.

In the context of SmoothIT, VPNs are skipped as not relevant to ISPs, since no high amount of traffic is generated by residential users' VPNs. It has to be noted however that VPNs might be useful to enable economic traffic management.

3 Overlay Application Examples

This section presents the most prominent examples of overlays of the service types described in the previous section. We point out significant characteristics of individual applications regarding their technical, QoE-related, and economic characteristics. Finally, we discuss relevant differences among these examples revealing their impact on the SmoothIT approach.

It has to be noted that in the deliverable D2.1 "Self-Organization Mechanisms for Economic Traffic Management (Initial Version)" the applied mechanisms for resource mediation and resource access control in the search overlay and the data exchange overlay are described for the different applications in detail, respectively.

3.1 File Sharing and Storing

In the following, we will discuss eDonkey, BitTorrent, and Gnutella which are the most popular file sharing systems. They have in common that there does not exist one single application for these systems, but a large variety of derivates speaking the basic protocol for each system. Additionally, we will take a look on Wuala used as permanent file storage system. All these applications implement MSD as fundamental resource access control mechanism, while for resource mediation and searching for contents centralized as well as distributed approaches are implemented. The idea of this section is to highlight the major differences between these applications which are of importance for SmoothIT. We start with eDonkey and go into more detail for this application, as basic observations are similar to BitTorrent and Gnutella. In the context of file sharing, Chordella as lookup system is of interest, since it uses a hierarchical DHT for lookup based on Chord and provides much optimization potential for SmoothIT, especially in heterogeneous environments.

3.1.1 eDonkey

eDonkey is a protocol for exchanging files which establishes a content delivery overlay, denoted as eDonkey network. There are several open-source applications available which implement the eDonkey protocol which is described *e.g.*, in [HB03,Go03,HLP04]. However, the actual implementations of resource access control mechanisms may differ. Popular derivates are for example eDonkey2000 [Ed08], eMule [Em08] or MLDonkey [Ml08] which are subsumed as *eDonkey application / client* in the remainder. eDonkey is one of the most popular file sharing applications and generates a large amount of traffic within the ISPs network. The popularity, however, strongly depends on the country which is closely related to the offered content. In Spain, for example, large index servers are operated offering among others a lot of Spanish spoken content. In US, however, BitTorrent is the dominant protocol and eDonkey only plays a minor role.

eDonkey belongs to the class of hybrid P2P architectures and uses two software implementations: the open-source eDonkey client and the closed-source eDonkey server. The eDonkey client stores, shares and downloads files. An eDonkey server operates as an index server for file locations and distributes addresses of other servers to clients. For joining the eDonkey network, a client connects to an arbitrary index server and signals his shared files. This means the index server knows all files available in the network. This index server is proprietary, but freely available and can be operated by any user.

Before an eDonkey client can download a file, it first gathers a list of all potential file providers from the index server. Then, the clients connect to each other directly for exchanging files. The MSD in eDonkey is enabled by dividing files into chunks with a size of approximately 10 MBytes. The consuming client can reassemble the file using the chunks or parts of chunks. When an eDonkey client starts to download a file, it asks the providing peers for an upload slot. Upon reception of a download request, the providing client places the request in its upload queue. The upload management of a peer maintains an upload queue which consists of two lists, the waiting list and the uploading list. The uploading list holds the requests which are currently served. A download request is served as soon as it obtains an upload slot, i.e. it moves from the waiting list to the uploading list. Each served request gets an amount of upload capacity according to its credits. eDonkey uses a credit system in order to foster a peer to share resources, i.e. upload bandwidth and storage capacity. A user earns credits for uploading to a client, which are then spent to advance faster in this client's queue. In particular, credits are earned for uploading any file. As a result, for eDonkey there exists a single content delivery overlay over which all files are exchanged. However, different modifications, called mods, exist which may change the upload management. As open source variants exist, each user has the possibility to modify the source code to its own needs, *e.g.*, prioritizing users with a certain IP prefix for any reason, and to release a new mod.

In eDonkey, TCP is mainly used for signaling data transfer. There is additional UDP traffic between peers and index servers; however, the amount is negligible at roughly 1% of the whole eDonkey traffic. In particular, the ratio of signaling traffic volume to data traffic volume is about 1:15. For file sharing, there are no strict delay and bandwidth requirements. However, 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 short time. Additionally, he might want to reduce its own upload traffic. The reduction of upload traffic is of interest, as the upload bandwidth of a user is typically much smaller than the available download bandwidth. Especially in mobile environments, the upload bandwidth is an expensive resource. While running eDonkey, the users does not want that other applications, like web browsing, are negatively influenced or experience a worse quality. An ISP transporting the eDonkey traffic follows two interests. First, he wants to improve the quality of experience for its customers as selling argument. Second, he wants to reduce the costs for transit traffic to other ISPs.

For satisfying the user's and the ISP's interest, the *adaptation of the overlay topology* to the underlying topology is a key solution. The download from close peers in the same domain reduces the download times as well as inter-domain traffic. As necessary instrument, topology promotion has to be implemented by the ISP. The user's and the ISP's benefit can be further fostered by caches. A cache might also reduce the inter-domain traffic as well as the upload traffic volume of peers while providing high download speeds.

3.1.2 BitTorrent

The BitTorrent protocol [Co03,Co08] was implemented with the objective to disseminate one large file (or a composition of large files) to a large number of users in an efficient way. Therefore, for each file an overlay network called swarm is created, in contrast to eDonkey or Gnutella. The file sharing is based on MSD, also denoted as swarming principle in this context.

According to the original BitTorrent specification, each overlay network consists of two different kinds of peers, the seeds and the leechers, and a so-called *tracker*. A seed holds the complete file and uploads to others altruistically, whereas a leecher is still downloading the file. The tracker is a centralized component which stores information about all peers. A new peer, who enters the network, asks the tracker for a list of active peers in the overlay. The tracker returns a random subset to the requesting peer. Furthermore, an active peer contacts the tracker from time to time to obtain information about new peers in the network. An extension of the protocol also incorporates the exchange of information about other peers in the network between connected peers. This is often stated as trackerless BitTorrent. Any user can operate freely such a tracker. A *torrent file* contains static meta-information about the location of the tracker or about peers to join a swarm with trackerless BitTorrent. Usually, a web server hosts such a torrent file, but of course, it can also be distributed by any other protocol like FTP.

BitTorrent specifies the messages between the tracker and a peer and between peers themselves. Furthermore, it implements two important algorithms which are run by each peer. These are the peer selection and the piece selection algorithm. The peer selection process is called choking/unchoking in BitTorrent. Each peer controls to whom it uploads data. Each peer uploads to a fixed number of other peers (the default value is four). Thereby, a peer chooses to upload to other peers from which it has the highest download rates. As a seed unchoking is based on the download rate of the connected peers rather than the upload rate. By default this tit-for-tat strategy is run every ten seconds by every peer, whereby the download rates are determined by a moving average over the last 20 seconds.

To discover new peers with better performance a so-called optimistic unchoke is done additionally. Here, one of the peers is unchoked independently of its rate. The optimistic unchoke is changed every 30 seconds to provide enough time to be possibly unchoked by the remote peer in return. Another rule in BitTorrent is to choke a peer when it has sent no data message in the last minute. This is called antisnubbing.

The piece selection algorithm determines which file fragment is requested when a peer is unchoked by a remote peer. The decision process follows the following rules: Firstly, when some bytes are received from a specific chunk the remaining parts of that chunk are requested. This scheme is called strict priority. When strict priority is not applicable, the rarest chunk is requested. Since a peer has only a local view of the network it can only estimate rarity based on the chunk information of its neighbors.

When a peer has no chunk at the beginning of the download, BitTorrent deviates from the rarest-first policy and the new peer requests a random chunk. This is intended to ensure a faster completion of the first chunk such that the upload bandwidth of that peer can be used by others. The default values in the original implementation are 256KB as chunk size and 16KB per requested block. To prevent that the sender runs out of requests and has to wait for a new request from another peer, the first requests after an unchoke are sent as a batch. By default the batch size is 5 requests. In normal mode a peer requests each part only once. This can become a problem at the end of the download. When the rest of the file is requested at a very slow peer, the downloading peer has to wait long although other peers may handle the request faster. Therefore, a peer can switch to the endgame mode, where it requests the same parts at multiple peers. Although, a peer can cancel requests at remote peers the endgame mode can consume additional bandwidth by transferring redundant parts.

All the involved mechanisms can be modified to follow the SmoothIT goals. *E.g.*, a peer can upload just enough to maximize the download and stay within the preferred list of as many peers as possible. Incentives for contribution of peers can be fostered by sophisticated sharing strategy variants of tit-for-tat that achieves co-operative behavior, similarly to repeated Prisoner's Dilemma. This may include to award seeds for uploading by allow them to discover better peers to download from whey they reciprocate. Again, topology promotion is a key issue. If the different roles, i.e. ISP, seed, leecher, and tracker, are given incentives to enhance locality awareness for example, then this could be considered as an *ETM mechanism*.

The BitTorrent protocol is neither standardized nor fixed and a large number of different applications, which use the BitTorrent protocol, are available. Especially, the implementation of the peer and the piece selection is implemented in different ways. Some clients additionally perform encryption making detection of BitTorrent traffic more complex. In this context, P2P NEXT [PN08] using basic Tribler [Tr08] have to be mentioned: "P2P NEXT: In a continued effort to support the development of P2P technology, the European Union has now invested $22 million in the development of an open-source Next-Generation BitTorrent client".

Tribler is an open source Peer-to-Peer client with various features for watching videos online and is available for Linux, Windows and Mac OS X. The first version of Tribler enhances the ABC client (Yet Another BitTorrent Client), the latest version 4.1.9 in the Tribler evolution is publicly available at [Tr08].

3.1.3 Gnutella

The Gnutella protocol still remains under development and many extensions exist for increasing its efficiency and robustness. There are several Gnutella clients available for free download, like BearShare, Morpheus, Gtk-Gnutella, Mutella, LimeWire, or Phex. This makes it even more difficult to describe "the" Gnutella applications. We focus therefore on the classical Gnutella version 0.4 as described in [Gn08]. It has to be noted that Gnutella is not so popular among users and is pushed away by eDonkey and BitTorrent.

The main difference of classical Gnutella to eDonkey and BitTorrent lies in the implementation of file lookup functionality and no use of MSD, however, in its current form Gnutella supports MSD by using a protocol named Partial File Sharing Protocol. Gnutella uses flooding mechanisms and time-to-live (TTL) counters to search for files and sources of files. In contrast to DHTs or centralized approaches, there is no guarantee that the file a peer is interested in is on any of the peers it can reach. Queries for files can take some time to get a complete response. It might be a minute or before all of the responses come with a typical TTL of seven hops. The exponential spread of request opens up the most likely source of disruption: denial-of-service attacks caused by flooding the system with request. As this leads to scalability problems, protocol extensions use super peers (also called Ultra-peers in this context) which build a hierarchical search overlay.

The file transfer is accomplished via HTTP. Downloading from a Gnutella host is technically equivalent to fetching a file from a web site. A major advantage of using HTTP is that two site can communicate even if one is behind a typical organization's firewall, assuming that this firewall allows traffic out to standard web servers on port 80.

Gnutella offers much optimization potential with respect to the traffic in the search overlay. In order to get the system scalable, DHT-based solutions for wide-area file search are proposed including mechanisms for flow control, dynamic topology adaptation and careful attention to node heterogeneity. In SmoothIT, schemes to optimize search traffic routing within the overlay can be tested. Furthermore, advanced network-aware schemes of search query expiration instead of simplistic TTL expiration schemes as used in classical Gnutella can be tested. Static and dynamically monitored information from this ISP on the network itself could be utilized.

3.1.4 Wuala

Wuala [Wu08] is different from file sharing applications, since it main functionality provides a free distributed online storage and backup based on P2P mechanisms. Additionally, files can be shared between friends or in groups of users. A right management is implemented to control for accessing files. All files and metadata are encrypted on the user's computer using AES 128bit, while for authentication purposes RSA 2048bit is used. All cryptographic operations are performed locally on the end-host which ensures that passwords never leave the local machines of a user. Anonymous publishing of files is not permitted.

Wuala is a hybrid P2P system with additional servers to support the system. These servers are owned and managed by the Wuala overlay provider. The business model behind this comprises online advertisements and additional services like photo prints. Power users can buy additional storage or backup solutions. Wuala itself offers licenses for its technology and provides an API.

The users of the system can take different roles. There are three different types of nodes, super nodes, storage nodes, and client nodes. Super nodes are responsible for routing. Storage nodes are used for storing (fragments of) contents and are connected to the closest super node. Client nodes have no responsibility and only consume data. To give the users incentives to participate in the system, the following rules regarding storage and download bandwidth are given. Users receive 1GB of initial storage. But they can get more online storage, if they provide local storage themselves. Users can trade local storage and gain online storage which is computed as follows: online storage of user = storage on local machine of this user * online time in % of this user. Users who are willing to provide storage and are online for at least 4 hours a day can get storage nodes. The download bandwidth of a user depends on its actually offered upload bandwidth using the fairness mechanisms Havelaar [GMS+06]. Its goal is a robust and efficient reputation system for active P2P systems. The more upload bandwidth a peer provides, the higher is its download speed.

Wuala is proprietary and offers its own mechanisms for resource mediation and resource access control. It is a proprietary P2P network based on a DHT which uses its own routing protocol. It is not a strictly structured overlay network, but uses links to direct neighbours plus random links learned by listening to messages piggy-backing routing information. The number of hops depends not only on the size of the network, but also on the amount of traffic, since regular messages contain network information.

Wuala allows a user to mark folders as public or private. In order to treat Wuala as a personal private storage space, a folder has to be marked private (which is default). Contents of a public folder can be accessed by anyone. For downloading public, shared contents, a BitTorrent-like mechanism for popular files is applied. Nodes build an overlay and download from other nodes that recently downloaded this file. The servers are installed for performance reasons to exchange file fragments, to provide meta data, and allow public search among others. For the transfer of data, an own transport protocol is developed which is based on UDP using port 7120 and is TCP-friendly. The transfer also allows streaming of video and music. This protocol (on the ISO/OSI transport layer) is used for file sharing, as well as for file storage or backup.

For the persistent storage of files, they are split into fragments and encoded by an erasure code (Reed-Solomon family). The fragments are uploaded to the P2P network and also to the Wuala servers. Thus, a high availability is achieved depending on the server storage capacities, as it is possible to retrieve the file from the servers without the P2P network. For example, 100 original fragments are encoded into 500, uploaded onto the storage node, and then 100 fragments are sufficient to reconstruct the file.

As optimization potential, a requesting peer could appropriately select nodes which store file fragments. Currently node selections are determined by respective hash keys. Locality and bandwidth can be used as selection criteria. However, from end-users' point of view bandwidth is more important than latency, since they normally deal with large files. Another option for optimization is reducing the usage of servers. The question arises how strong the performance can be improved. Regarding content delivery, servers are usually only used for missing fragments. In case of searching or providing meta data, servers are needed due to performance requirements.

3.1.5 Chordella

Chordella [ZDK08, ZHD+08] is different than the above introduced applications, as it provides a framework for efficient and robust resource mediation and lookup functionality. Therefore, Chordella can be used for any application to offer this functionality. However, the current proprietary implementation of Chordella is used for a file sharing application which runs in a local test bed. Chordella and also the related file sharing application are developed by DoCoMo Euro-Labs. However, the source might be open to the project partners on request in the context of SmoothIT.

As the name already suggest, Chordella is based on the Chord DHT. However, in order to cope with heterogeneous peers, especially in a mobile environment, super peers are selected which form the Chord ring itself. As super peer nodes with better capabilities are chosen. The other nodes are denoted as leaf nodes which connect to the closest super peers, but do not participate in the lookup itself. Only the super peers perform the lookup and store the required meta-information. To get the system work, they have to accept connections from leaf nodes. In this context, it is necessary to provide incentives to super peers which will be covered by the SmoothIT approach. Otherwise, they may misreport their available resources pretending to act as leaf nodes.

There exist already a lot of optimization mechanisms which are, however, not integrated yet into the current implementation. An algorithm is developed to select an optimal DHT configuration, which means a proper selection of the ratio between super peers and leaf nodes. The system dynamically checks the load and tunes the number of super peers as to keep their load as close as possible to the maximum acceptable. This way the total network traffic can be minimized.

Other optimization algorithms deal with caching and load balancing. The latter mechanism dynamically assigns leaf nodes to least loaded super peers in order to maintain load balance in the super peer overlay. The caching algorithm allows caching the queries along the query paths in order to further reduce the search traffic.

Further optimization potential includes the selection of super peers according to information from the ISP, *e.g.*, availability or bandwidth of users. This means that the DHT is fitted to the physical network in order to reduce the search traffic and improve the performance of the lookup system in terms of search delays. The hierarchy of Chordella and the above mentioned algorithm to select minimum DHT size help here, as the super peer network is much smaller than if the architecture is flat. An interesting observation is that an ISP might benefit from super peer selection because of incoming traffic from other ISPs while keeping the traffic in the own network.

3.1.6 Relevance to SmoothIT

In summary, the observations and considerations outlined that file sharing applications are of interest for SmoothIT due to the high optimization potential and many opportunities. In particular, open-source variants for the popular eDonkey and BitTorrent protocol can be used in internal and external trials, including the option to use Chordella for lookup. However, the file sharing applications differ significantly in their characteristics and implementations which will affect strongly the SmoothIT solution.

3.2 P2P VoD and Live TV

Due to the increasing demand for video services over the Internet, traditional client-server delivery seems to have reached a limit where it is becoming prohibitively expensive to serve more and more users. As such, P2P approaches for delivering video are gaining traction. These approaches offer rapid deployment and low const since they are not based exclusively on operator-owned infrastructure and resources. In this case, deciding on the overlay network architecture seems to be the main challenge [LRL+08]. Two different methods for pure P2P streaming have emerged thus far: mesh-based pull and tree-based push delivery.

In the *mesh-based approach*, a peer constructs a list of nearby peers, connects to some of them and receives the streaming media by requesting and combining stream parts from different peers ("pull"). As a result, the stream delivery is essentially a swarm of packets from different peers. At the same time, the peer shares the media with other peers that request stream parts. There is no clear relationship between peers – all of them are considered equal, apart from the peer that originally streams the content for the first time of course. This approach is based on a simple design principle that is inherently robust. It has seen a number of successful commercial deployments, like CoolStreaming and PPLive, but there are few open-source implementations.

In the *tree-based approach*, the peers are organized in a structure, such as a tree, and they have a well-defined relationship (parent-child). A peer in this structure forwards each media packet it receives to the child peers and as a result the stream of packets flows from the root peer to the leaf peers ("push"). This approach is not as robust because each node is a single point of failure for the branch that starts from that point (a critical characteristic for nodes high in the tree), it is not as adaptable to the very dynamic nature of a P2P overlay and it does not take advantage of the outgoing bandwidth of a majority of nodes, namely the leafs. However, this approach does not require sophisticated video encoding algorithms. Implementation examples for this approach include the End System Multicast (ESM) and PeerCast. As an improvement on this approach that combines elements from the previous method, peers might request media packets from their siblings on the tree and combine it with the stream they receive from their parent. In this way, the leaf bandwidth is utilized for improving the media distribution.

Some successful commercial implementations (like Joost) employ a *hybrid method*. This method is not purely P2P, as the overlay provider places a number of proxies to strategic locations on the network in the same fashion as a CDN (*e.g.*, Akamai). These proxies form a backbone network and perform the initial media distribution to peers. Then, the peers can use one of the other approaches to deliver the media among themselves. This approach tries to combine the best aspects of traditional and P2P media distribution.

P2P streaming applications are bandwidth-intensive. Currently, the state-of-the-art [SMG+07] applications feature media with approximately 300 Kbps bit rate for video streaming or 56-192 Kbps for audio streaming. Note that this bit rate is aggregate; this means that it is not necessarily delivered by a single peer. Start up time i.e. the time it takes from the selection of a channel or program until the user is able to watch video or listen to audio varies from about 20 seconds to a few minutes for pure P2P applications. Finally, the playout lag i.e. the time difference of the media between the source and the consumer is about one minute in the best case.

Table 1 gives an overview on the P2P VoD and Live TV applications discussed in this section. It shows the homepage from which the software can be downloaded, the popularity of the application, the offered service (VoD, Live TV, radio, file sharing) and the source code availability (GNU General Public License - GPL, BSD license, proprietary). BSD licenses represent a family of permissive free software licenses, i.e. free software licenses for a copyrighted work. In contrast, GNU GPL means "copyleft" licenses which require copies and derivatives of the source code to be made available on terms no more restrictive than those of the original license.

Table 1: Available software implementations for P2P VoD and Live TV

Name	URL	Popularity	Service	License
Vuze (Azureus)	http://www.vuze.com	Very popular	VoD, File sharing	GNU GPL
PeerCast	http://www.peercast.org	Limited	Live TV, Radio	GNU GPL
ESM	http://esm.cs.cmu.edu/	Very limited	Live TV, Radio	BSD
Freecast	http://www.freecast.org/	Very limited	Live TV, Radio	GNU GPL
Nodezilla	http://www.nodezilla.net/	Extremely limited	Live TV, Radio	GNU GPL
Joost	http://www.joost.com/	Very popular	VOD	Proprietary
PPLive	http://www.pplive.com/en/	Very popular (esp. in Asia)	Live TV, VoD	Proprietary
Zattoo	http://zattoo.com/	Popular	Live TV	Proprietary
SopCast	http://www.sopcast.com/	Fairly popular	Live TV, VoD	Proprietary

3.2.1 Vuze

Vuze (formerly Azureus) is an open-source BitTorrent client that also features limited VoD P2P streaming functionality over the Bittorrent protocol. During our tests, we found out that some video files on the *Vuze network*, a collection of torrent files distributed by Vuze Inc., sport a *Play* button (as well as the normal *Download* button). When the user clicks on *Play*, the video file starts downloading via the BitTorrent protocol. However, the chunks are not downloaded in random order (as with most BitTorent files) but from the beginning of the file to the end. Moveover, the Vuze client is capable of starting the playback of these files before they finish downloading. This qualifies as (limited) P2P video streaming.

Note that all the files that were streamable over the BitTorrent protocol were files tracked by Vuze trackers and were aided by Vuze servers, which functioned as super peers. However, many other (normal) peers provided the rest of the chunks of the file and those chunks were downloaded in sequence from the beginning of the file to the end as well. Normal torrent files (event files containing video) were not successfully streamed over the BitTorrent protocol: they had to be fully downloaded as ordinary torrents before they could be played. Only certain torrent files on the Vuze network were streamable.

The Vuze client is written in Java, and as a result it runs on most desktop operating systems. It features a very user-friendly GUI (especially for the VoD TV content), has plenty of configuration options available and it is very popular among end users. It also features a plug-in system that can be used by developers looking to augment the functionality of the client.

3.2.2 PeerCast

PeerCast is an open-source live TV and radio application that uses the tree topology. It is written in C++ and has been ported to Windows, Linux and Mac OS X. It has limited popularity and focuses mainly on radio content that the users create. As far as end-users are concerned, it has very simple channel selection but the configuration options are complicated. We were able to successfully test video and radio streaming.

3.2.3 End System Multicast

ESM is an open-source P2P streaming client from Carnegie Mellon University. It is a fairly new project and appears to have an active development team. It uses the tree topology and supports live TV and radio content that is usually provided by its users. It is written in C++, it runs on Windows and Mac OS X and it features a very user-friendly User Interface. It also appears to support third party development as it has limited developer documentation and resources.

3.2.4 Freecast

Freecast is an open-source Java-based P2P streaming client that supports live TV and radio. It uses the tree topology and has an extremely simple (almost Spartan) User Interface. It is designed so that it can run within a browser with a simple click. We were able to successfully radio streaming with this client.

3.2.5 Nodezilla

Nodezilla is an experimental, open-source client for grid networking that supports RTP (Real Time Protocol) streaming (among other features). With this module, it can play video (live TV) and audio (radio) content. It uses the tree topology and has a very complicated UI for end-users (and as a result, it is hardly ever used). It is also very configurable.

3.2.6 Joost

Joost is a commercial, proprietary application for VoD. It is one of the most popular and easy to use P2P streaming applications and it is available for Windows and Mac OS X. It uses a hybrid topology, combining distribution servers (like a CDN) and P2P mesh topology. The servers provide about 50% of the bandwidth. The client is free to use and the project is financed by commercials. The available content depends on the user's location (as there are a number of agreements between Joost and content providers). However, the client does not appear to be location or ISP aware.

3.2.7 PPLive

PPLive is a commercial, proprietary P2P TV application that uses mesh as the overlay topology. It features mainly Chinese content but it is user world-wide, sometimes in order to overcome license agreements of content providers. For example, Primera División matches are delivered from Spain to China (this could be an example where locality-aware optimizations might be applicable). It is more aggressive than other video applications. Measurements show that it uses 200 kbps upstream, 500 kbps downstream and about 5000 IP (Internet Protocol) addresses are contacted within 30 minutes of operation.

3.2.8 Zattoo

Zattoo is a proprietary application that features a pure live TV streaming architecture that broadcasts public TV channels. Service availability is location-dependant (and inferred from the IP address of the user). It uses TCP (Transmission Control Protocol) for video delivery and has similar architecture and traffic requirements with other applications.

3.2.9 SopCast

Sopcast is a proprietary application that uses server-assisted P2P multicast. It supports live TV. VOD development is at an early stage. SopCast limits the number of connections to 30, while in Joost up to 500 contacts were measured.

3.2.10 Relevance to SmoothIT

Usage and popularity of P2P live TV and video-on-demand applications is expected to increase in coming years. Thus the P2P streaming service type is considered as a future-proof application class. Moreover, compared to download completion time with file sharing, a visible performance improvement is easy to display and evaluate for QoE when using a SmoothIT optimization approach. Therefore, P2P video streaming offers even more optimization potential than file sharing, as topology promotion as well as QoS provisioning may be utilized as shown in Section 2.However, most popular applications are still proprietary and it might be hard to influence them to achieve a TripleWin situation.

3.3 P2P VoIP – Skype

The most popular overlay network for Voice-over-IP is Skype [Sk08] with more than 11 million concurrent users. The main attractions of the system are free Skype-to-Skype calls between its users and the high success rate in penetrating firewalls and NAT boxes. Another benefit of Skype is that it does not require any additional hardware, a simple PC with an Internet connection is enough. Based on the free basic service (plus chat and file transfer function) additional services are offered and charged, mainly Skype-to-PSTN and PSTN-to-Skype. Additionally, the software offers an API for third-party plugins, such as a shared white board and some simple games.

While Skype is a proprietary software it applies a hybrid approach: Centralized components are used for authentication, public key approval, billing and the connection to PSTN. The rest of the system is organized in a super-peer network. This means that the actual cost of communication is shifted to the users, and therefore, to the ISP networks too. This includes Skype internal calls, the forwarding of traffic close to the PSTN bridges and traffic relaying for NAT traversal, if required.

The whole Skype traffic is encrypted (using 256-bit AES and 1536-2048 bit RSA) and specific verification measures prevent the reverse engineering of Skype's program code [BS06]. Further, Skype provider claims that the system use a super-peer overlay which allows peers to communicate directly with each other without involving centralized components. A peer can be elected to become a super-peer if it is not hidden behind a NAT and has an upload bandwidth connection larger than 56 Kbps. These super-nodes are also responsible for relaying which requires about 60 Kbps per super-node in the median [GDJ06].

Another feature of Skype is the phone conference mode which is free for up to 10 users. In this case the conference initiator is responsible for receiving the voice streams from all hosts involved, mixing it and redistributing it again [BS06]. Further, Skype offers a simple video call function.

Skype bandwidth requirements are quite low [HB08], about 20-30 Kbps, and mostly UDP is used to transport the voice data while TCP is used for signaling [BS06]. Despite the low bandwidth requirements a typical Skype user will prefer a flat rate tariff in order to benefit from cheap Skype telephony. Therefore, for ISPs Skype usage is a selling argument for Internet flat rates. On the other hand, Skype competes with VoIP and Triple-play bundles offered by many ISPs as an additional service to their customers.

Depending on its own VoIP business ambitions, an ISP will be either interested in replacing Skype usage with its own service or optimizing the Skype traffic in the own network. Here, competing with Skype can be solved either by throttling Skype traffic (which is difficult due to the traffic being encrypted and to the network neutrality issues) or to offer a better VoIP service to ISP customers.

On the other hand, optimizing opportunities for Skype traffic include mostly traffic amount reduction. This could be done by offering ”real time“ connections to Skype users to increase the QoE, but the monetary opportunities are low because Skype users are used to the free service. Moreover, the possible traffic reduction for the inter-domain traffic is very limited because of the application's low bandwidth requirements. Other cooperation possibilities include an information API for an optimized super-peer selection, NAT relay selection and conference mixer election. But again the benefit to the ISP is limited.

Based on these considerations together with the fact that Skype uses a proprietary protocol and requires the access to the central login server, we conclude that this application is not suitable for the use in the internal or external trial in SmoothIT.

3.4 Content Distribution Network – Akamai

Akamai maintains the biggest commercial Content Delivery Network, claiming to have about 20.000 servers distributed over 71 countries and servers about 20% of the global

Internet traffic [Ak08]. The Akamai network is used by many major web enterprises to deliver large amount of content with a good user perception, i.e. mostly optimizing the end-user's latency. This is done by caching the customer's content close to the users and by redirecting the users' requests from the web sites to the caches running on Akamai servers. The content hosted mostly includes large objects found in web pages, such as images and videos. Therefore, Akamai allows customers to dynamically adjust the service capacity of their content. An important feature is Akamai's measurement service which is proved to be able to find almost optimal path depending on the current state of the Internet [SCK+06].

The user's requests are redirected via DNS redirection to Akamai servers. A hierarchical system of DNS servers is used to make the redirection decision based on the user's location, network topology, dynamic link characteristics, server load, bandwidth available and the content characteristics. The Akamai network is able to deal not only with failures inside of its own data centers but also with Internet wide path failures and congestions. Therefore, Akamai constitutes a resilient overlay network being able to cope with failures, bottlenecks and flash crowd events.

As Akamai moves the source of the content close to the requesting user the deployment of Akamai servers has a positive effect on the ISP network, the amount of inter-domain traffic is reduced. Hence, it seems that ISPs are willing to offer Akamai price reductions for bandwidth usage, while Akamai does not own a network by its own. So its main economical characteristic is the ability to work to the benefit of all three parties: content provider, user and ISP. The (global) optimization of the Akamai's network is done internally, however, additional state information from an ISP might help to further improve the network utilization and user perception.

Akamai shortly started to offer a streaming service to their clients. This additional feature and the fact that Akamai is able to provide service guarantees to their customers lead to the assumption that the network and the data amount served by it will grow further.

Akamai builds a central-managed global proprietary overlay network and it does not appear suitable for the SmoothIT approach and the succeeding trials. The reasons include the fact that Akamai hardly requires more cooperation with ISPs than it already has. On the other hand, Akamai is a paying customer to ISPs and therefore they cannot throttle its traffic. The only opportunity would be for ISPs to offer an additional service, *e.g.*, to prioritize traffic coming from Akamai servers. Finally, the opportunity for trials is very limited as it requires cooperation with Akamai.

3.5 VPN Used for Gaming: Hamachi

In this section, the Hamachi [Ha08] application will be analyzed as a very common application in the Internet world to play online emulating a private network over the Internet.

The Hamachi application/service provides two kinds of VPN services: free and premium. Both versions allow to set up a LAN over the Internet and arrange multiple computers into their own secure network, just as if they were connected by a physical cable with zero-configuration (the application works without having to adjust a firewall or router, it is able to work with almost all the NAT configurations) and in a secured way. This allows the end users to, *e.g.*, remote access to your home machines (with just Windows Remote Desktop or via Virtual Network Computing (VNC) programs, such as tightvnc).

The architecture used to achieve these features is based on a centrally managed VPN freeware: one server cluster that is managed by the vendor of the system and all the client software installed in the end-users computers (that will have a new network interface to the computer to intercept outbound and inbound traffic). Direct links between computers that are behind a NAT or firewalls can be established, resulting in a peer-to-peer based overlay network, supported by the server provided by the vendor of the system. As commented before, the software client creates a virtual network interface which captures the outbound and inbound traffic that is encrypted and authenticated. The traffic is then sent via a UDP connection. The amount of traffic sent and/or received will depend on the application used over the VPN (*e.g.*, for a gaming application, the amount of traffic is around 200 kbps since just updates must be sent).

In order to influence the overlay performance, it is important to analyze how the tunnels (connections between peers) are created and how data and control packets are sent.

VPN setup with Hamachi. From the user perspective the procedure to create and use a VPN is very simple since he/she only needs to download and install the Hamachi client software. To create a VPN, the end user just needs to provide a network name and a password. To join the VPN, the end users just need to provide the network name and the password. After that, a new network interface is created.
From the technical point of view, each client establishes and maintains a control connection to the server cluster. When the end users want to join a VPN, an authentication process starts; this authentication process is also used to infer the NAT characteristics. When a new member joins/leaves the VPN, the server asks the other members of the network (that is characterized by a name) to establish/release tunnels to the former.

For NAT traversal, Hamachi uses a technique similar to UDP hole punching but the technical characteristics have not been provided since this is “key business achievement”. A special procedure is also provided to maintain the status of the server and the clients in order to avoid transient networks problems.

Data transfer. The peers keep multiple upload/download connections to other peers and servers. The connections among peers are used for data transfer and the connections between each peer and servers are used for control traffic. The amount of data traffic depends on the application. In order to assure security and privacy, the traffic is encrypted (in fact, it is able to manage IP traffic).

Once we have analyzed the behavior of the application, the following points show the optimization potential:

1. The selection of the peers is deterministic since the overlay is composed of all the peers that want to be in contact each other. Therefore, an optimal peer selection will not be an optimization potential.
2. An important problem in current Hamachi service is its lack of reliability (the service is sometimes down). So a good optimization potential would be the fulfillment of interconnection agreements with the server cluster provided by Hamachi.
3. Taking into account that the end users aim to emulate a LAN scenario, they will be willing to use applications that requires LAN network performances in terms of delay, jitter and packet losses. Therefore, the provisioning of QoS guarantees to Hamachi tunnels for, *e.g.*, premium services could be an option.

Hamachi has not been selected as a relevant application, since it is not open source and SmoothIT could not influence on the overlay setup.

4 Classification Criteria for the Relevance of Overlay Applications for SmoothIT

The goal of this section is to introduce and explain the classification criteria which are used within SmoothIT. These criteria will be applied later on in Section 7 for the classification of the general overlay applications classes, like file sharing or video-on-demand (see Section 2), and the presented overlay application examples (see Section 3). We distinguish between major criteria and supplementary criteria regarding the selection process for the application which will be used for the further studies and work packages in SmoothIT. Additionally, we discuss the impact of the technical environments besides classical ISPs operating fixed networks which have to be taken into account when classifying overlay applications.

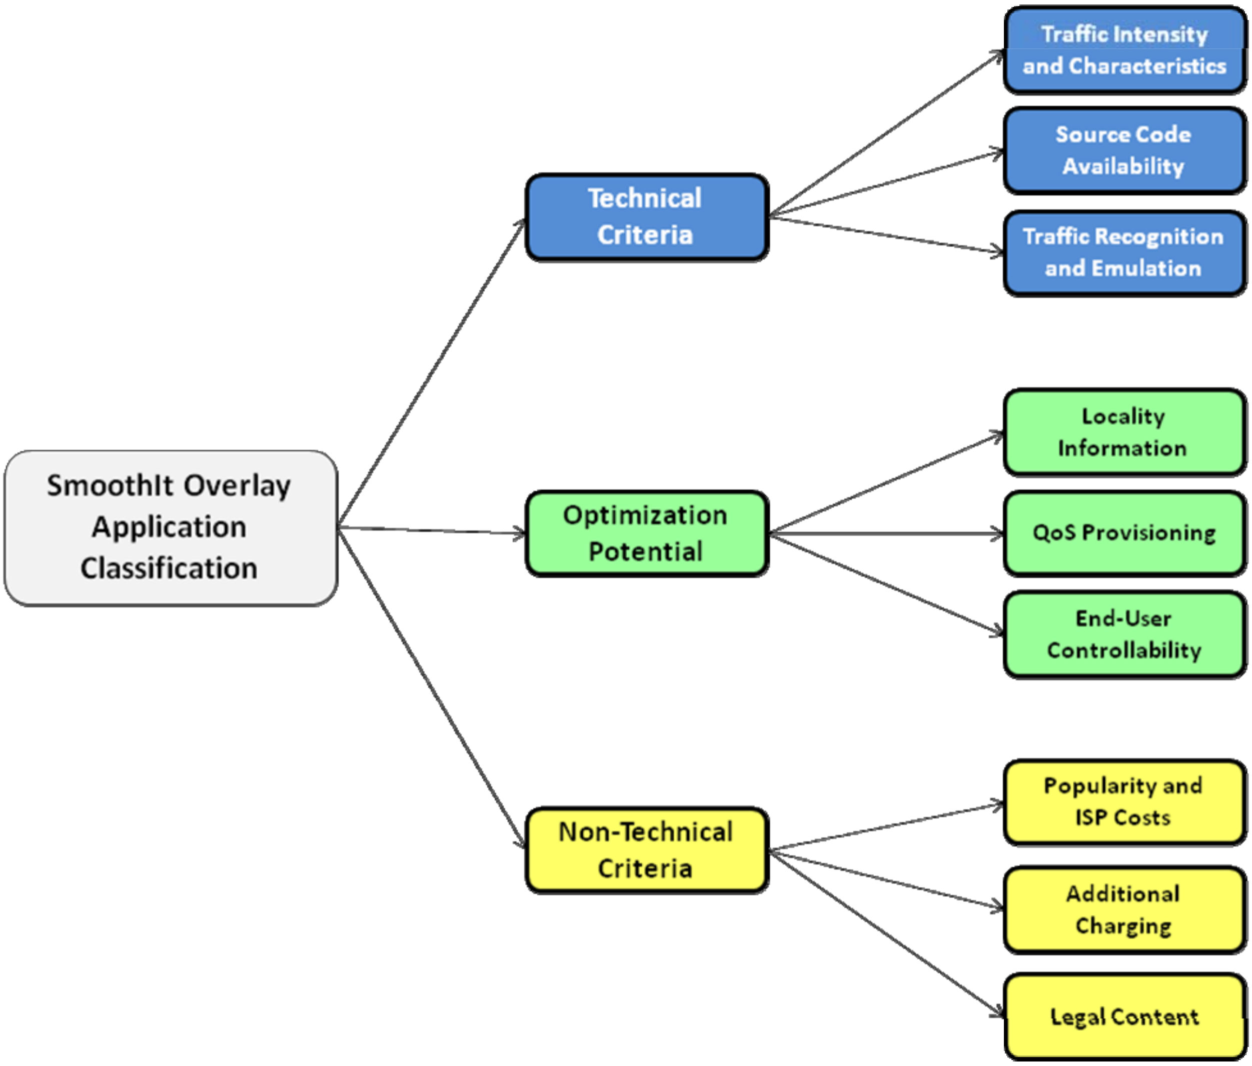

Figure 4.1 - Major classification criteria applied in SmoothIT

4.1 Major Classification Criteria Related to SmoothIT

From a technical point of view, SmoothIT aims at adapting the structure of the established overlay network such that the costs for ISPs due to inter-domain traffic are minimized, while the quality of experience of end-users is maximized. In order to achieve this goal, different optimization possibilities exist. Therefore, the applications have to be classified whether they can utilize the existing optimization potential. The actual solution and implementation of mechanisms utilizing the existing optimization potential of a particular application strongly depends, however, on technical criteria. These technical criteria include among others the traffic intensity and traffic characteristics, as well the source code availability which is especially important for the internal and external trials in SmoothIT. Since the ability to recognize traffic is important for the traffic characterization and traffic engineering, it also has an impact on the particular SmoothIT solution. The capability to emulate (by means of network traffic generators) and simulate the observed traffic patterns is crucial for the performance evaluation of the SmoohtIT approach including the internal trial as well as theoretical studies in WP2.

Additional non-technical criteria are of interest for the application selection process in SmoothIT, as the popularity of the overlay application and, thus, the expected costs for ISPs and telcos are crucial. The popularity of applications might also give a hint whether the user is willing to pay an extra fee for a better application's performance. Hence, the classification also takes into account whether such additional charging is applied for specific applications. In respect with the planned testing of the SmoothIT solution in an external trial, the legality of content has also been addressed when classifying the overlay applications. Figure 4.1 gives an overview on these major classification criteria applied in SmoothIT. These are a) technical criteria, b) optimization potential, and c) non-technical criteria which will be explained in more detail in the following.

4.1.1 Technical Criteria

The major technical criteria for the SmoothIT specific application selection process within in SmoothIT cover the source code availability**,** the traffic intensity and characteristics, as well as traffic recognition and emulation.

In order to have a full flexibility for implementing and testing the adaptation of the overlay to the underlay and to utilize information from the underlying network in the overlay (*e.g.*,the existing network topology or additional performance measures for particular links and nodes), it is required that the application is **open-source** allowing to modify the overlay application itself for our implementation and evaluation.

In case of a proprietary overlay application, we still see some options to use them for SmoothIT: a) some "indirect" interaction between the overlay network and the underlay network, *e.g.*, shaping the overlay traffic by the ISP such that the overlay application reacts accordingly and restructures the overlay; b) cooperation with the overlay application provider such that the proposed solutions and mechanisms derived in SmoothIT are implemented by the overlay provider; and c) reverse-engineering and re-implementation of the overlay application which of course would require a large amount of time.

In reality a) and b) would also lead to an intended change in the overlay application implementation by the application developers themselves, which is hard to trace and understand in detail. As the SmoothIT project should evaluate different overlay behavior and reactions to incentives, we prefer open-source applications instead of proprietary ones.

It has to be noted that the examined overlay applications presented in this deliverable are all freely available. Therefore, we distinguish only between proprietary and open-source applications. However, overlay applications might consist of several software pieces. An example is the eDonkey file sharing application. For eDonkey, proprietary software for the index servers exist, while the client application like eMule, being installed on the consuming end-user's machine, is open-source. This will be taken into account when evaluating the source code availability.

Another major technical criterion is **traffic intensity**. This includes the amount of totally generated traffic of the considered overlay application, especially the amount of inter-domain traffic as this will cause costs for ISPs. It has to be noted that the ISP's costs is separately considered as major non-technical criteria. Traffic intensity as a technical criteria addresses whether the resulting traffic from such an application is elastic or has a constant pattern, whether the traffic is continuous or bursty, or whether it is bandwidth demanding. Thus, the traffic intensity implicitly determines the solution approaches in SmoothIT and, hence, the achievable benefit.

As mentioned in the beginning of this section, the ability to recognize traffic is necessary for characterizing the traffic emerging from an overlay application. In addition, sophisticated traffic engineering mechanisms need to identify and recognize the traffic. For the application selection process in SmoothIT, the emulation of these observed traffic patterns is important, as the performance evaluation of the SmoothIT solution resembles these patterns. More details on **traffic recognition and emulation** in general are given in the following paragraph.

As an important part of OAM (Operation, Administration and Maintenance) mechanisms, operators usually carry out extensive measurements in order to collect and analyze the new traffic profile (*e.g.*, which are the most important contributors to the global traffic, the traffic streams characteristics in terms of duration, inter-arrival time of the packets, etc.).

One common used methodology to characterize the traffic is by means of identifying the ports used. Usually each application uses a specific port, so it could be easy to process, *e.g.*, NetFlow traces to extract the amount of traffic associated to one port and indeed to one application.

This methodology has an important inconvenience to really identify the traffic associated to P2P file sharing applications, since users behind firewalls usually change the port used by the P2P file sharing applications to traverse the firewalls or to address possible ISP filtering. Therefore, after the identification of the ports, it is important to analyze other traffic patterns, such as the symmetry of the traffic or inter-arrival packet time, in order to avoid, *e.g.*, the classification of P2P File Sharing traffic into Web traffic.

4.1.2 Optimization Potential

The possible gain which SmoothIT can reach for a particular overlay application is expressed in the optimization potential. While for the ISP it is important to reduce its costs, *e.g.*, by reducing the inter-domain traffic, the user's interest is reflected in its current QoE. Therefore, an overlay application looks promising for SmoothIT if it is able to utilize topology information and QoS provisioning. The control knobs for realizing this are taken into account in what we call the end-user controllability of an overlay application.

The **end-user controllability** shows how the end-user can influence the performance of an overlay application. There exists a continuous graduation on the different possibilities ordered according to the degree of controllability of the overlay application: 1) the adjustment of parameters within the application software, 2) the utilization of incentive mechanisms, 3) the selection of particular nodes in the overlay, and 4) the modification of the source code of the application. The highest value indicates the highest controllability. An example for the adjustment of parameters for an overlay application is the number of parallel uploading peers of a file sharing application using multi-source download. This parameter has a major impact on the performance of the entire system [SHT06]. In a heterogeneous environment like Beyond 3G systems, the performance of the system can be drastically improved when dynamically setting this number according to the available bandwidth resources of a user [DHS07]. This might be done by changing the source code and releasing a new software version or by some additional script running in the background which monitors the available resources and modifies appropriately the parameter saved in an XML file for eMule.

The selection of nodes as possibility for end-user controllability can also be illustrated on the example of eDonkey. Here, the end-user chooses the index server to which it connects to the eDonkey network. As the index server knows all peers connected to it and thus the available files and sources for these files, the selection of an appropriate index server results in an improved performance. The reason for this lies a) in the clustering of users into groups of interest and b) smaller delays for “near” index servers. For example, an index server run in Italy attracts Italian eDonkey users, as the near index server answers file requests faster and as there might be information and online documentation available also in Italian language. As a result, more Italian content is registered at this index server and a user interested in this content will connect to that index server. In general, this selection of nodes influences the emerging overlay topology and inherently the user perceived quality of that service.

The next major classification criterion is the **utilization of locality information** by an overlay application. In this context, two questions have to be answered for each overlay application. Can the overlay application itself benefit from available locality information? Can the ISP or telco reduce costs when providing this information? This shows whether locality promotion can be utilized to improve the QoE while reducing the costs for ISPs due to traffic crossing the ISP’s borders. For a video streaming application the utilization of locality information is possible, as it does not matter for the end-user from whom it downloads the data, as long as the available throughput is large enough and the perceived end-to-end delays and jitter values are small. However, for a VoIP application, this is only partially possible, as the voice call has to be established between the caller and the callee.

In the context of optimization potential, the **utilization of QoS provisioning** is the next key issue. Overlay applications which are able to exploit QoS mechanisms will provide a better QoE to its users. Although an ISP cannot reduce costs when offering the technical environment to guarantee a certain QoS, its customers will be satisfied leading to a win/win situation (cf. Section 5).

The degree to which QoS provisioning is beneficial depends on the QoE sensitivity of the application. As regards the sensitivity of QoE evaluation the applications can be classified more or less the same as presented in chapter 2 of D1.1. The most sensitive applications would be voice and video applications. Users are usually very critical while assessing such applications since video, image and voice quality perception is very important. QoE must be evaluated during the trials in case of video-on-demand, voice-over-IP as well as on-line gaming applications. A QoE evaluation would be also interesting in the case of VPNs. People are usually less critical in assessment of file sharing or CDN applications. *E.g.*, they are usually likely to tolerate long waiting time for the content download. In the following, more details are given on the utilization of QoS provisioning and QoE sensitivity.

QoE Sensitivity and Utilization of QoS Provisioning

Customers assess the quality of application by its perceivable performance. The performance of all applications depends on QoS and QoR, which is network condition and parameters. Also technical aspects of the application itself, including codecs, algorithms, overlay architecture, etc., are important. However, the final user assessment usually depends on several additional factors not necessarily connected with technical aspects such as psychological, sociological and environmental conditions, pricing policy, content etc. In general, a very broad set of parameters and conditions affecting the final QoE evaluation can be established. In case of a particular application concrete parameters can be defined but it would be only a subset of the whole broad set. The subset would differ not only between the application classes but between similar applications. For example, two similar applications of which one is free while the other not may receive different ratings in terms of QoE.

Analysis of QoE and sensitivity of users' assessment of the applications' quality is very important for the network operator. Although the operator is not a service provider, the mechanisms implemented in its network domain influence the user perceived applications' quality. QoE may be either improved or degraded. Finally, the end user will assess the network provider basing on QoE of the applications he uses. Thus, the additional criteria encompassing QoE sensitiveness of applications is important for SmoothIT. Since QoE is, among the other things, influenced by QoS and QoR parameters, the connections between those parameters should be established, and their inter-relations should be analysed.

The QoS parameters reflecting the network condition and traffic characteristics are commonly known. The most often mentioned in the context of application class sensitivity to them are: bandwidth, delay, jitter, packet loss probability. Additional parameters that may affect the user perceived quality are: blocking probability (if the call admission control is considered), changing the packet order, effect of a "forced traffic locality". The last parameter is new in the context of ETM mechanisms used by operator to force the overlay application to reduce the inter-domain traffic.

In the literature [CMH+07a] the following group of QoR (Quality of Resilience) parameters (features of a network that affect the QoS observed by the users, and are related to resilience) is enumerated:

- Reliability attributes (attributes adopted from the classical reliability theory): continuity, downtime, availability.
- Recovery-related features (features of communication networks that without being identical with any of the reliability attributes, still strongly influence them): quality of the recovery path, affected traffic (traffic lost due to failures), resilience to multiple failures, preemption sensitivity, fault coverage.

All of them have a good rationale for connection-oriented networks where a service is related to a well-defined communications chain (i.e. where we can easily determine working and alternative path). However, in the context of overlay networks the situation is not the same. Thus, we claim that only a subset of those parameters/features can be used to effectively describe the overlay behavior in failure-related situations:

- availability of selected resources – the probability of finding the resources (*e.g.*, the enough number of file chunks in a population of overlay nodes) in an operating state at any time we want their service (*e.g.*, download the file, taking into account a situation that a population of overlay nodes can be faulty or can fail during the file download),
- affected traffic – the amount of traffic lost or disturbed due to failures in inter-nodal paths or due to node faults (an accumulated unfinished work,

- resilience to multiple failures – flexibility to connectivity changes generated by multiple simultaneous faults,
- fault coverage – the fraction of traffic or connections which is recovered in a given failure scenario.

On the other hand, it is important to be aware of the QoE parameters that can be affected and are significant to application rating. For voice applications (such as IP telephony) the QoE parameters are, but not limited to: communication quality, presence of echo, interruptions/silence, cracks. In case of video applications (including video conferencing or video streaming) the perceived quality may be assessed with respect to: image distortions (visible blocks, slices, blur, smearing), non-continuous streaming, picture freezing. The most important QoE parameter depending on network performance in the case of on-line gaming can be called interactivity. In turn, the file sharing download may be assessed in terms of download speed/duration, availability of content, probability of unsuccessful download.

In general, QoE evaluation by the end user is significant in the case of video and voice applications. The potential bad effect of any mechanisms (including ETM) affecting the network performance and, what follows, influencing the perceived quality of application, would be easily noticed by the user. People are usually more critical while assessing those applications. A separate case is on-line gaming. QoE depend on the type of the game in this case. Games usually do not require high bandwidth. However, the action games require the information to be delivered from one peer to the other in a very strict time regime and without packet loss. File sharing application seems to be less sensitive to QoE assessment. Users usually can tolerate longer waiting time for the operation to complete. But the situation may change of they pay for fast download.

The above general description of applications' QoE sensitivity is not sufficient to take a decision on the selection of application for trials. More detailed analysis of the application is needed. Several non-technical aspects such as pricing policy, target group of users etc. must be taken into account.

The following questions should be answered while classifying the overlay applications with respect to QoE issues:

- What are the perceivable characteristics of the application that contribute to the final QoE assessment by the user?
- How sensitive is the QoE evaluation of the application to the changing network conditions?
- What QoS/QoR parameters is the application sensitive to?
- Can we obtain the information of QoS/QoR parameters and any network conditions meaningful for the QoE evaluation of the application?
- Is it possible to build a QoE model for the application that allows calculation of MOS as a derivative of technical parameters of the network and traffic that affects the QoE?

4.1.3 Non-Technical Criteria

After discussing the technical criteria and the optimization potential for the classification of overlay applications with respect to SmoothIT, we finally consider non-technical criteria which have a major impact on the overlay application selection. The **popularity** of an overlay application determines the impact on the underlying networks and the emerging traffic volume. In this context, it might be necessary to include a country differentiation, as some applications might be more spread in some countries than in others. For example, eDonkey is the dominant file sharing application in France, while in US BitTorrent is mainly used. The popularity of an application is also important for the external trial in order to have a broad user basis for demonstrating the SmoothIT approach. Another facet of popularity covers the willingness of users to pay an extra fee for improved QoS and QoE of an application. In addition, **additional charging** might be accepted for specific applications. The user's interest are important to consider in the selection criteria, as the popularity only reflects the current situation, while the additional charging allows to get a glimpse on future popularity. In order to quantify the popularity of overlay applications and the willingness for additional charging, AGH proposes a user survey for this investigation. The user survey handed out to the test person is evaluated in Section 7.1.

Strongly correlated to the popularity and the traffic intensity of an application are the **ISP costs** which is one of the key objectives in SmoothIT to reduce them. The ISP costs also include CAPEX as well as OPEX. CAPEX include costs for buying and upgrading network infrastructure while OPEX include costs for operation and maintenance of the ISP's network. The traffic intensity affects the ISP costs as augmented intra-domain traffic might result in higher CAPEX and OPEX (details are given in Section 5). Therefore, we explicitly consider this criterion in the application selection process.

Another factor which has to be taken into account, especially for the external trial, are copyright protected contents. In order to respect copyrights, supporting mechanisms have to be implemented. If, however, there is only legal data available and exchanged among the entities of the overlay network, no special arrangements are needed. Therefore, a major non-technical criterion discusses **legal contents**.

Finally, **other opportunities for trials** have to be considered for the decision which overlay application will be further investigated during the course of the SmoothIT project. They include among others if there are any other projects for liaisons, *e.g.*, integrating the SmoothIT approach in an existing test bed or offering a basic overlay applications which SmoothIT adapts to its own needs. In this context, it is also important whether open-source variants for a particular overlay application exist which can be used in the trials to mimic the application behavior.

In addition to the economic view in Section 5, we briefly take a closer look at the additional charging. In a most common case users of peer-to-peer applications are charged by ISPs only for internet access (IA subscription fee). In order to compensate the producers and providers of the content different payment schemes are considered, *e.g.*, subscription based peer-to-peer networks, pay per access, or pay per view.

Different charging models for the content are considered. An example of approach, taken into account is legalization of peer-to-peer sharing by enforcing a mandatory monthly fee on all Internet users. The revenues would be split according to the popularity of downloads. The proposed revenue sharing model was drafted by the Songwriters Association of Canada. Another approach was proposed by EMI Music. This world's largest independent music company has agreed to make its entire directory available through P2P service (Mashboxx). This solution would allow to preview and buy legal content from within existing P2P networks.

Clients would have the ability to perform search for the content. The results would be offered to the user who can either evaluate a sample for free or pay and download the required content. In the case of the sampler option users could play the full length track up to a predefined number of times. According to NPD Digital Music study, approximately 75 % of peer-to-peer users consider the try-before-you-buy option crucial before making purchase decision.

Sony BMG also considers making its catalogue available through Mashboxx as soon as the service is operating. The related issue is the payment methods that would be feasible and appropriate for particular cases: a) micro payments – the order of a fraction of Euro, b) macro payments - payments which are large enough amounts as to be processed via credit card payments or other traditional payment instruments. The question that appears in this context is if and how these approaches could be applied to a differentiated content exchanged in peer-to-peer networks.

4.2 Supplementary Classification Criteria

In the previous section, the major criteria for the classification and the selection of an overlay application are explained and discussed. There are, however, much more supplementary classification criteria. Although, they will have an impact on the particular algorithms and solution approaches for reaching SmoothIT's objectives, they are less important for the classification targeting the application selection than the criteria introduced in Section 4.1. The main idea of the classification is to use as little selection criteria as possible while still providing all necessary information related to the SmoothIT project. The following list shows example supplementary classification criteria and shortly explains whether already considered implicitly in major criteria or why skipped:

- Traffic requirements, like bandwidth, delay, or jitter. This is implicitly considered in the optimization potential 'utilization of QoS provisioning', as well as in the technical criteria 'traffic intensity'.
- Traffic characteristic: TCP/UDP (User Datagram Protocol), ratio signaling to content, bandwidth consumption, traffic patterns, etc. This is considered again in 'traffic intensity'.
- Type of overlay architecture for signaling and/or for data exchange like central/distributed/hybrid. This is covered in the 'utilization of locality information'.
- Overlay topology: unstructured, small-world property, structured, ring, mesh, tree, forest, etc. This is addressed by the three 'optimization potential' criteria.
- Sensitivity with respect to network neutrality. This is covered in the discussion of economic and regulatory view in Section5.
- 'Business vs Residential' addressed the question whether some applications are specific for certain type of users (*e.g*., only for residential, or both). This is partially covered by the 'popularity' of applications.

At present almost all peer-to-peer applications are targeted at residential users. However some applications, *e.g*., conversational services (voice and video communications), can be considered for business clients as well. Features specific for business users comprise: increased security, easy deployment to multiple machines in the client's company, extended control features for IT administrators.

For example Skype provides a Business Control Panel which is a free web-based tool that allows controlling user's Skype Credit expenditures. Centralized management of employees' credit balances is available. This enables flexible credit allocation. Automatic recharge can be set up in case when an employee's account falls below a set threshold. Monitoring and reporting functions are possible as well.

Apart from features making the application directly applicable for business usage, they can have attributes and functionality for providing 'value added' business services. Again, an example can be Skype Prime™ which can be used for providing charged consulting services.

4.3 Influence of Mobile Environments

Due to the present growth of 3G networks (3^{rd} Generation) and accompanying data transfer, mobile operators can be considered a special kind of ISPs. The networks run by mobile operators have two important properties that are not normally found in the fixed networks (or do not have such severe impact). These are **heterogeneity** and **mobility**. They have to be taken into consideration when classifying the listed applications.

To make these points clearer, consider heterogeneity as an example. Some nodes in a typical cellular network have extremely small upload capacity, which is at the same time extremely expensive. So it makes sense to let these nodes only download content. The key question is now how easily can the proposed applications be adapted to take this fact into account. A more general question is also interesting and relevant: how easily can they be changed to accommodate various models in which costs are associated with upload capacities of the involved nodes. As a first conclusion, it is impossible to make such extensions without being able to modify the code, i.e., **heterogeneity puts strong emphasis on source code-openness**.

Locations and mobility patterns of the participating nodes play an important role as well. For example, if the employed solution to reduce the total traffic in the underlay network implies deploying operator-owned peer nodes to help content distribution, then mobility becomes important in making decision on where exactly to deploy these nodes. At the same time, **mobility** may have impact on upload capacity of a mobile node, so it is **unclear what changes are needed to account for this fact** when making efficient self-organization strategies (such as BitTorrent’s tit-for-tat and choking strategy).

However, our main focus in the project is “classical” ISPs operating fixed networks, in which the problems stemming from mobility and heterogeneity are not so severe that we would need to push them to test-bed and trials. On the other hand, mobile operators and their **(future) problems cannot be neglected** and this is why we want to **address them in the project as well**.

5 Economic and Regulatory View

The purpose of this section is to describe all the effects (positive or negative) and new opportunities that overlay applications may introduce to the operator, the end-user and the overlay provider. We analyze the different objectives for each stakeholder and discuss what the appropriate incentives are for the stakeholders so that they make good use of the overlay application, leading to a win situation for all. Finally, we conclude by mentioning the (direct or indirect) “conflicts” that such actions may raise, along with some remarks on the possibility for further optimization.

5.1 Observations

The emergence of high-speed access networks has shifted the use of Internet into a new era: new applications have emerged that are bandwidth intensive, require some quality of service levels to be met, and have several other quality requirements. On the other hand, users and applications try to find ways to overcome the network problems that they often come up with, *e.g.*, delays or use of congested paths. Overlay applications are an example of such applications, where the user-defined (or application-defined) routing affects the way traffic flows in the Internet.

The most popular applications up to now are file sharing applications that serve as a new medium for content distribution. Recently, new types of overlay applications appear and gain popularity, such as P2P voice and video applications. However, due to the high volume of data that circulates due to P2P file sharing applications, we will examine the effect of those applications in the business relationships between ISPs. We expect that the new types of overlay applications will have the same effects, especially as the volume of the generated traffic increases.

A very common observation is that, due to overlay traffic, ISPs have witnessed a change in their traffic patterns. As discussed later on, such changes affect their business relationships. On the other hand, overlay applications offer new opportunities to end users. A conflict between ISPs and overlay providers arises, each one trying to acquire the most benefits for him. Hence, it is necessary to provide the correct incentives to all the stakeholders, so that we reach a situation where all the conflicts are resolved.

Since the central notion of this section is the incentives for each stakeholder, we have to define the term as considered within the SmoothIT project:

> *An incentive determines a monetary or non-monetary factor which provides a motivation for a particular course of action or counts as a reason for preferring one choice to another.*

5.2 Incentives for End-users

Overlay applications have emerged when the need for applications to control their traffic became an objective in order to offer new features that are not offered by the network such as multicast, file sharing, etc. From a user’s perspective, overlay applications were initially offering a better-than-best-effort service above Internet’s best effort communication services.

As overlay application technologies advanced and became more popular, it was noted that such systems were based on users' altruistic behavior, i.e. sharing content without any direct gains. At the same time, many users started exploiting altruistic peers, by not offering content to the 'community' (free-riding). These phenomena dictated the need of rules for the user of overlay applications to comply with, along with incentives so that users would make a good use of the overlay services. In this section we summarize the most important incentives for the end-users.

5.2.1 Performance Improvements

As already mentioned, one of the most attractive incentive to be offered to the end-users is the performance improvements introduced when using the correct mechanisms. By doing so, users are expected to experience better performance of P2P applications. Due to an enhanced peer selection process, possibly assisted by the ISP, higher bandwidth and/or lower transmission delay can be achieved. Higher bandwidth is especially relevant for applications transferring large amount of data, *e.g.*, file sharing applications. Low transmission delay is essential for real-time applications, *e.g.*, telephony or live streaming, since high delay can lead to discernible performance degradation. Moreover, efficient incentive mechanisms can also reduce packet loss rate, since they can decrease link load, resulting in a less congested network. Such mechanisms result in better quality of experience (QoE) for users and elevate perceived quality, since users experience shorter download time in file sharing applications and experience better quality in real-time applications.

5.2.2 Availability of Peers

Peer selection is a process in which a P2P application chooses from which peer to request a service, considering the same service is offered by a variety of peers. SmoothIT intends to introduce mechanisms which can provide enhanced information about peers, for example concerning availability and locality, which can be used by the P2P application to perform a more efficient peer selection. Based on this information, P2P applications can select peers that remain online longer as routing nodes for data transmission, which improves the stability of the overall P2P network. Another factor that plays an important role is locality of peers. Such a mechanism may deliver locality information to P2P applications that may be used to improve peer selection. For the stability of a P2P network an important aspect is path diversity. Diverse paths can be achieved by selecting peers that are located at different places in the network. The advantage in deploying different paths is an improved reliability. A failure in one connection may have a minimized effect on the whole P2P network, since there are redundant paths that can be used as alternatives. Another advantage is the fact that a peer node can download from several other nodes at the same time over independent data paths, which reduces his download time considerably.

Additionally, locality information enables that nodes located near the requesting node are preferably chosen over very distant ones to provide a service. The reasons for handling requests this way are manifold; for example, download times can be reduced when transferring data from a near-located node and the network as a whole is less probable to get congested due to the fact that nodes distributed across the whole network are less likely to connect to each other.

However, there is a trade-off concerning locality and path diversity. Path diversity implies that different paths are being chosen, which lies in contradiction to the selection of near-located peers. This means that peer selection cannot only consider peers located close to a peer node and having a high capacity link but shall select also peers concerning path diversity even if these peers are far away.

In general, better routing can be achieved by applying location information. Nodes close to each other in the network are not necessarily close as well in the overlay topology. With information about the location of the nodes, faster routes can be created by selections peers that are near to each other in the underlying network topology.

5.2.3 Reputation and Trustworthiness of Peers

For several overlay systems, information about the reputation and trustworthiness of peers is very important since it dictates how mechanisms deal with them. Whether a peer is trustworthy or not is decided on its previous behavior. This depends on two factors, first related to the content and second related to the peer's capabilities. Every peer is expected to comply with the protocol and contribute to the network as much as it consumes. Peers with a greater amount of resources are more capable to support the network than peers that share only little of their resources. Therefore, such peers are crucial to the overall performance of the network and because of their supportive function they are rated higher than others. P2P applications may be able to choose peers that have a high reputation and connect to these peers rather than connecting to peers that cannot fully be trusted. Proper incentives should be given to peers so that they actively participate in the overlay community.

5.2.4 Monetary Benefits

In some cases, incentives of economic value may be given to peers in order to join and participate in an overlay 'community'. Such incentives may be in the form of lower overall price, since the ISP will have lower cost, charging a lower price especially for traffic circulating inside the ISP's domain, or in terms of discounts for lower level services.

5.3 Incentives for Overlay Providers

As already mentioned, due to the large volumes of data exchanged in the overlay networks, ISPs see their traffic patterns changing drastically and in an unpredicted manner. Due to the conflicts that arise (see next section for more details), overlay traffic is sometimes treated as unwanted traffic. In order to avoid such tussles that may deteriorate the QoS of overlay services, overlay providers should be given incentives to adapt their traffic to some of the underlay requirements.

5.3.1 Monetary Benefits

Collaboration with ISPs might lead to some economic gains for the overlay provider. For example, some providers might deploy their own peers in the ISP's premises so that they act as super-peers and improve the overlay service provisioning. Cost for network traffic can, in this case, be decreased due to the locality preference. In the case traffic is charged differently for intra-domain and extra-domain, fees from the ISP to the overlay provider will decrease. Intra-domain traffic might be favored, which results in lower costs for both ISPs and application providers.

5.3.2 Performance Improvements and Better Service for End-users

Another incentive, even though indirect, for the application providers is the better service performance that they can offer to their customers. Bandwidth increase and lower delay times are enhancements that augment the received performance for consumers.

5.3.3 Reputation and Loyalty of Users

The main concern of an overlay provider is to keep and increase his customer base. Mechanisms can be deployed and proper incentives for using them can be provided that address this issue. Due to the large number of overlay applications that provide almost the same services, providers want to differentiate in some way. Better performance of the overlay network leads to an increase in customer satisfaction which itself is crucial for building up a strong loyalty on the side of the user.

5.4 Incentives for Operators

The role of the Operator is very crucial to an environment where underlay and overlay networks have conflicting interests. The operator, being the owner of the underlay, sees his traffic following an unpredictable pattern when overlay applications are active. One not so desired option for the ISP is to disregard and dump overlay traffic. In order to avoid such extreme actions, it should hold that ISPs and Overlay Providers collaborate so that they both attain some benefit.

It is important thus to find the appropriate mechanisms and incentives for both sides to cooperate in order to reach a stable point of operation. The user-driven routing technologies render fragile the business relationships in the interconnection market. The need to rationalize the economic interests of the ISPs with the desires of the users is a crucial issue to be considered by any incentive mechanism. SmoothIT’s vision is to provide the framework that enables information to flow between overlay and underlay in both directions, allowing mechanisms in both layers to consider multiple cross-layer criteria for its decisions. Such mechanisms can either require some changes in overlay and underlay protocols or can be fully transparent.

5.4.1 Monetary Benefits

The driving force for Internet Service Providers to use new technologies, offer new services and expand their business relationships is the maximization of profits, achieved either by the increase of revenues or the decrease of costs. Hence, incentives with a positive economic impact are also applicable here, as in the case of end-users and overlay providers. In the case of ISPs though, this kind of incentive has much stronger impact.

5.4.2 Traffic Management

Regarding traffic management, ISPs can profit from less congested links. Links with a tendency to congestion can be detected and avoided on time. Data volumes can be shifted towards links that are free or under-utilized. In addition, data traffic can be predicted much easier.

5.4.3 Performance Improvement and Better Service for End-users

The aforementioned advantages for users, namely higher bandwidth and lower delay, hold in the case of ISPs as well. Customers who experience better service quality are more likely to keep a contract with an ISP that provides good service.

5.4.4 Reputation and Loyalty of Users

This is especially interesting in terms of marketing reasons, since ISPs want to attract a great customer group by delivering better performance than other Internet Service Providers.

5.5 Transport Aspects of Overlay Traffic

In order to understand better how overlays affect the operators' network traffic, we briefly describe the structure of an ISP's network and the type of interconnections between ISPs.

5.5.1 Intra-domain Transport

Currently, a very high percentage of Internet traffic comes from overlay applications. For example, according to a recent study about the Internet traffic distribution in Germany [IS07], around 74% of total Internet traffic in Germany is generated by P2P applications as shown in Figure 5.1.

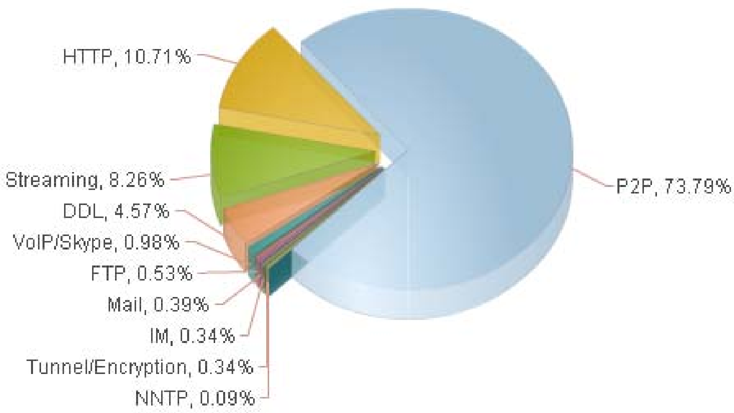

Figure 5.1 – Protocol Type Distribution in Germany, 2007

The impact of this huge amount of traffic on the network costs strongly depends on its distribution. For example, as shown in Figure 5.2, if an ISP customer is exchanging P2P traffic with a customer of another ISP, then such traffic is consuming resources in the whole network: access, aggregation, core and interconnection.

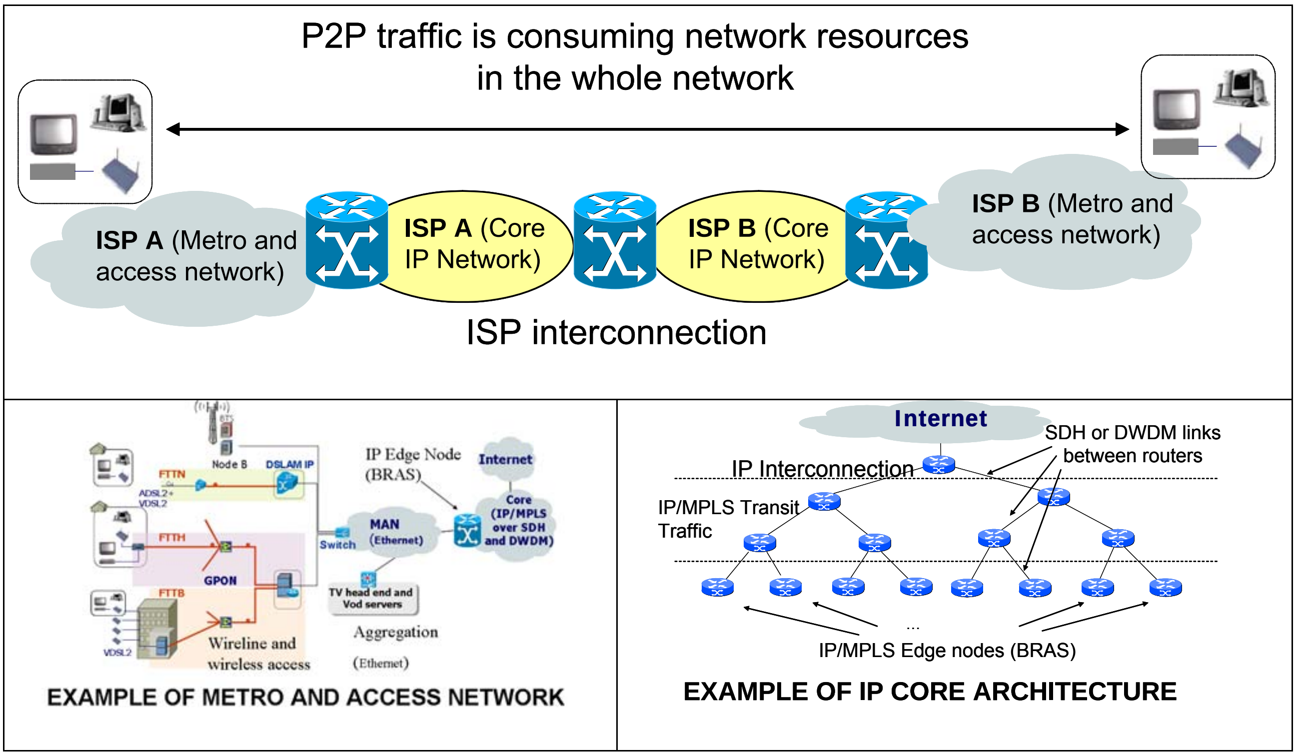

Figure 5.2 - Network architecture for overlay traffic transport

The end to end path followed by an inter-carrier overlay flow starts in the access network. End users of overlay applications (peers) can use either wire-line (such as fiber, cable, xDSL (Digital Subscriber Line) or wireless (*e.g.*, WIMAX, UMTS [Universal Mobile Telecommunication System], GPRS [General Packet Radio Service]) access connections to the Internet. Traffic from multiple end users is aggregated in the operator's access edge node (such as DSLAMs [DSL Access Multiplexer], GPON OLTs (Gigabit Passive Optical Network Optical Line Termination), UMTS Node B), so that aggregated traffic flows from multiple access nodes are transported over Layer 2 networks (*e.g.*, Ethernet) towards the IP edge router (BRAS, Broadband Remote access Aggregation Router) which inspects users' packets in order to check their destination address. Afterwards, IP packets are aggregated in MPLS (Multi Protocol Label Switching) 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.

5.5.2 Inter-domain transport

ISPs' networks are interconnected as autonomous routing domains. Global routing and network reachability among these Autonomous Systems (AS) is managed by eBGP (external Border Gateway Protocol) as illustrated in Figure 5.3.

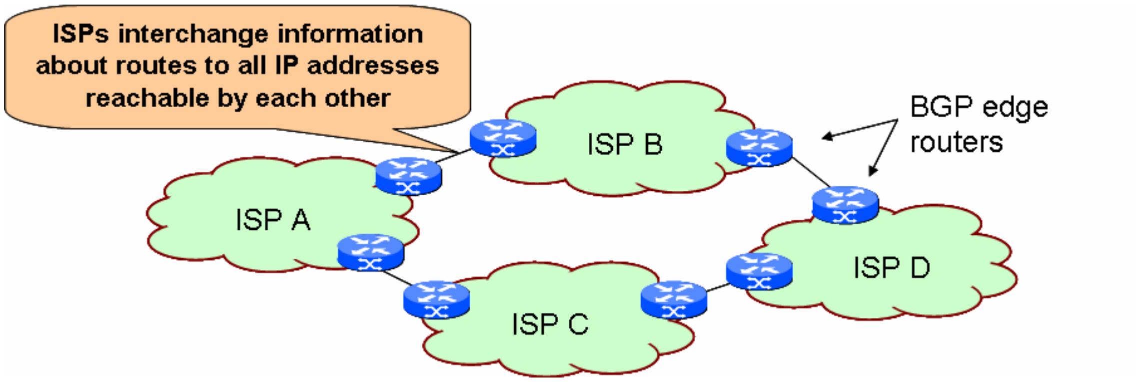

Figure 5.3 – ISPs interconnections

The eBGP is a routing protocol used on the edge of Autonomous Systems (AS). It calculates loop-free (or direct) paths across the Internet by tracking the path in terms of which AS it passes through. However, it does not track the "route" through individual routers within an AS. To use eBGP, an operator must have a router that supports BGP (Border Gateway Protocol) and a registered public AS number. Routes learned via BGP use associated properties to determine the best route to a destination. These properties are referred to as BGP attributes, and are used in the route selection process.

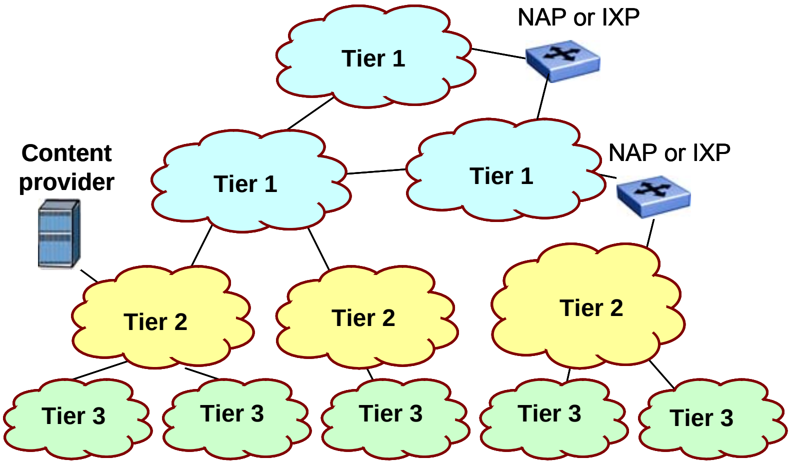

Figure 5.4 - The hierarchy of the Internet

As there are so many networks comprising the Internet, it is impractical for networks to directly interconnect with each other. As a result, larger networks may choose to offer a "transit service". This is a service for the delivery of packets across a network to IP addresses which that network can "see". A transit service provider configures the routing table of the BGP router to advertise the IP addresses of the network to which it is interconnected. In particular, this hierarchy is described by reference to tiers of operators and is depicted in Figure 5.4:

- Tier 1 ISPs (sometimes known as "backbone operators") are large telecommunications operators which have significant numbers of points of presence (PoPs) and do not use transit providers.
- Tier 2 ISPs usually have some network of their own, although limited to a geographic region (*e.g.*, a European country) and rely on purchasing some level of transit from Tier 1 ISPs to exchange messages with out of region networks and content providers.

- Tier 3 ISPs are purely re-sellers of internet access services and purchase transit from Tier 2 ISPs.

The physical interconnections between two ISPs are categorized into two types, cf. Figure 5.5:

- **Public interconnection:** Interconnection utilizing a multi-party shared switch fabric such as an Ethernet switch. Public interconnections are typically done in multi-party shared locations called Neutral Access Points (NAP) or Internet Exchange Points (IXP)
- **Private interconnection:** Interconnection utilizing a point-to-point interconnection such as a patch-cable or dark fiber between two parties. Private interconnections can be done either between individual carrier-owned facilities or at carrier neutral collocation facilities
-

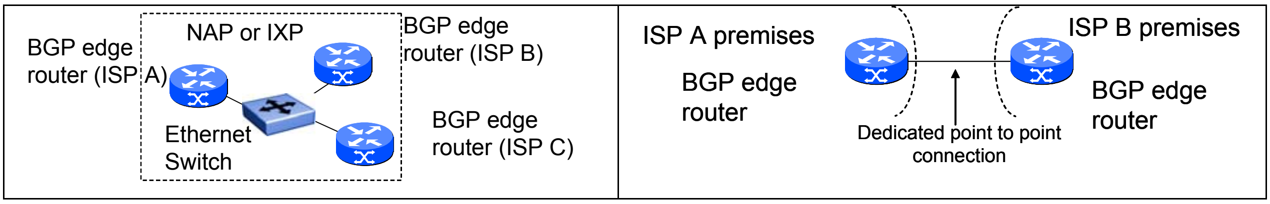

Figure 5.5 - ISPs interconnections- Interconnections-Physical connectivity

5.6 Overlay-related Costs

Operators' costs generated by multi-domain overlay traffic can be divided in to three categories:

- Investments on the internal IP core network
- Investments international links
- Transit costs

This section aims to describe the rationale behind these costs.

5.6.1 Investments on the Internal IP Core Network

As mentioned in Section 5.5, overlay applications are consuming a huge amount of bandwidth in operators' access, metro and core networks. Actually, if we translate such bandwidth into network resources consumption (*e.g.*, network nodes switching and transmission capacity) then we can easily realize that both the amount and distribution of overlay traffic is strongly impacting on total network costs, namely CAPEX and OPEX. CAPEX include costs for buying and upgrading network infrastructure while OPEX include costs for operation and maintenance of the ISP's network.

In particular, according to TID calculations, the required investment (CAPEX) in internal IP core routers is typically increased in **15-20 K€ per Gbps**[1]. In fact, a high percentage of this cost (around a 74%) is related to overlay applications.

[1] Estimation based on public Cisco prices

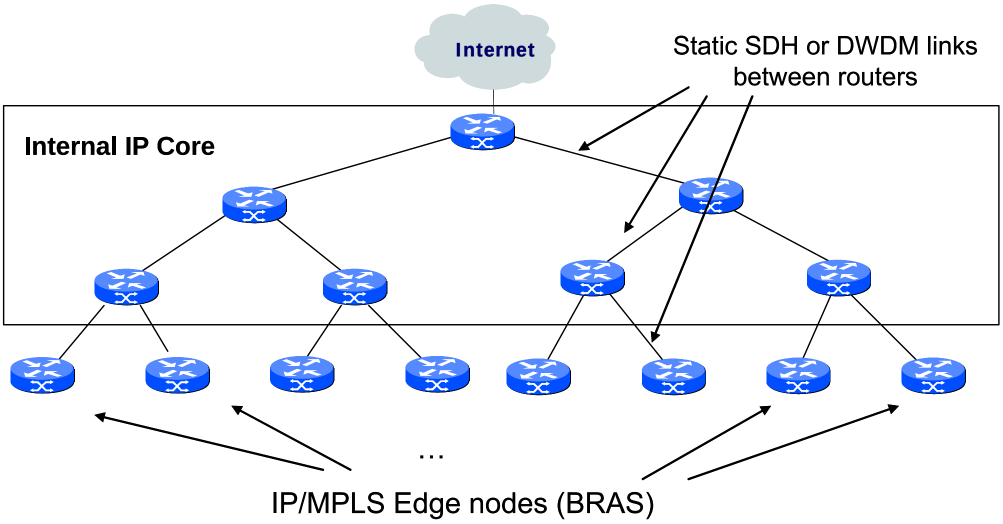

Figure 5.6 - Internal IP core

5.6.2 Investments on inter-carrier infrastructure

Local/national ISPs often need to use international links for multi-ISP interconnection purposes (see section 5.5). This situation is depicted in Figure 5.7. An Internet exchange point (IXP) is a physical infrastructure that allows different Internet service providers to exchange Internet traffic between their AS. A typical IXP consists of one or more network switches, to which each of the participating ISPs connect.

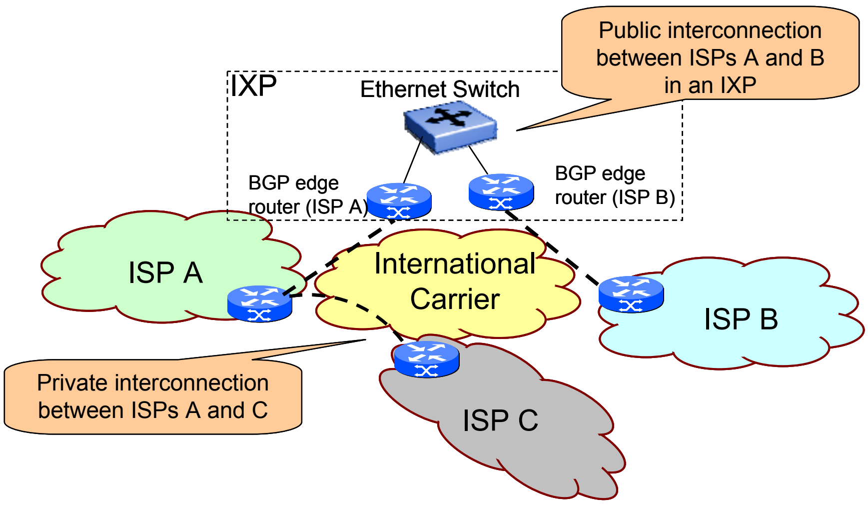

Figure 5.7 - Inter-carrier Infrastructure

This infrastructure often belongs to international carriers: Level 3, TIWS (Telefonica International Whole Sale Services), etc. Therefore, local/national ISPs (*e.g.*, in Cyprus or Spain) should pay to these carriers for the international links. According to it, if the inter-domain traffic increases, then local/national ISPs should buy more capacity (*e.g.*, STM-64 POS [Packet-over-SONET] links over DWDM [Dense Wavelength Division Multiplexing]). For example, Telefonica of Spain (local ISP-Tier2) periodically buys new connections (*e.g.*, at 10 Gbps or 40 Gbps) between Spain and the Internet Local Exchanges in USA, to international carriers.

However, it is important to highlight that the required investments on international infrastructure strongly depend on each ISP situation.

5.6.3 Transit Costs

As described in section 5.5.2, the ISP interconnection consists of the advertisement between ISPs of routes to their customer's IP addresses. Such exchange of reachability information and traffic could be done freely or not depending on the ISPs business relationship. The relationships between ISPs are generally described by one of the following categories:

- **Peering agreement:** Two ISPs (*e.g.*, two Tier 1s) agree on exchanging traffic between each other without any cost.
- **Transit agreement:** An ISP (*e.g.*, a Tier 2) pays to another ISP (Tier 1) for the traffic exchange

Transit agreements are typically done between Tier1-Tier2 and between Tier2-Tier3.

Interconnection between Tier 1 and Tier 2 ISPs: Tier2s have to pay to the Tier1s 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 95th percentile method. According to this method, the net imbalance 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 thrown away, and that next measurement becomes the billable utilization for the month.

Interconnection between Tier 2 and Tier 3 ISPs: The Tier 3 ISP will always pay for data it downloads from the Tier 2 ISP. Data flowing from the Tier 2 ISP to the Tier 3 ISP is always likely to exceed traffic flowing in the other direction, given that the Tier 3 ISP has no content to host, and merely sends retail customer requests for data or applications. For this reason, a Tier 3 ISP has no opportunity to offset charges for any data it uploads to the Tier 2 ISP, against its charges for downloading from the Tier 2 ISP. However, in the traditional internet environment, the Tier 3 ISP usually recovers the entire retail charge for the service from the retail customer.

5.6.4 PrimeTel’s Perspective

Within the SmoothIT project, we have decided that the test bed for the trials will be PrimeTel’s network. We therefore briefly present PrimeTel’s network, business relationships and strategies, which we think are very important for the identification of the correct incentives that would be the most appropriate.

PrimeTel’s case introduces some particularities since it operates in Cyprus, which is a small island national with population of less than 1 million. The amount of traffic generated by local content providers is very small compared to external sources on the Internet, and cost of international traffic is very high because due to small population the island does not require as much international bandwidth but still needs to pay the same or higher costs for submarine fibber optic cables, and the amount of competition on international bandwidth is lower because size of the market does not allow for many players.

The end result is that 98% of Internet traffic in PrimeTel's network is international, and that one-off cost per 1 gigabit of international connectivity on IRU basis is several millions Euro. At the same time, the cost of domestic (intra-Cyprus) backbone network is quite low due to the small geographic area & population covered by the network.

5.7 Interconnection Economics

As already mentioned, ISPs make bilateral agreements on the volumes of traffic exchanged between them. Such agreements are not regulated, strongly depend on negotiations between ISPs and are either of free exchange of data (peering agreements) or volume-charged (transit agreements). ISPs base their interconnection agreements on the patterns of the traffic exchanged. As a result of such agreements, ISPs have the obligation to police their traffic in order to conform to them. Traffic engineering is used by ISPs for traffic policing and shaping.

5.7.1 Overlays and Interconnection Agreements

The emergence of overlay applications directly affects the policing of traffic and makes it more difficult for the providers to change the way outgoing traffic flows, i.e. how traffic is routed. Overlay applications introduce user-driven routing technologies and the effects are twofold: they change the routing of P2P traffic in order to meet some quality requirements and they alter the volume of traffic exchanged between ISPs, as well as the destination for some routes.

It is obvious that agreements between ISPs do not remain unaffected by the existence of overlay applications. Such changes in routing behavior can introduce economic gains or losses to ISPs, depending on the type of agreements they have made and how the flow of traffic affects the flow of money. This conflict between ISPs and overlay networks is known as “tussle”.

In peer-to-peer (P2P) networks, nodes act simultaneously as clients and servers. Due to the various mechanisms and rules, demand symmetry is introduced, since nodes are given the incentive to balance demand with supply. Hence, ISPs who exchange P2P traffic can make peering agreements, due to this demand symmetry. But this might not be the only possibility, as shown in the examples to follow.

In [AW04] the authors present some examples of how overlay traffic affects the interconnection relationships of various ISPs. We will provide some short description of those examples so as to clarify the tussle in the core of the Internet due to the existence of overlay applications.

5.7.2 Tussle – Examples

Consider four ISPs and their interconnections, as depicted in Figure 5.8. ISP A has two transit agreements with ISPs B and C, defined by a unit price of €50 per Mbps and €100 per Mbps respectively. ISP A chooses to reach ISP D through ISP B, in order to face lower transit costs. Assume that customers of ISP A use overlay application. Due to various quality metrics and peer placements, the overlay routing protocol may dictate that flows originating from an overlay node in ISP A to an overlay node in ISP D, should travel through an overlay node located in ISP C. This will incur an increase of costs for ISP A, which now has three alternatives: either to not interfere with overlay traffic or try to re-route overlay traffic (if possible) through ISP B or totally block outgoing overlay traffic. In any of the above cases, either customers or ISPs are not content with the situation.

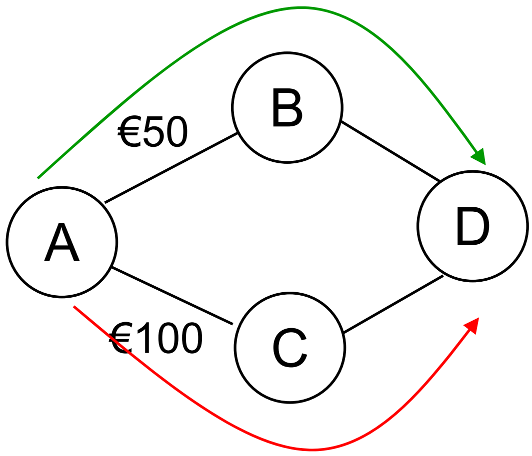

Figure 5.8 - Example 1: Transit agreements

5.7.3 The Case of CDNs – Stakeholders & Relationships

CDNs are an example of an overlay network that is widely used to bring seamlessly web content closer to the end-users. As already mentioned in previous sections, Akamai is one of the leader companies in this market. Based on Akamai's example, the work of [F07] defines the stakeholders of the market and focus on the monetary/business relationships that exist and how these relationships are affected by the existence of overlay networks.

On the one side we have *Content Users (CUs)* and on the other side we have *Content Providers (CPs)*. CPs determine customers of *Content Delivery Networks (CDNs)* and can be distinguished in "cost sensitive/delay insensitive" and "cost insensitive/delay sensitive", depending on which is the nature of their business and type of content. Furthermore, from the underlay network side, we have the *Internet Service Providers (ISPs)* and *Internet Backbone Providers (IBPs)*.

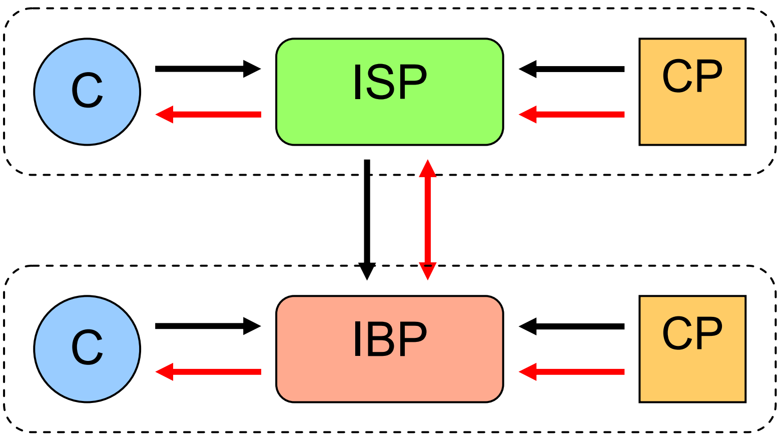

Figure 5.9 - Data (red) and money (black) flows between stakeholders (without a CDN)

Without the presence of CDNs (see Fig. 5.11), content users and providers pay Internet providers (either ISPs or IBPs) for accessing the Internet and either reaching or offering content. ISPs pay IBPs in order to connect to the core Internet.

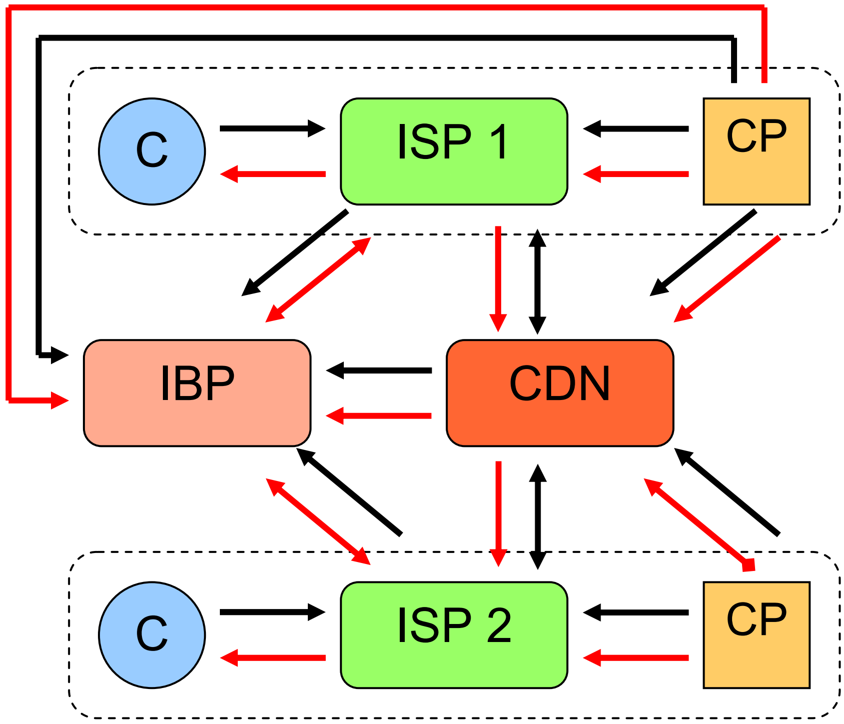

Figure 5.10 - Data (red) and money (black) flows between stakeholders (with a CDN)

When CDNs are present, new flows of data and money occur. Content providers can either distribute their content through ISPs, CDNs or both, depending on their type, i.e. if they are cost sensitive or delay sensitive. In all cases though, providers pay ISPs, CDNs or both. IBPs treat CDNs like any other network with large asymmetric outbound traffic characteristics. A transit agreement therefore is made between them. CDNs may pay or not ISPs depending if they are collocated or CDNs are accessible by the Internet through ISPs. IBPs continue to charge ISPs, as mentioned in the previous case, and ISPs continue to charge end users. Thus, a case of double billing arises where IBPs charge CUs, CPs and CDNs for the same service (either directly or indirectly), since both CPs and CUs are charged due to IBPs' bargaining power.

It becomes hence obvious that the existence of an overlay network changes the structure of the business model and the flow of data. Even if the example of proprietary CDNs is not the case for SmoothIT, it is still interesting to study the issues that arise since it might be that other overlay applications have the same impact on the interconnection market.

5.7.4 Conclusions

We have displayed how interconnection costs can be affected by the presence of overlay traffic. This observation provides the main motivation to search for and design mechanisms that consider all the economic implications and resolve all the conflicts that may arise due to decisions taken at the underlay and overlay layers.

5.8 Locality-awareness

According to the contents of the previous sections, we can conclude that, as higher the amount of inter-carrier overlay traffic as higher the total network costs. This is due to two main reasons: on the one hand, inter-carrier traffic flows should pass through the whole network and consumes more transmission and switching resources than internal traffic. On the other hand, in case of having an IP transit agreement, then inter-carrier traffic should be paid to another ISP. However, as shown in Figure 5.13, intra-carrier overlay traffic doesn't consume interconnection bandwidth.

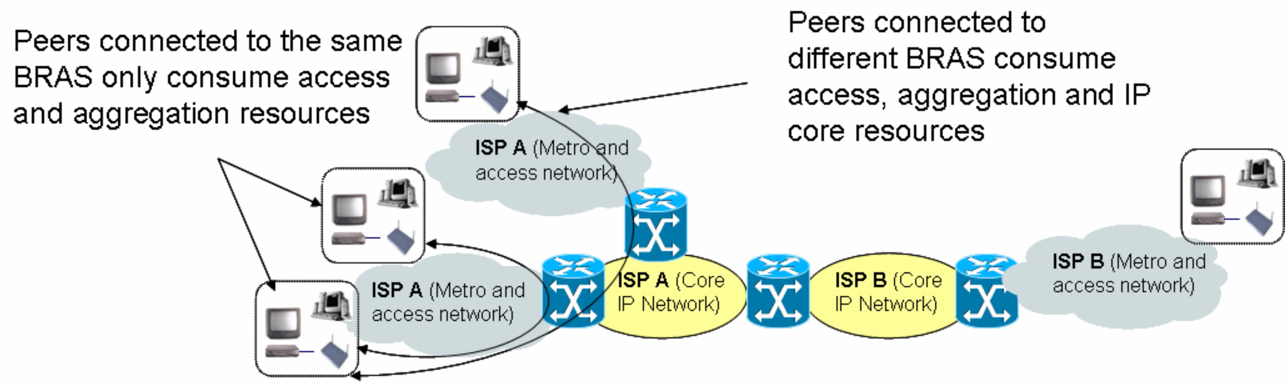

Figure 5.11 - ISPs interconnections

Therefore, the promotion of overlay traffic locality is a key issue under an operator's perspective since it may reduce both network investments and transit costs.

In the case of PrimeTel's network, the company interested to maximize the amount of domestic traffic, and minimize the amount of international traffic, due to its particularities as well. Implementing an effective scheme which would allow stimulating overlays to retain as much traffic locally as possible would result in clear economic benefits to PrimeTel, which can be shared back with customers by means of a proper incentive mechanism which stimulates them to generate traffic locally.

Video-on-demand is recently becoming one of the top growers among various types of Internet traffic and is causing concern among ISPs which have to upgrade their networks to be able to cope with the pace of growth and to be able to continue providing high QoE to end user - compared to file transfer, real-time video-on-demand streaming is much more demanding to network QoS, requiring stable throughput and low packet loss. An overlay which would be capable to reduce the amount of times the same video content is fetched from outside the domain can bring very considerable effects on reduction of transport costs end-to-end. It should be noted that in case of video overlays, relationship between overlay players may be based on a more organized Producer-Transporter-Consumer top-down approach, not ad-hoc Consumer-Consumer as in most of file sharing overlays. Both Producer and Transporter will clearly benefit from reduced content delivery costs and ability to offer additional services to end-user. This makes it interesting to consider a generic Producer/Transporter-independent overlay solution, where any of the parties would be able to setup a respective node and use the common pool of overlay nodes for delivering of content. At the same time, use of traditional Consumer-Consumer overlay designs for delivering of streaming video content does not seem to be very scalable due to generally low upstream rates available to Consumers.

5.9 New Business Opportunities

As explained in the previous section, ISPs could minimize their network costs by increasing the percentage of local overlay traffic. However, this might not be the only economic advantage of overlay traffic management.

According to some IP traffic forecasts, Internet streaming applications are expected to play an important role in the short term. In fact, as shown in the next figure, a 35% of total

Internet traffic in 2010 is expected to be generated by real-time applications (*e.g.*, streaming, gaming, VoIP, or videoconferencing) [CIS08].

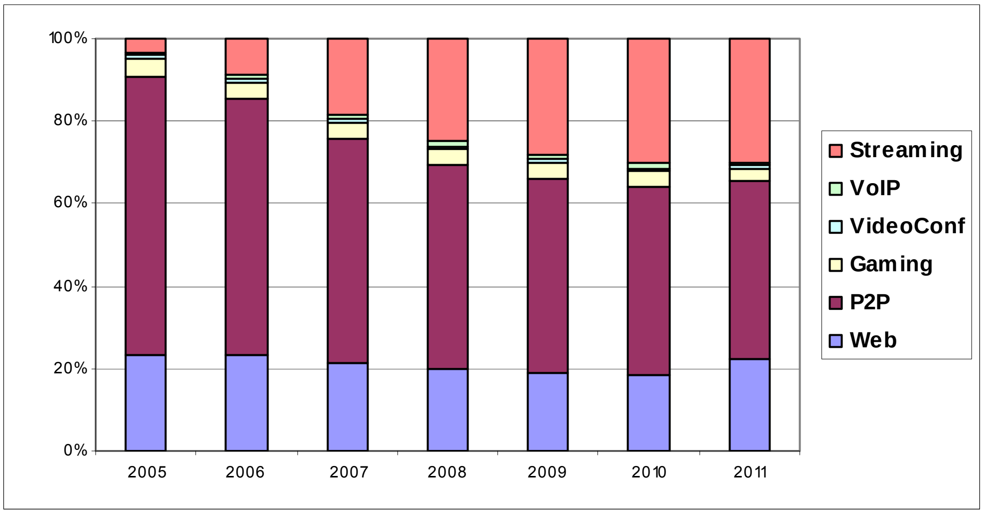

Figure 5.12 - Evolution of Internet traffic distribution by application [CIS08]

Therefore, a new potential business opportunity for ISPs would be related to QoS differentiation aspects. Currently, Internet traffic (*e.g.*, HTTP [Hypertext Transfer Protocol] or overlay) is transported according to a "best effort" approach. However, some overlay applications such as IP-TV, VoD, VoIP, videoconference or gaming present strict requirements in terms of delay and packet loss.

	Bandwidth Mbps						
Type of service	peak down	peak up	mean down	mean up	Max delay	Max jitter	Packet loss r
Video Broadcast 0 (Mobility TV)	0,384	0	0,256	0	< 2 s	< 40 ms	< 3 E-3
Video Broadcast 1 (SDTV mpeg2)	6	0	6	0	< 2 s	< 40 ms	< 3 E-3
Video Broadcast 2 (SDTV mpeg4)	3	0	3	0	< 2 s	< 40 ms	< 3 E-3
Video Broadcast 3 (HDTV mpeg2)	20	0	20	0	< 2 s	< 40 ms	< 3 E-3
Video Broadcast 4 (HDTV mpeg4)	10	0	10	0	< 2 s	< 40 ms	< 3 E-3
VoIP	0,008	0,008	0,008	0,008	< 70 ms	< 20 ms	< 3 E-3
Grid computing	0,512	0,128	0,0358	0,009	< 200 ms	< 50 ms	< 1 E-4
Gaming 1 Mobility	0,04	0,04	0,03	0,03	< 50 ms	< 10 ms	< 5 E-2
Gaming 2	0,25	0,25	0,20	0,20	< 50 ms	< 10 ms	< 5 E-2
Videoconference 1 Mobility	0,03	0,03	0,026	0,026	< 100 ms	< 10 ms	< 3 E-3
Videoconference 2	0,128	0,128	0,1	0,1	< 100 ms	< 10 ms	< 3 E-3
Videoconference 3	3	3	2,2	2,2	< 100 ms	< 10 ms	< 3 E-3

Figure 5.13 - Typical bandwidth and QoS requirements of real time and streaming application

Recent studies carried out in the NOBEL-2 project evaluated the quality degradation of multimedia stream caused by interruptions [NOB]. These studies, which were based on a subjective measurement system, showed that the quality degradation depends basically on the packet loss, but besides this it is also depends on the distribution and length of the outages.

Therefore, the QoE perceived by the end users of real time or streaming overlay applications, such as Skype, PPlive, or Joost, could be significantly improved by introducing application-aware transport services able to provide low packet loss rates and required jitter performance. Developing and implementing such application-aware traffic management mechanisms would require new network investments. However, operators could also be indirectly profited by improving the QoE of overlay applications since they could increase the broadband customers' fidelity as well as selling new broadband connectivity services specially adapted to Internet real-time and streaming applications.

Figure 5.16 shows an example of a potential commercial offer based on QoS differentiation.

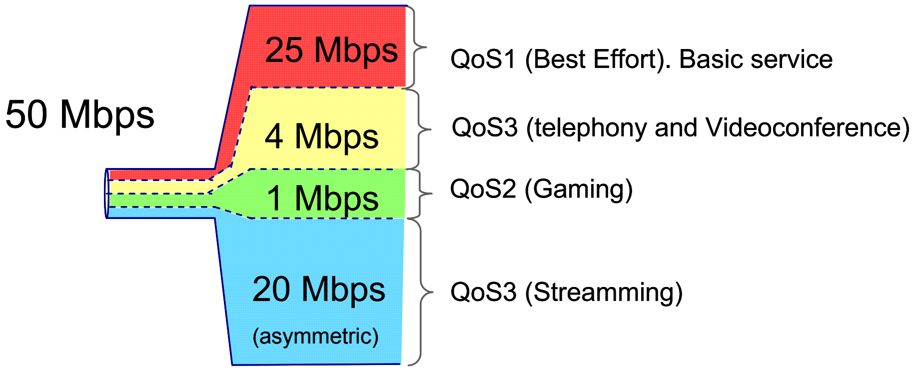

Figure 5.14 - Different connectivity services over a broadband access connection

According to it, ISPs could offer connectivity services especially adapted to streaming, videoconference or gaming applications. The QoS differentiation could be charged by adding an extra fee to the regular Internet flat rate.

5.10 Network Neutrality Issues

Games, voice and video applications, IP television services or even file sharing need QoS assurances to operate properly. The values of traffic parameters, like packet delay or packet loss have to be minimized. QoS requirements may be ensured by providing traffic policies by ISPs or by the well known mechanisms like IntServ (Integrated Services) or DiffServ (Differentiated Services). In fact, in many cases, ISPs have to block packets of some applications to ensure proper transmission parameters of some other traffic. Such an attitude discriminates selected applications which may be perceived as unfair. Moreover, it is possible to guarantee different quality of the services based on, *e.g.*, source or destination addresses or network device ports. ISPs may exert this possibility to prioritize some network applications, therefore, assuring better Quality-of-Service (QoS) to the selected traffic.

Network neutrality (sometimes referred to as net neutrality or NN) represents the legal concept which, for example, forbids such behavior. The scope of the discussion about NN, however, is huge and the presented examples cover just one aspect of the issue. More specifically, we present the main idea of NN and the issues that arise, so that we consider them adequately while designing mechanisms that offer the correct incentives to operators, overlay providers and end-users.

It has to be noted that this section is based on the debate in the US. In Europe, this kind of scenarios are not under discussion. Therefore, the term telco or ISP used in this Section 5.10 refers to the main actors of the net neutrality discussion in US.

5.10.1 Net Neutrality – Overview

In general, the idea of net neutrality is that a user's traffic is not discriminated at all in relation to a traffic generated by other network users. The legal conditions of the net neutrality are discussed all over the world, including the United States Congress. Even though the first legislative approach failed, the debate still continues.

Although there are many views on net neutrality, the main three perspectives can be summarized as follows:

- strict NN (Network Neutrality) – no QoS allowed whatsoever,
- reasonable NN – allowed QoS, but restricted,
- no NN regulations/law – just as today.

The last option is forced by the telecom companies and ISPs, who, obviously, do not want any new restrictions [C06]. Their interest is directly connected with the money earning capabilities (given that, the lobbying process to maintain the current status quo is strong). Telecom operators argue that they have no incentives to develop the networks, if direct revenues are not possible. Such an approach, as the history taught us, is questionable [O08]. The reason for that has the nature of a ripple effect. Telcos would like to extra charge content providers, like YouTube for their extensive bandwidth consumption, which might seem fair at a first glance. However, the upbringing of broadband services (or bandwidth consuming for that matter), in fact, increased the telcos' revenues, as the customers started to migrate towards broadband Internet access.

Strict NN is proposed mainly by non-technical parties who simply state that, currently, Internet works sufficiently well and there is no need to change anything. This extremely rigid version of NN has little chance of success, as the necessity to provide QoS in the networks is undeniable.

The third option, the consensus between two extremes, presents how, most likely, net neutrality will look like. Service differentiation will be allowed but it must remain fair. For the operator, it is fair to prioritize all VoIP calls within the network, however focusing on a certain application (like Skype or Vonage) violates the rules. Consequently, an operator might want to degrade P2P file sharing traffic volume to enhance the performance of the remaining applications. Again, this conforms to reasonable NN, as long as identical restrictions are imposed on every single P2P file sharing application. However, the technical feasibility of this scenario is unclear. A different realization of this consensus is referred to as QoS aware net neutrality. In this scenario, the user itself selects the QoS level which will be provided accordingly by the net operator.

There are many situations in which NN regulations come to the rescue of the end user. In [G06], four nightmare scenarios are presented. The first one, called "inequity nightmare", assumes that companies with a substantial market power may offer the higher tier, where the profit margins will be more lucrative. The net neutrality followers show that the investment in the upper tier may lead to an advanced Internet that will be available to only a small part of users. The second scenario: "corporate bureaucracy nightmare" represents the possibilities for charging extra money by the large corporate broadband firms and telephone or cable companies. The former may require a special charge for the access to the upper tier, *e.g.*, for a new adventuresome web site. At the same time the latter may insist that the users connected to that site are moved to a more expensive tier. It may cause that the network will become more expensive for the users. The third scenario, named "bad incentive nightmare" describes the situation where ISPs have their own services (like VoIP) and may block or discriminate their competition. The last presented scenario, called “less innovative content nightmare” involves worries that companies may produce new applications and services and protect their interests in those applications giving no chance for using and developing them by other providers. In [W07], an additional nightmare scenario is presented. It is possible that ISPs will charge application providers twice (firstly to its own ISP and secondly to the ISP of every single user who wants access to that application – a case that is confirmed in economics by the Two-Sided Market theory). It may begin to break the unique many-to-many nature of how information is linked in the Internet. The net neutrality problem and the complexity of its definition reflect a conflict of interests between content or application providers, Internet users, and ISPs.

Application and content providers usually second the net neutrality motion in the most rigid version believing that any service differentiation whatsoever should be prohibited and that all traffic in the Internet should be handled in the best effort manner. They argue that current network links have capacities high enough to carry all traffic with proper guarantees. In most cases it is really true because network resources are often over-dimensioned, but we have to be aware of an enormous progress in the telecommunications, especially in the area of access networks. New applications and services grow rapidly and the number of Internet users rises significantly as well. It may cause that in the future, the network link capacities will not be sufficient to carry all the traffic with a proper QoS. It is one of the key arguments raised by ISPs to allow for service differentiation and provide traffic priorities. They claim that the data transmission in the networks without QoS mechanisms will be unacceptable from the user perspective in the future. ISPs are also concerned that inducing net neutrality may impede network investments and development of new services and applications, since the abilities of investment cost returns are severely lessened.

The future of network neutrality is unclear. Andrew Odlyzko in [O08] states that: “*(…) inability to predict the development of technology and its impact on society mean that no fixed set of rules can work indefinitely*”, which is true. Huge progress in the telecommunication field may quickly render any regulation obsolete. It does not mean, however, that the issue is not important. As national telecom companies are naturally regulated when it comes to standard telephony, in the future, similar restrictions might affect carriers and Internet service providers. Following [O08] “*the general conclusion is that some form of government intervention, to set the rules, is inevitable*”.

5.10.2 Evaluation of Influence of NN on P2P Applications

Internet Service Providers and other telecommunication companies argue that the net neutrality regulation may prevent them from performing good enough network management. They suggest that it is necessary to block, filter or inspect for malicious viruses, spam or illegal content or just to manage the traffic differently in order to address different requirements (of course, the network performance requirements for Real Time applications are not the same as in the case of Web browsing or file downloading).

ISPs clarification that P2P traffic contains a high proportion of malware may be easily rejected. E-mail spam and web surfing are the vectors for malware, but the ISPs do not block such traffic as this would go against the user privacy to check the content of mails or web sessions. Nevertheless, this is the argument for the net neutrality flatterers. It means that packet inspection may be used by ISPs for unfair managing the traffic in their networks.

It is probable that in the accepted version of net neutrality concept the network operators will be prohibited from censoring the content, or sources of content, which travel across their networks. Some censorship actions are obviously required to deal with very real problems such as child pornography, spam and viruses, however, it is very important to separate this issue from possible censoring the content based on commercial or political interests. On the other hand, P2P users may intensify the process of encrypting all the traffic, to protect themselves from the consequences of their behavior (sharing the illegal content). It may lead to a type of "arms race" between users and ISPs [M07].

Now, there is no consensus as to the type and extent of traffic shaping and other forms of locking and throttling P2P traffic. However, in the near future, it may change. When planning the incentive mechanisms, we have to be aware of limitations described in this section. In particular, we should not propose and promote any solution, which requires or is based on deep packet inspection.

5.11 Intermediate Conclusions

So far, we have analyzed the types of incentives that are applicable to the end-users, the overlay providers and the network operators. We have also described the various implications of overlay traffic to the cost structure of an ISP and we have given examples of the tussle between ISPs and overlay networks.

From the analysis, we have seen that incentives provided to one stakeholder may introduce negative effects to another one. For example, the performance improvements that an overlay provider may want to introduce may come in direct conflict with the economic incentives for the operator (ISP), since such improvements may change the traffic patterns, affecting the interconnection agreements and charges for the specific ISP.

There are some other shortcomings that should also be taken into consideration. It can be anticipated that the number of P2P users will rise if all incentives are in place. While initially the impact will be positive due to minimization of inter-domain traffic, the increase in user base may, in the long run, provoke also a significant increase in inter-domain traffic. The trade-off requires further investigation.

The disclosure of information by the ISP concerning topology characteristics, link capacities, user population and distribution etc may also be a concern. The solution is to find the balance and to disclose generic information that is still useful.

Furthermore, a wide adoption of locality-aware techniques might lead to isolated content "islands", which is not a desired effect for the overlay. Thus, another point of conflict may arise. Moreover, changing the ratio of outbound and inbound traffic might as well trigger other types of charging which eventually might not be beneficial for the provider. Augmented intra-domain traffic might result in higher CAPEX and OPEX, leading to economic loss for the ISP. It is unclear yet if there are incentives for providers not to follow locality-aware rules when others do. Situations may arise that it is beneficial for providers to change their "strategy". Taking into consideration all these issues, it is therefore not wise to assume that locality-awareness is the only solution for the tussle between overlay and underlay.

Finally, there are legal issues that might be challenging to handle. If ISPs support networks that are means of sharing illegal content, they might be called to account for fostering illegal content distribution. However, there are many P2P networks that are legal and support licensed or free content.

From all the above, it is becoming obvious that not all kind of incentives and combinations may lead us to a situation where all stakeholders are benefited by incentives offered in underlay and overlay level. One of the main objectives in SmoothIT is to identify those incentives that are compatible and which, when combined, have non-negative effects in all stakeholders. These issues are going to be thoroughly discussed in Section 7.2.

6 Requirements for the SmoothIT System Design

Peer-to-peer overlay networks, in which different end users share their resources (CPU, storage, bandwidth, etc.), have become today's main option to distribute contents in the Internet, due to their inherent scalability. This has resulted in an important traffic growth in ISP networks, as it has been explained in the previous section. This issue is especially important if a scenario is considered where Content Providers are using peer-to-peer applications to distribute their contents.

This section identifies the key information that should be shared or exchanged between overlay and underlay, in order to address the current information asymmetry. Furthermore, it also discusses possible ways to share this information.

6.1 Information Asymmetry and Interaction Possibilities

Most algorithms for selecting peers in P2P applications are based on random selection due to the lack of locality information about peers and due to its potential to maintain a robust overlay network. The information asymmetry resulting from the **lack of communication between overlay and underlay** leads to an **increase of provider costs** and a **decrease of end user's Quality of Experience (QoE)**.

Therefore, it is desired that the underlay provides some kind of information and/or prioritization regarding peer selection to the overlay application. The aim is to support traffic management of the overlay application and to prevent any negative effects on both parties caused by the information asymmetry. In any case, any information exchange must be able to lead to a "win-win" scenario for all parties involved. The **prioritization** is the **result of an economic decision function** which takes into account both requirements: reduction of provider costs and improvement on users' QoE.

The following subsections review the main problems caused by the lack of interactions between the overlay and the underlay networks, and discuss about information to be exchanged to address these problems.

6.1.1 Provider Costs

P2P traffic has an important impact on the economics of Autonomous Systems (AS) as it was described in Section 5. By summarizing all the problems, the following main issues can be identified:

1. Intra-domain traffic is growing up due to P2P applications. Due to the lack of interaction between overlay and underlay networks (*e.g.*, peers select other peers from other PoPs instead of from their neighborhood), more traffic than expected must be managed in ISPs' networks.
2. Overlay traffic has an important impact on the inter-domain traffic. In particular, two clear scenarios can be drawn:
    a. An ISP (Tier-2 or 3) could experience an important growth in their interconnection costs. Especially, if it has a transit agreement and if the inter-domain traffic grows, it will have to pay more, and in case of a peering agreement, if the symmetry of the traffic is not maintained, this initial free peering agreement could evolve to a charged peering agreement.

b. A Tier-1 carrier could see how the changes in the traffic matrix could lead to violations in their interconnection agreements.

Taking into account the above scenarios and the different roles of ASs in the Internet hierarchy, a prioritization in terms of cost can therefore be made as follows:

1. Tier-2 and 3 ISPs' goals and tentative ways to achieve these goals:

   a. Optimization of the interconnection costs that will strongly depend on the type of interconnection agreement:

      i. If a ***transit interconnection agreement*** is in place, the ISP's goal will be to minimize the traffic that uses this IP transport interconnection.

      ii. If a ***peering interconnection agreement*** is in place, the ISP aims to maintain the symmetry in the traffic carried through this interconnection.

   In order to optimize the interconnection costs, the candidate information to be used is the BGP information correlated with the type of interconnection agreement. The following figure shows a scenario where the AS 65508 has just one peering agreement with the AS 65509 and the rest of the Internet locations are reached through a transit interconnection agreement (hub model) with the AS 65504, which is a Tier-1 provider.

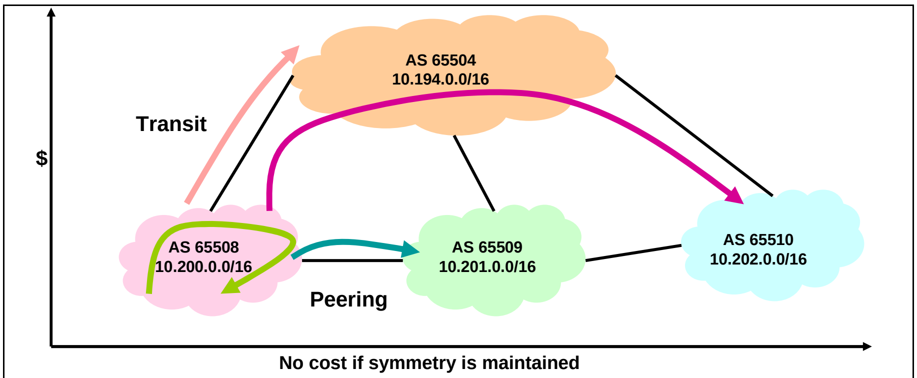

BGP Information of the AS 65508

Network	Next Hop	AS Path
10.194.0.0/16	10.194.0.1	65504 i
10.200.0.0/16	10.200.0.1	i
10.201.0.0/16	10.201.0.1	65509 i
10.202.0.0/16	10.194.0.1	65510 65504 i

Interconnection Agreements

Last AS Path	Interconnection Agreement
65504 i	Transit
i	Intra-domain
65509	Peering

Figure 6.1: Interconnection agreement information

   In this scenario, the AS 65508 will be interested in minimizing the traffic to be carried through the AS 65504 and will try to maintain the symmetry between the traffic sent to and received from the AS 65509.

   b. Minimization of the traffic within the AS: ISPs will try to minimize the number of hops needed for the overlay network formed by peers within the network of an ISP. In order to do this, one possible solution is to use the information maintained in the RADIUS (Remote Access Dial-up Service) servers typically used in xDSL access networks.

2. Tier-1 ISPs' goals:

   Tier-1 operators will focus on avoiding violations in their interconnection agreements. This issue will be strongly linked to the Operations, Administration and Maintenance processes of the Tier-1 carriers and will be dealt with in Section 6.1.3.

The following example illustrates this issue in a P2P based file sharing application. The initiating peer (New Peer) looks for a file which is available on nine other peers (from P1 to P9). Peer P1, P2, and P3 can be found in the local AS (and, moreover, P1 is accessible without crossing the IP/MPLS backbone of the ISP), peers P4 to P6 in a Peering AS and peers P7 to P9 in a Transit AS. Without knowing this information, the file sharing overlay application of the New Peer (or the end user) may choose (this choice is more or less random, if the bandwidth of the different peers is the same) to download the desired file from peers P4, P5, P7, and P9, which negatively influences the cost of the ISP of AS65508.

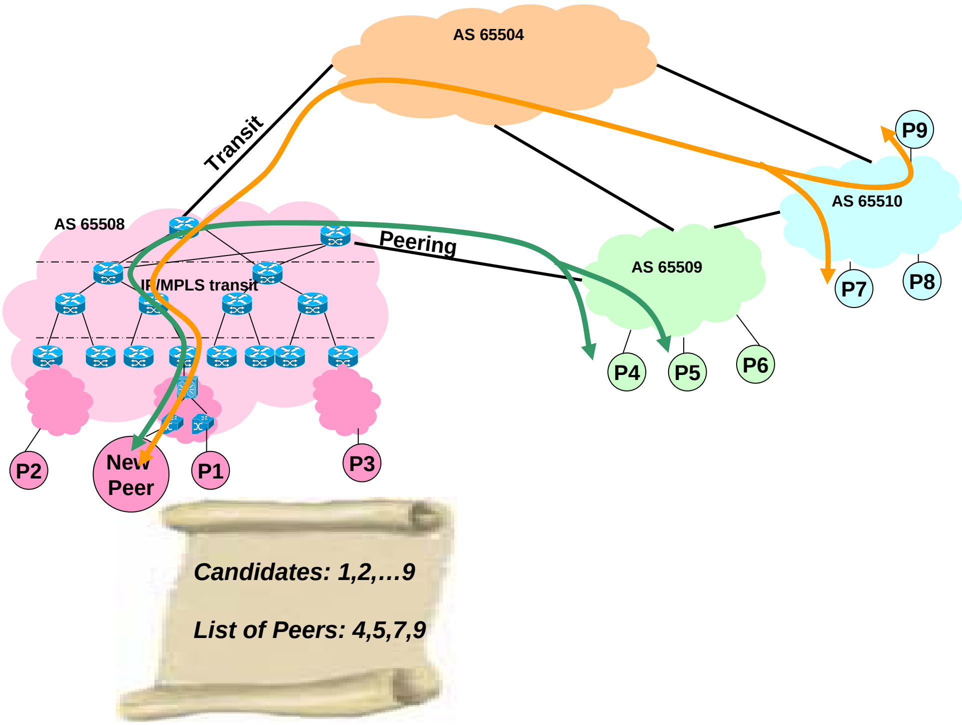

Figure 6.2: Peers selection without knowledge of the underlay network

6.1.2 Introduction of SIS to the Network

Taking into account the previous example, a new element called ***SmoothIT Information Service (SIS)*** element is proposed to provide a possibility for the overlay application to communicate with the underlying network. In this case, the incentive for the end user to use this SIS service is to improve its QoE, while the ISP aims to reduce the costs. Basically, the type of interactions needed has a request-response pattern. A peer requests information (attributes) for a given list of peers and specifies optionally the type of overlay application it is using. And the SIS sends a response back with the requested information (attributes) assigned to each peer in the list. Possible attributes include priorities and more detail information like locality, peer link capacity, peer availability, peer reputation, etc.

As a first step, when the SIS provides both inter-domain and intra-domain locality awareness, the SIS will suggest the New Peer to use peers P1, P2, P3, and P4 as depicted in *Figure 6.3*,. The first three peers are within the same AS and are sorted according to the number of hops in the IP/MPLS network of the AS; while P4 is in an AS with a Peering agreement. This would be optimal in terms of cost; however, in order to provide a real incentive to the end user, it is necessary to really improve the performance of the recommended peer-to-peer links by assuring that these links are, *e.g.*, not overloaded or by prioritizing the traffic. These issues are linked to the next sections, where OAM and QoS/QoE are detailed.

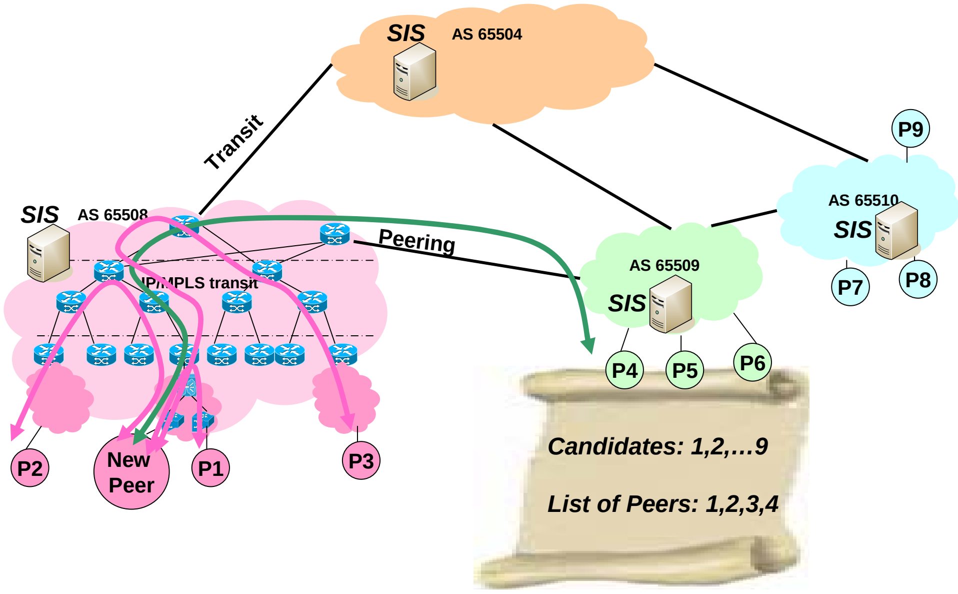

Figure 6.3: File sharing peers selection with SIS

6.1.3 Providers OAM

OAM (Operation, Administration and Maintenance) processes are very important in the management of interconnections, since they are responsible for monitoring of the fulfillment of the interconnection agreements by, *e.g.*, checking that the traffic does not exceed a defined threshold in inter-domain links. This issue will be especially important for a Tier-1 provider, since the penalizations for this kind of carriers are very important when a service violation is detected.

Moreover, another important issue is the traffic matrix estimation. Usually, network operators are interested in determining the traffic matrix in order to introduce changes in the underlay routing by, *e.g.*, reconfiguring the link costs or by provisioning a new LSP (Label Switched Paths) in an IP/MPLS based backbone. If these traffic matrices changes continuously, the routing decisions could lead to non-optimal traffic management, causing network congestion.

Therefore it is proposed to link the SIS to the OAM processes running in the different ISPs, in such a way that, in case of recommendation of a list of peers, the load in intra-domain links are taken into account. As an important starting point, it is recommended that the SIS becomes aware of the alarms generated in the network in order to avoid the recommendation of peers that have associated paths with detected failures. In this way, the SIS could be only aware of a specific set of alarms, *e.g.*, alarm type and level. The following figure shows an example of the interaction between the Alarm Manager (an important part of the OAM system) and the SIS; in this figure, it is shown that the SIS could offer a push interface to the Alarm Manager in order to receive the notification of the different alarms.

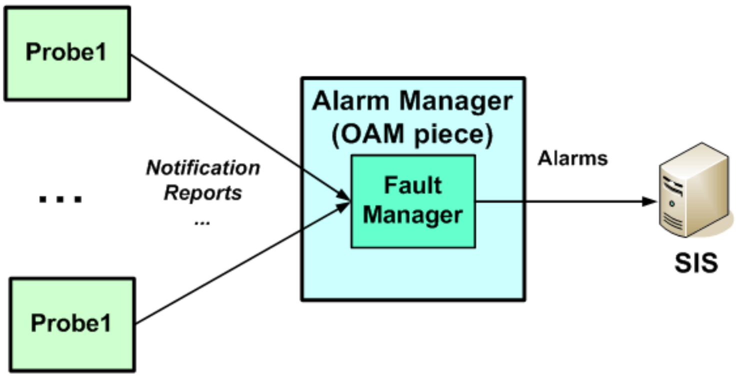

Figure 6.4: SIS and OAM integration

6.1.4 QoS and QoE

The cost based perspective on SIS has shown that also QoS and QoE are relevant to the providers' decision on how to prioritize peers of a communication. This subsection adds the need of quality aspects in the view of the end user. The end user typically has a specific application scenario where certain quality aspects are relevant. In a Voice over IP (VoIP) scenario the delay parameter may be relevant for the communication; in a file sharing this might not be as critical. Therefore the prioritization of peers in the end user perspective heavily depends on the specific application scenario in dependency of network parameters like delay, jitter, and others. This enhances the requirements of the ETM element as follows:

- SIS knows statistical values of each important parameter for its own AS.
- SIS elements in different ASes may communicate with each other to get the overall view of a communication in respect of the parameters specified.
- The SIS Peer Interface provides ability to specify the application scenario and the respective parameters.

The overall QoE of the end user in the optimal case should equal the prioritization. The QoS and QoE discussion has revealed the following further requirements to the SIS architecture:

- Ability to measure and store statistical values of the most important network characteristic properties. Note that if network status is aligned with admission control, the admission control has to be based on measurements. There have been a lot of work on this issue but it seems that solutions do not scale well in large scenarios where several measurements should be correlated in short time or have a lot of deployment problems. Therefore the specification of a measurement-based CAC able to work with different traffic types would be a real technical challenge.
- Ability to configure network resources.

6.1.5 Interaction Between Different ASs: Interconnection Planes

An important issue in the provisioning of QoS is the type of interconnection agreement that could exist between the different Autonomous Systems. The following figure shows an abstraction of the Interconnection agreements that could exist in today's networks.

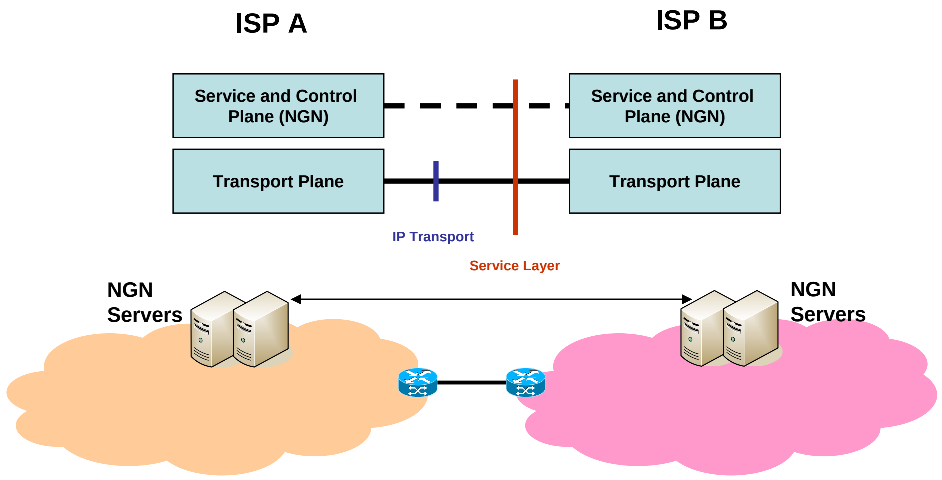

Figure 6.5: IP Interconnection levels

As it can be seen in *Figure 6.5*, there are 2 levels in the current interconnection framework:

- The Transport level provides interconnection at the IP level: it is able to offer interconnection to Internet and different Classes of services are supported in the inter-domain links. The types of agreements are peering and transit and the charging models are based on the difference between the traffic sent and received.
- The interconnection at the Service and Control Plane interconnection would mean that the different domains involved in an end-to-end path will be aware of the services that are being provided; this would allow, *e.g.*, cascading payments.
    - o An example of this kind of interconnection is the framework specified in the IP Packet Exchange (IPX) project carried out by the GSM Association to define an interconnection framework in NGN (Next Generation Networks) that is currently under testing. It aims at enabling customers in any network to contact each other using reliable and secure IP multimedia services. The IPX network will offer a number of key advantages over the open Internet, including guaranteed Quality of Service [GSMA08]. This actual interconnection framework is focused on the provisioning of VoIP services using SIP protocol following a model similar to that followed by PSTN (Public Switched Telephone Network).
    - o In the ITU-T (International Telecommunications Union – Telecommunications Sector) NGN, a specification of the interaction between the different RACFs (Resource and Admission Control Function) deployed in the different NGNs. This interaction seems to be implemented with RCIP (Resource Connection Initiation Protocol).

Taking into account these two layers, the SIS should be aware of the interconnection planes available in order to be able to interface with the different layers to, *e.g.*, provide QoS incentives; and, if interaction of the different SIS at different domains is defined, it could be seen as an extension of the interconnection framework.

6.1.6 Other Considerations

The following four subsections outline additional aspects to be considered.

6.1.6.1 Security and Privacy

One important topic in today's networks is security and privacy. Even though the SIS will not be an entity in charge of checking the integrity of the contents, taking into account that the SIS will provide a service to the end users, a minimal set of requirements must be considered to provide a secured service and to preserve the privacy of the end users.

6.1.6.2 Differentiated Pricing

The support of different CoS (Class-of-Service) implies differentiated pricing, since otherwise all users will request and use best service quality, if they pay the same for all CoS. This leads to a proposal to also specify prices in addition to priorities assigned to peers in a peer selection process.

6.1.6.3 Mobility

Mobility causes changes in locality and other peer related information. Therefore, information returned by SIS has a certain lifetime. Supporting mobility may lead to the need to support a notification pattern for interaction. Otherwise, a peer may have to invoke a SIS service periodically. There is a disadvantage of using a notification pattern, namely a SIS must be statefull.

6.1.6.4 Business Plane: IPSphere Work

The IPSphere forum is defining a general framework that could ease the agreements between network providers, service providers, content providers, etc. This general framework is known as the SSS (Service Structuring Stratum) that could support the negotiation process between the different players in the provisioning of a new service.

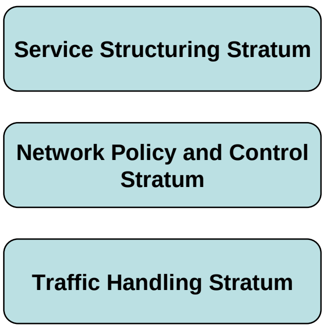

Figure 6.6: IPSphere stratums

Figure 6.6 shows the main stratum being defined in IPSphere. As it can be seen, what is being defined is a SSS that could influence on the lower stratum that are related to the Interconnection planes defined in the previous sections.

If SIS shall be integrated in the IPSphere framework, a possible solution would be to integrate it in the SSS: effectively, the SIS could be seen as an interface that could be used by a content provider, such as, *e.g.*, a P2P based TV video streaming distributor. By means of the usage of the SIS service (a business level agreement) the ISP could enforce QoS policies in its networks to assure a proper service delivery. Of course, this would be a step forward the current interconnection models based on basic connectivity provisioning.

6.2 Requirements for the SmoothIT Architecture

Based on the discussion above, the following main functional requirements are identified for the SmoothIT architecture:

R.1. Incentive-compatibility: The SIS shall provide incentives for ISPs, overlay providers, and end-users that result in a win-win situation for all three parties.

R.2. Support of different overlay applications: The SIS shall provide an open service that is accessible by different P2P applications.

R.3. Interface supporting various optimization schemes: The interface between the SIS and the overlay application shall provide means to specify the application scenario and the respective parameters. Due to the various incentives of ISPs, overlay providers, and end-users, the SIS shall provide several services (*e.g.*, "Throughput Optimization", "QoS enhancement") that could be classified into *free* and *premium* (charged) network services.

R.4. Different mode of operation: The SIS shall be able to operate in two different modes: user anonymity mode for free services and user aware mode for premium services.

R.5. Network status gathering: The SIS shall be able to measure and store the most important network characteristics and properties. The SIS shall know statistical information of each important parameter for its own AS.

R.6. Inter-domain support: The SIS deployed in different ISPs shall be able to interact with each other. SIS elements in different ASs may communicate with each other in order to get the overall view of a communication in respect of the optimization parameters specified.

R.7. QoS support: The SIS shall support QoS for network services and it shall be able to configure network resources.

R.8. OAM support: The SIS shall be able to interact with the OAM processes of the ISP.

R.9. Mobile network support: The above requirements should also be valid in the context of a cellular network operator, which is characterized by the following key properties: node mobility, heterogeneity of nodes and link capacities, and presence of shared medium.

The following non-functional requirements are identified for the SmoothIT architecture:

R.10. Traffic optimization: The SIS shall optimize traffic of the overlay network taking into account the underlying network information.

R.11. Easy deployment: It shall be easy to extend existing overlay applications with the functionality of the SIS and it shall be easy for ISPs to deploy the SIS in their networks.

R.12. Extensibility: The SIS shall be extendible to support new overlay applications, new optimization attributes, and new metrics (both application-driven and provider-driven).

R.13. Scalability: The SIS shall be scalable to support a large end-user population.

R.14. Efficiency: The operation of SIS shall be efficient in terms of communication (bandwidth) overhead, storage consumption, and processing requirement.

R.15. Robustness: The SIS shall be robust against malicious behavior and against dynamic behavior (churn of peers). It shall be also fault tolerant.

R.16. Security: Secure communication between SIS entities and between SIS and overlay application shall be supported, providing message origin authentication, data integrity, and data confidentiality. Any data storage in the system shall provide data integrity, confidentiality, and authentication.

R.17. Data privacy and legislation/regulation: The SmoothIT architecture needs to provide interfaces for regulation aspects, such as data retention, and it has to address data privacy concerns, which are determined by the European Directives on Security.

R.18. Standard compliance: The SIS shall use and based on standard protocols where applicable.

7 Overlay Application Classification for SmoothIT

This section shows the overlay application classification for SmoothIT and finally reasons for the overlay application selection for the internal trial. In Section 7.1, we apply the technical classification criteria and evaluate them for the application classes introduced in Section 2 and application examples of Section 3. In Section 7.2, we discuss the incentives that will lead to a situation where all players benefit from the mechanisms employed by the overlay applications, i.e. a TripleWin situation. To this end, we consider the various overlay applications and characteristics and propose compatible incentives per each class of applications. Finally Section 7.3 represents the documentation and reasoning on the final decision on the overlay application used in the internal trial (milestone M1.1) and takes into account additional aspects neither captured by technical nor economic criteria alone.

7.1 Application of Classification Criteria

In this section, the classification criteria introduced in Section 4 are applied to find the most promising application for SmoothIT. To back up the criteria discussion, we start to evaluate the preliminary results of the user survey to initially assess the popularity of applications and the acceptance of users for additional charging. All classification discussions are summarized in table form based on the criteria of Section 4. To capture the expected evolution of these applications in the current Internet, we apply the criteria from today's view point in the year 2008, as well as in the near-future in the year 2010, when the SmoothIT final results become available. This investigation allows deriving the application service type of interest, which has the largest optimization potential, to demonstrate the SmoothIT approach and.

The goal is to identify a particular example application which will be further used during the course of this project. This is especially important for WP3 and WP4 which require an actual implementation of the overlay application.

For visualization, we use a matrix where the columns represent the classification criteria and where the rows depict the application service class and the overlay application examples, respectively. In each cell of the matrix, a numerical value between 1 and 5 is used. The minimum value of 1 means that that the criteria is not or marginally addressed, while the maximum value of 5 means that the criteria is fully addressed.

The particular values for the evaluation were worked out and agreed upon by the SmoothIT partners in the SmoothIT project meetings in Athens, March 11-13, 2008, and in Munich, June 12-13, 2008. It has to be noted that the notion of *BitTorrent* in this section summarizes the existing BitTorrent clients and modifications, like for example the Vuze application. After evaluation of the user survey, the SmoothIT consortium decided to use Vuze as an alternative in the internal trial of SmoothIT (see Section 7.3). Therefore, Vuze is explicitly mentioned in Section 7.1.3 "Classification of Overlay Application Examples".

7.1.1 Popularity and Potential Additional Charging

In order to assess the popularity of the P2P applications a questionnaire was prepared, as introduced in Section 4. The questionnaire was distributed to over 100 persons. The surveyed persons were students at AGH-UST.[2] The choice of the target group was motivated by their accessibility and the fact that they can be accounted to the most frequent users of the considered applications.

The surveyed persons were asked to answer questions respecting overlay applications:

- "do you know the application ?"
- "do you use the application ?"
- "would you use the application if you should pay for it ?"
- "would you pay for enhanced quality ?"

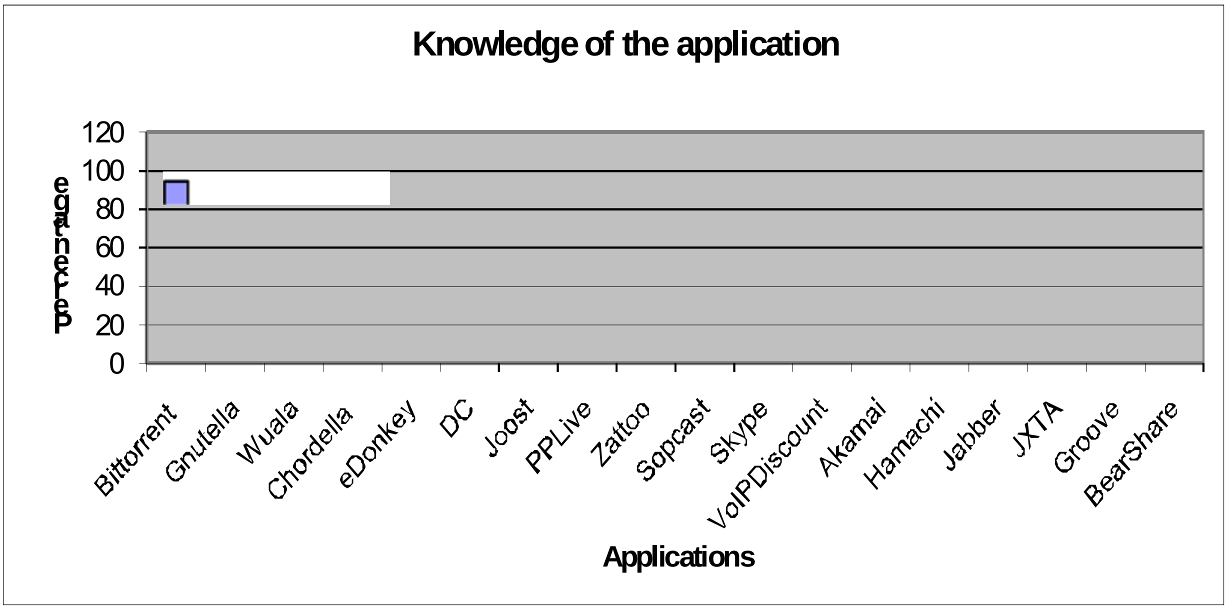

Figure 7.1: Knowledge of the application

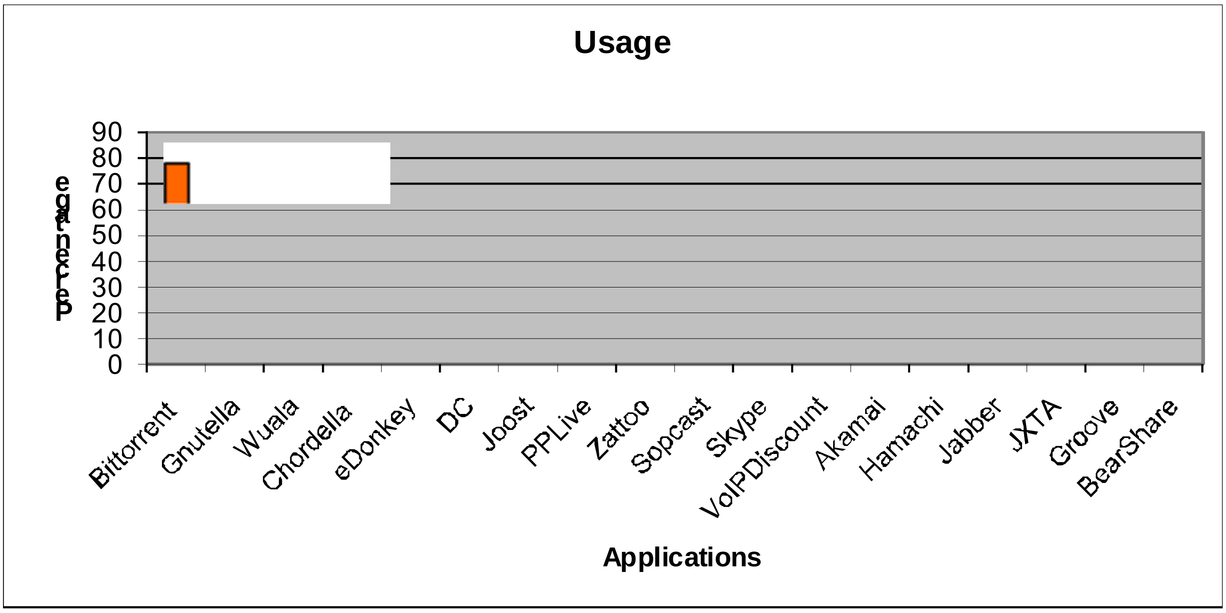

Figure 7.2: Application Usage

[2] In this deliverable, preliminary results of the user survey are provided. In the future, we will hand out the survey to test persons / students of the participating partners in SmoothIt and present the results later on.

Figure 7.1 shows the percentage of persons who know relevant applications. Results confirm that the most widely known (80-100%) are applications used for file sharing (BitTorrent, eDonkey and DirectConnect) and VoIP (Skype). Second group, in the middle range (40-60%), includes Gnutella (file sharing), Sopcast (P2P-based live TV), Hamachi and Jabber. However there are applications that are practically unknown for the target group: Wuala, Chordella, Zattoo or Akamai.

Figure 7.2 indicates fractions of persons using overlay applications. As it can be seen these results are strictly related to the percentage of persons who know about respective applications shown in Figure 7.1.

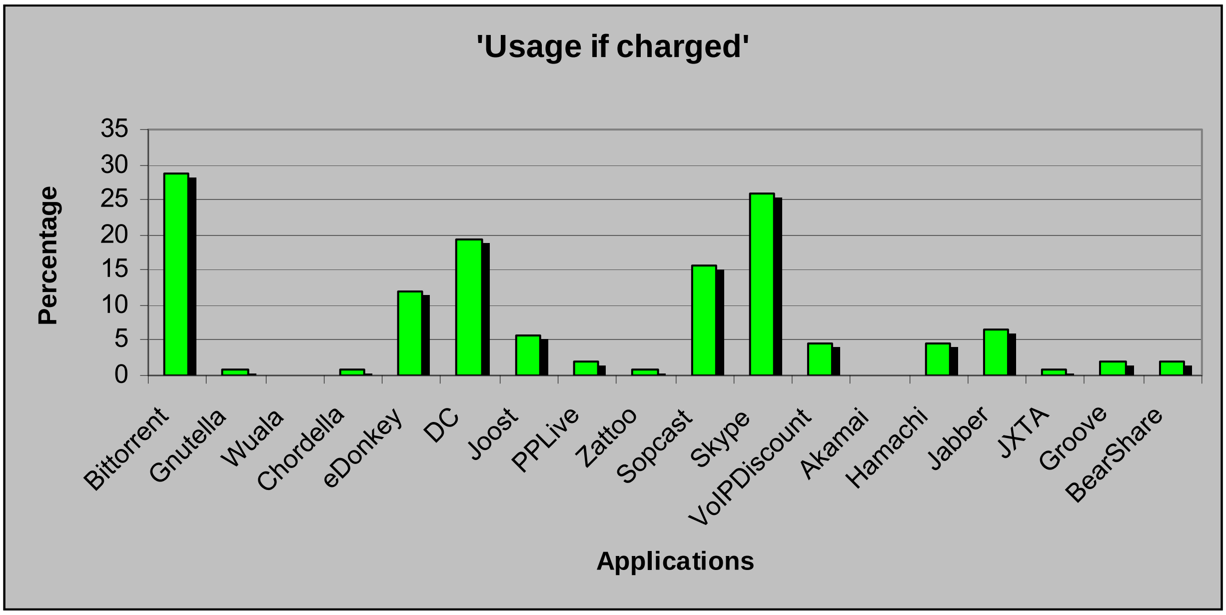

Figure 7.3: Acceptance for charging applications

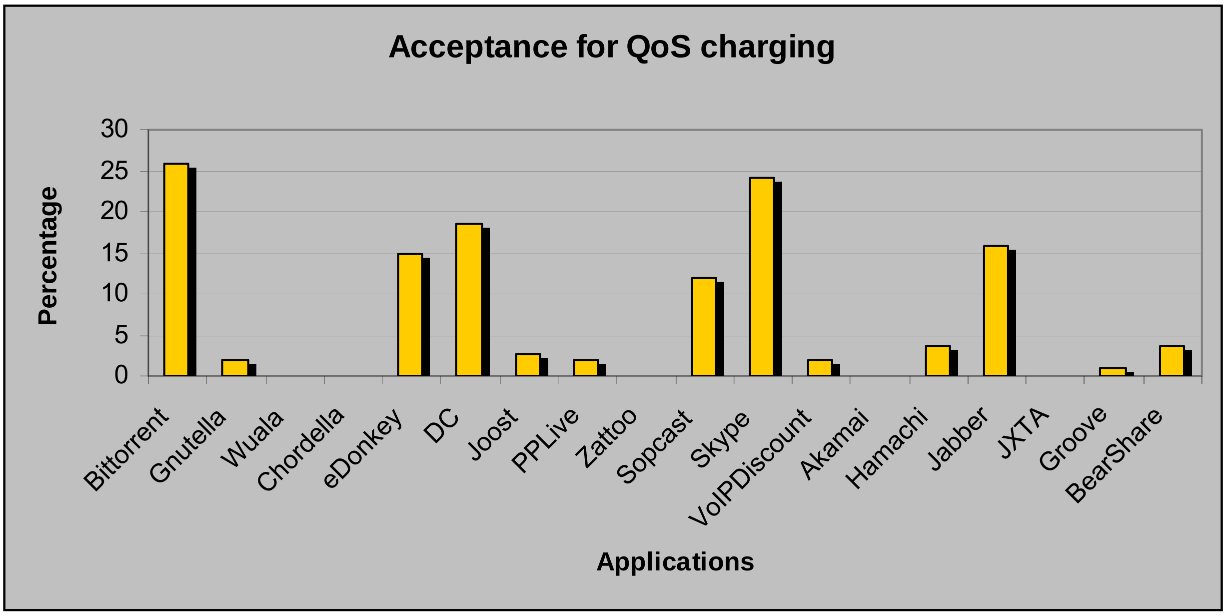

Figure 7.4: Acceptance for QoS charging

According to answers represented in *Figure 7.3* there exist quite significant percentage of persons who would use some overlay applications even if they were charged:

- File sharing: BitTorrent, eDonkey and DirectConnect
- P2P-based live TV: Sopcast
- VoIP: Skype

Figure 7.4 shows potential acceptance for additional charging under condition that the user would be provided with enhanced QoS for particular overlay application. When compared to Figure 7.3 it can be stated that the number of persons ready to pay for application itself and for the enhanced quality of the application do no differ much. BitTorrent and Skype are the leading applications in these categories again.

Apart from applications considered by SmoothIT the surveyed persons could indicate other applications that they use. These most frequently indicated were as follows:

- File Sharing: BearShare, Azureus and Mtorrent
- IP TV: AntsTV, PPStream
- VoIP: GG, Tlen (local country-wide applications)

Only BearShare was included in the results shown in the figures.

7.1.2 Classification of Overlay Application w.r.t. Application Service Types

Table 2 shows the filled matrix for the classification of overlay application with respect to the application service types as discussed in Section 2. To capture the expected evolution of these applications in the current Internet, we apply the criteria from today's view point in the year 2008, as well as in the near-future in the year 2010.

Although one may argue about an exact value for a particular cell, the overall picture shows a clear trend. Currently, file sharing applications seems to be very promising to realize the different tasks in SmoothIT, including engineering and modification of the overlay application as well as for the internal trial. However, P2P VoD and P2P live TV are the next interesting candidates within SmoothIT from today's viewpoint. In the near future, the picture slightly changes and it is expected that P2P live TV will be most promising. However, P2P VoD and file sharing are in the same order of magnitude comparing the overall score.

Table 2 - Classification of Overlay Application w.r.t Application Service Types

	source code availability	traffic intensity	traffic recognition	end-user controllability	utilization of QoS provisioning	utilization of locality	popularity	legal contents	ISP costs	additional charging	opportunities	overall
Year 2008												
File sharing	4	5	3	5	2	5	5	2	5	1	5	**42**
P2P VoD	2	3	3	2	4	5	3	4	3	3	3	35
P2P Live TV	2	3	3	2	5	5	3	5	3	3	2	36
P2P VoIP	1	2	3	1	4	1	5	5	1	1	1	25
P2P Gaming	1	1	3	1	1	1	1	1	1	5	1	17
CDN	1	4	5	1	4	4	5	3	3	2	1	33
Year 2010												
File sharing	4	5	3	5	2	5	5	2	5	1	5	42
P2P VoD	4	4	3	3	4	5	4	4	4	3	4	42
P2P Live TV	4	5	3	3	5	5	5	5	5	3	4	**47**
P2P VoIP	1	2	3	1	4	1	5	5	1	1	1	25
P2P Gaming	2	3	3	2	4	2	3	5	3	5	2	34
CDN	1	4	5	1	4	4	5	3	3	2	1	33
Weights	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	
	technical criteria			**optimization**			**non-technical criteria**					

7.1.3 Classification of Overlay Application Examples

The actual overlay application examples discussed in Section 3 are now classified, but from today's viewpoint, and shown in Table 3. Again, file sharing seems most appropriate for SmoothIT. In this context, the two most popular file sharing networks, eDonkey and BitTorrent, get the highest scores. Due to the scientific interest in BitTorrent, the opportunities to use this file sharing protocol are the highest. In addition, modifications of BitTorrent clients, like the open source Vuze application, exist which try to provide video streaming additionally. For this reason, it is expected that future versions of Vuze will integrate enhanced P2P mechanisms for supporting video streaming. Thus, Vuze has a higher optimization potential regarding the utilization of QoS provisioning than pure file-sharing BitTorrent clients. It has to be noted that the remaining BitTorrent clients are summarized under the point BitTorrent in Table 3. Therefore, the BitTorrent protocol and in particular the Vuze application seems the most promising candidate from today's point of view.

In order to take care of future interesting applications, P2P live TV and P2P VoD applications are especially considered. From the existing software implementations, the differences are only marginal. However, PeerCast seems most promising as it is open-source. However, the popularity is rather limited due to the offered content. Vice versa, the popular P2P video streaming applications are proprietary. One exception is Vuze, which is a popular Bittorrent client that also supports (limited) P2P video streaming (VoD) over Bittorrent. Since it's open-source, it's possible that it could be adapted to be used in SmoothIT.

Table 3 - Classification of Overlay Application Examples

	source code availability	traffic intensity	traffic recognition	end-user controllability	QoS provisioning	utilization of locality	popularity	legal contents	ISP costs	opportunities	overall
eDonkey	3	5	3	5	1	5	5	3	4	4	38
BitTorrent	3	5	3	5	1	5	5	3	4	5	39
Vuze	5	5	3	5	3	5	5	3	4	5	**43**
Gnutella	3	3	3	5	1	5	2	3	4	3	32
Wuala	3	3	3	3	1	5	2	3	4	1	28
Chordella	5	1	5	4	1	5	1	5	1	5	33
Joost	1	4	3	1	3	5	3	5	3	3	31
PPLive	1	4	3	1	3	5	5	3	5	3	33
PeerCast	5	4	3	3	5	5	2	3	2	3	35
End System	5	4	3	3	5	5	1	3	1	3	33
FreeCast	5	4	3	3	5	5	1	3	1	3	33
Nodezilla	5	4	3	3	5	5	1	3	1	3	33
Zattoo	1	4	3	1	5	5	4	5	4	3	35
SopCast	1	4	3	1	5	5	4	3	4	3	33
Skype	1	1	2	1	5	1	5	5	1	1	23
Hamachi	1	1	5	1	4	1	2	4	2	1	22
Akamai	1	4	4	1	5	5	4	5	3	1	33
Weights	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	*1*	
	technical criteria			**optimization potential**			**non-technical criteria**				

7.2 Discussion about Incentives Leading to a TripleWin Situation

This section discusses incentives that will lead to a situation where all players benefit from the mechanisms deployed in the underlay and/or employed by overlay applications, i.e. leading to a TripleWin (i.e win/win/win) situation. By considering the various overlay applications and their characteristics, incentives per each application class are proposed that can be compatible with the incentives of all players. This section concludes by pointing out the necessity of a framework that offers incentives and provides mechanisms for all players so that it is economically beneficial for them to participate.

7.2.1 Classification According to Economic and Business Aspects

Prior to looking into the incentives that can lead to a win situation for all the stakeholders in the underlay and overlay, i.e. ISPs, end users and overlay providers, applications are categorized according to their inherent economic/business characteristics. These characteristics can be summarized as follows:

- Management of overlay structure (centralized or distributed)
- Distribution of content (proprietary, open or hybrid)
- Participation/sharing mechanisms
- Traffic requirements (delay-sensitive, bandwidth-demanding, or combination)
- Service differentiation (network or application level)

One of the key issues in an overlay network is who controls the structure of the overlay network. In one case, there can be a single entity controlling the formation and structure of the overlay, which can either reside in one domain or span into multiple domains. An example of such an overlay is the overlay network in the case of Joost. On the other hand, the overlay can be self-maintained in a distributed manner by all the participating nodes, i.e. nodes where end-users and super-peers are located, as in the case of a BitTorrent distribution overlay network.

Complementary to the above distinction is who distributes the content, i.e. who makes the content available over the overlay. A proprietary overlay consists of nodes that offer content belonging to a single entity, namely the overlay provider, who places the content to the nodes. Thus, such an overlay can be either centrally controlled by a single entity in the backbone or it can be formed in a distributed manner by the nodes. In an open overlay content belongs to the end-users but can be either offered by a centrally controlled overlay (the case of the early P2P systems) or by a distributed, self-maintained overlay. There is also a mix of approaches, a hybrid system where there exist content servers owned by the overlay provider and edge nodes that help in further distributing the content.

The previous classification criteria had to do with the provision of the overlay services. Another important aspect of overlay networks is the mechanisms that are employed by the overlay network concerning the participation of node to the content distribution and the sharing of resources. There are cases where there is no monitoring or “policing” of participation, i.e. peers are free to share and/or download as many resources as they want. This is mainly the case for proprietary overlays since users pay so as to have access to licensed content and do not have to comply with participation and sharing rules, apart from the fact that they cannot re-distribute the content outside the overlay. Usually, in such cases there exist some mechanisms that are transparent to the end users and help the distribution of popular content so as not to overload the content servers. On the other hand, there are overlays where users should conform to a set of rules regarding requirements for participation and sharing. Participation rules refer to the prerequisites for a node to become member of an overlay node. Early file sharing systems have introduced relevant rules that required a volume of data (in Megabyte) to be shared in order for the node to be accepted. Sharing rules usually come in terms of an upload/download bandwidth relationship. Other applications require that a certain number of upload connections is open in order for a peer to download content from others.

Another important aspect that characterizes an overlay is the nature of the traffic that circulates in it. For some applications, download bandwidth is the key issue (*e.g.*, for file sharing applications), while for others the minimization of delay is critical (*e.g.*, VoIP applications). Of course, there can be a combination of characteristics. For example, VoD applications require both an acceptable downlink rate and a specific delay in order to deliver the required QoE.

One last aspect of overlay applications that is of great interest is the possibility of providing service differentiation, either in the network layer compared to the other types of traffic inside a domain (*e.g.*, path differentiation, better QoS) or in the application layer (*e.g.*, enhanced software for more options, new features).

Based on the characteristics above, a classification of overlay applications is summarized in Table 4.

Table 4 – Classification of overlay applications

	Ownership and control	**Overlay mechanisms**	**Traffic req.**	**Service differentiation**
Akamai	Proprietary/centralized	No policing	Combination	Network
Skype	Proprietary/distributed	No policing	Delay	Application
Joost	Proprietary/hybrid	No policing	Combination	Application
Wuala	Proprietary/hybrid	Participation/Sharing	Bandwidth	Application
SopCast	Proprietary/distributed	No policing	Combination	Application
BitTorrent	Open/distributed	Sharing	Bandwidth	Application
Vuze	Open/distributed	Sharing	Bandwidth	Application
eDonkey	Open/distributed	Participation/Sharing	Bandwidth	Application
eMule	Open/distributed	Participation/Sharing	Bandwidth	Application
Gnuntella	Open/distributed	Sharing	Bandwidth	Application
PPLive	Proprietary/hybrid	No policing	Combination	Application
Hamachi	Proprietary/distributed	No policing	Combination	Network
Zattoo	Proprietary/hybrid	No policing	Combination	Application

7.2.2 Incentive Compatibility for Classes of Applications

As already mentioned in Section 5, the various incentives applicable to each stakeholder may not be compatible with each other and can lead to conflicts between the players. Therefore, the objective is to find which types of incentives can be compatible, according to the previous classification, so that a situation can be reached, where all stakeholders are satisfied, thus leading to a TripleWin situation.

At this point, it is necessary to make an important observation: although the stakeholders in this environment are three, namely the end-user, the overlay provider and the network operator, conflicts may appear only between the underlay (network operator) and the overlay (end-users and overlay provider) entities. Indeed, conflicts between the end-user and the overlay provider can only occur in the (improbable) case that the provider makes some drastic changes to the overlay application that alter the nature of the service provided, rendering it not beneficial for the end users. Henceforth only conflicts between these overlay and underlay sides are considered.

In the following, compatible incentives are discussed based on the criteria mentioned in the previous subsection. Regarding the ownership of content, it is obvious that when the systems are open and free content is exchanged, monetary incentives and pricing mechanisms cannot be employed. End-users and the overlay provider are not willing to start paying for a service that is supposed to be free. Hence, the underlay cannot recover directly a part of its cost, by charging the overlay. On the other hand, there is an indirect way to share costs, *e.g.*, by offering service differentiation on the underlay, as mentioned in the paragraphs that follow. In the case of proprietary content though, since users pay to acquire the content and overlay providers have probably made special agreements with the operators, pricing mechanisms can be employed in the form of discounts so as to affect the user/overlay choices.

Centralized or distributed approaches do not introduce any specific problems and conflicts. In the case of centralized provision, where the overlay is controlled by a single entity, it is easier to impose certain incentives or rules. In the case of distributed provision, smarter and scalable mechanisms should be introduced. Closely related with this issue is the existence of participation and sharing rules in the overlay. If such mechanisms already exist, then it is easier to enhance them with economic notions in order to achieve the desired outcome. On the other hand, when such mechanisms are not in place, indirect methods should be used in order to affect the operation of the overlay.

The characteristics of the traffic that flows in the overlay network do not introduce any limitations to the mechanisms that can be applied. They just determine what performance improvements can be made and which traffic characteristics are most important.

On the contrary, the applicability of differentiation, either in the network or in the application layer, greatly affects the way incentives can be introduced. As mentioned above, monetary incentives can be provided in the case of proprietary content. But, for example, even if this is not the case, a provider of an open overlay could design and offer a second version of the same application with some enhancements that increase the QoE. Such enhancements/optimizations should be charged, even though access to the content still remains free. Some legal issues arise here, though, that render this option inapplicable, like DRM (Digital Rights Management) conflicts in case the exchanged content is illegally obtained. Network Neutrality issues also arise here since access for the same content can be charged or not. Differentiation could also be realized in the network layer, in terms of traffic characterization and appropriate QoS techniques. Of course, combination of both application and network layer differentiations are possible, as long as no conflicts arise.

7.2.3 Compatibility of incentives

Section 5 identified several incentives for end-users, overlay providers and network operators. The two most important incentives, common for all the stakeholders are the *monetary benefits* and the *performance improvements*.

Monetary benefits, in the form of cost reduction or profit increase, are desirable in the case of network operators since the effect of overlay applications is directly translated into loss/gain of money in the interconnection market. But for such monetary benefits to occur, the network provider should be able to recover a portion of his costs (due to the overlay traffic) back to the overlay providers and/or overlay users, by charging them accordingly, or to reduce the inter-domain traffic resulting in cost reduction as well.

Monetary incentives cannot be directly applied to overlay applications that are open, unless the overlay provider wants to offer two versions of the same software; namely, one free version that provides the necessary functionality and one enhanced version that is charged and has extended features (offered in both the underlay and overlay). Even in this case, these incentives are not necessarily compatible with those of the underlay provider. For proprietary applications, it is easier to employ monetary incentives, but the way money flows between the stakeholders should be clear, together with how this flow of money can be altered. Monetary incentives could affect the relation of the overlay provider with both the network operator (in case there exists a business agreement between them, *e.g.*, if the operator hosts some of the proprietary servers of the overlay provider) and the end-user (since users pay to have access in proprietary content). Such benefits should be provided to the end-users and application providers in the cases they conform to the objectives of the network operator for reducing his costs.

Performance improvements can substitute or complement the monetary benefits, depending on the status of the overlay network. They can substitute monetary incentives when the latter cannot be applied (*e.g.*, for open overlays). They can also be a complement to monetary incentives, when actions affecting the network performance can be translated into monetary gain/loss for one of the stakeholders. In particular, the target here is to shape the overlay traffic in such a way that both end-users and overlay providers benefit, while the network operator does not risk violating his interconnection agreements and keeps the performance of his network at an acceptable level.

For an example of how monetary and performance incentives can co-exist, consider the case where the traffic of an overlay application increases the transit costs of an ISP. The ISP has two alternatives: either to block/limit the overlay traffic at the inter-domain link, or try to co-operate with the overlay provider and recover a part of the costs, by transferring it to him through charging. If the overlay application and content are proprietary, the overlay can easily include those costs to his provision costs. Parts of those costs will be transferred to the end-users by the overlay provider. But an increase in the provision costs is never welcome by the end-users, unless there are other benefits. In our case, the ISP could guarantee some performance improvements for the end-users, in order to justify the option for increased payments. For the case of open overlays, since there cannot be a direct charging for content, the ISP and overlay provider can co-operate in order to provide the end-users with an enhanced version of the software that offers better QoE while considering the costs of the ISP. The free version could continue to exist but the ISP would give lower priority to its traffic while the volume of exchanged data should stay below a certain limit, so as not to violate any interconnections agreements of the ISP.

Concluding this section, we provide a summary of how incentives can be combined in order to reach a TripleWin situation:

- Monetary benefits can offer the desired outcome, if there is a possibility of transferring/recovering costs through charging schemes.
- Monetary benefits can also apply when combined with performance improvements or with service differentiation in general.
- Performance improvements should not be considered as substitutes but as complementary to monetary benefits.
- Performance improvements for one stakeholder can provide monetary benefits for another one or vice versa. In other words, the type of incentives provided may not be the same for all stakeholders.

It is also important to point out the necessity for a general framework that takes into consideration all the aforementioned incentives, analyzes the economic effects of the overlays applications, identifies the conflict of interests and proposes the mechanisms that solve those conflicts. It is the main objective of SmoothIT to propose such mechanisms that achieve the collaboration of all stakeholders and lead to a win-win situation for the overlay and underlay.

7.3 Selection on Application for Internal Trial

The basic reasoning for narrowing down the decision for the selection of an overlay application service class in one of the a) P2P VoD, b) P2P Live TV and c) File Sharing is that all three service-classes seem to be future-proof, as showed explicitly in Section 7.1.1.

The selected overlay service class for the internal trial therefore needed to be relevant to the needs of users in one and a half year time, but it also needed to be capable and scalable to cope with heavily increased resource demands in both user and traffic requirements.

Towards that aim the SmoothIT consortium decided in unison to devise not one but two action plans concerning the decision on the selection of the application to be utilized for the internal trial.

Plan A, is driven by a strong desire of the SmoothIT project to engage in a bilateral beneficial agreement with the P2P NEXT, also an EU-funded research project[3]. P2P NEXT will build on the Tribler technology.

If we establish a cooperation agreement with P2P NEXT, the SmoothIT consortium will be clearly targeting the P2P VoD and/or the P2P Live Streaming overlay application service class. This decision brings the focus on mesh architecture and in effect on a client such as Tribler that is based on the BitTorrent protocol.

The reasoning for seeking an agreement with P2P NEXT will be the actual prospect of seizing the mutual benefits from a potential collaboration amongst two FP7-ICT projects.

Another factor will be the fact that the source code of Tribler is publicly available under the GPL Open Source license.

Concluding P2P VoD is especially interesting for SmoothIT, due to its high optimization potential.

Plan B suggests selecting Vuze (formerly called Azureus) as the overlay application for the internal trial. More specifically Vuze is based on the popular Azureus Java BitTorrent client implementation with optional integrated adaptations for video streaming; in effect enabling both file sharing and video streaming.

Moreover Vuze, like Tribler share some common characteristics that are vital for the application decision for the internal trial: a) both are published under the GPL Open Source license and b) both are mesh based.

The decision to select Vuze, is particularly important and interesting, since it provides to the SmoothIT consortium a backup solution in the essence of using a different overlay service class namely file sharing, if eventually streaming does not work effectively.

The basic reasoning for choosing P2P Video Streaming as our main overlay application service class and in consequence devising Plan A and backup plan B, is that the potential benefits from overlay/underlay interaction are extremely high. Both plans aim to tackle the

[3] Tribler is an open source P2P client with various features for watching videos online and is available for Linux, Windows and Mac OS X. P2P NEXT might extend it to also perform P2P Live. Streaming.

increased bandwidth problems that operators currently face and will continue to experience even more strongly in the near future, while the positive demand for such high bandwidth applications continues to exist. Another major factor for these choices will be nonetheless than the high capability for improvements for all major stakeholders, in effect seeking additional benefits for the end-users, the overlay providers and the network operators.

8 Summary and Conclusions

Objective of the project SmoothIT is to define, develop and test Economic Traffic Management (ETM) mechanisms to optimize the traffic impact of overlay applications on ISP and telecommunication operator networks based on a cooperation of network operators, overlay providers and application users. This deliverable summarizes the initial investigations in respect with overlay applications, their characteristics, and their classification according to their significance for ETM and explains initial ETM approaches. As main outcome for the further work in SmoothIT, we have derived a set of 18 requirements for the system design in SmoothIT. P2P-based streaming has finally been selected as the target reference overlay application for our investigation and system development. In particular for the internal trial the P2P VoD streaming application currently developed in the ICT project P2P Next has been selected, with the overlay application Vuze as a backup solution.

SmoothIT participants have performed an in-depth investigation of overlay applications related to the classes file sharing, P2P video on demand, P2P live TV, VoIP, P2P gaming, CDN and VPN. Their evaluation was based on technical and non-technical criteria in order to judge their relevance for the SmoothIT objectives. Evaluation criteria include traffic intensity, traffic recognition, optimization potential, popularity, legal content and charging possibilities. In addition, overlay applications are characterized by their design parameters such as overlay algorithms, overlay topology, or QoS requirements, which are implicitly contained in the optimization potential discussion and have more impact on WP2 algorithm development then. Heterogeneity of network systems (including wireless networks) and user mobility have an impact on all evaluated overlay applications. As they are not addressed in most of the currently available applications, their impact is considered as a general requirement for the project work. As a result of this evaluation P2P-based video streaming and P2P (BitTorrent-style) based file sharing have been identified as the most relevant and important applications for SmoothIT. This was motivated by their high traffic impact, their popularity and their optimization potentials such as locality promotion. In particular, P2P streaming was prioritized as its popularity is expected to overtake file sharing and it shows a higher optimization potential, i.e., sensitivity in respect with QoS and QoE. This is also reflected in the trial application selection where the P2P VoD overlay application of the ICT project P2P Next (Plan A) and the overlay application Vuze (Plan B) were selected for the internal trial. Here especially the availability of open source software was taken into account additionally.

In order to derive the requirements, not only overlay application characteristics have been analyzed, but also the potential incentives for the different stakeholders and their effects have been investigated from an economic and regulatory viewpoint. Incentives for end users include performance improvement, peer availability and peer reputation. For overlay providers incentives are performance improvement and user loyalty toward the overlay application. Operator could benefit from traffic optimization and related general performance improvement and loyalty of users selecting and staying with the ISP. Monetary benefits may also be applicable strongly depending on the application. Based on an investigation of overlay related costs, possible ETM mechanisms such as interconnection agreements and locality promotion were discussed initially. A more detailed description and evaluation can be found in deliverable D2.1.

The overlay application and incentive discussion led to a set of requirements for the SmoothIT system design manifesting around the lack of information exchange (information asymmetry) between overlay providers and network operators. An information service (referred to as SmoothIT Information Service, SIS) should be provided by network operators to optimize overlay traffic taking into account the underlying network. Possibly distributed SISs should provide an open, reliable, scalable service that can be differentiated in free and premium services provided anonymously respectively customer-aware.

This deliverable documents all initial project investigations in respect with high-level requirements and application selections. These findings will be updated at M1.2 in D1.2

9 References

[Ak08] Akamai, http://www.akamai.com

[AW04] M. Afergan, J. Wroclawski, “On the benefits and feasibility of incentive based routing infrastructure,” PINS `04, Portland, Oregon

[BS06] S. Baset, H. Schulzrinne, “An analysis of the skype peer-to-peer internet telephony protocol”, Infocom 2006

[C06] P. Cochrane, “Net Neutrality or Suicide?”, Proceedings of the IEEE, vol. 94, pp. 1779-1780, October 2006.

[CIS08] “Global IP Traffic Forecast and Methodology”, 2006-2011, Cisco 2008

[CMH+07a] Piotr Chołda, Anders Mykkeltveit, Bjarne E. Helvik, Otto J. Wittner, Andrzej Jajszczyk, “A Survey of Resilience Differentiation Frameworks in Communication Networks”, IEEE Communications Surveys & Tutorials, vol. 9, no. 4, 4th Quarter 2007.

[Co03] B. Cohen, “BitTorrent protocol specification”, First Workshop on Economics of Peer-to-Peer Systems (P2P’03), Berkeley, CA, 2003

[Co08] B. Cohen, “BitTorrent Protocol Specification”, http://www.bittorrent.org/beps/bep_0003.html, 2008

[DLH+05] C. Dana, D. Li, D. Harrision, C.-N. Chuah, “BASS: BitTorrent Assisted Streaming Systems for Video-on-Demand”, International Workshop on Multimedia Signal Processing (MMSP), Shanghai, China, 2005

[Ed08] eDonkey2000 homepage: http://www.edonkey2000.com

[Em08] eMule homepage: http://www.emule-project.net

[F07] P. Faratin, ”Economics of Overlay Networks: An Industrial Organization Perspective on Network Economics,” (NetEcon+IBC), in conjunction with ACM Conference on Electronic Commerce (EC'07) San Diego, California, June `07

[G06] S. Greenstein, “Four nightmares for net neutrality”, Micro, IEEE, vol. 26, pp. 12–13, November – December 2006.

[GDJ06] S. Guha, N. Daswani, R. Jain, “An experimental study of the skype peer-to-peer voip system”, IPTPS 2006

[Gn08] The Gnutella group, “The Gnutella Protocol Specifiaction v0.4”, http://www9.limewire.com/developer/gnutella_protocol_0.4.pdf

[Go03] I. Gosling: “eDonkey/ed2k: Study of a Young File Sharing Protocol”, SANS Institute Technical Report available at http://www.gia.org/practical/GCIH/Ian_gosling_GCIH.pdf, 2003

[GSMA08] The GSMA’s IP Packet Exchange (IPX) project, http://www.gsmworld.com/about/people/emc.shtml

[Ha08] Hamachi homepage, https://secure.logmein.com/products/hamachi/

[HB02] O. Heckmann, A. Bock, “The eDonkey 2000 Protocol”, Tech. Rep. KOM-TR-08-2002, Multimedia Communications Lab, Darmstadt University of Technology, Dec. 2002

[HB08] T. Hoßfeld, A. Binzenhöfer, „Analysis of Skype VoIP Traffic in UMTS: End-to-End QoS and QoE Measurements”, Computer Networks, online http://dx.doi.org/10.1016/j.comnet.2007.10.008, Vol 52/3 pp 650-666,, 2008.

[HL08] T. Hoßfeld, K. Leibnitz, “A Qualitative Measurement Survey of Popular Internet-based IPTV Systems”, Second International Conference on Communications and Electronics (HUT-ICCE 2008), Hoi An, Vietnam, 2008.

[HLP+04] T. Hoßfeld, K. Leibnitz, R. Pries, K. Tutschku, P. Tran-Gia, K. Pawlikowski, "Information Diffusion in eDonkey Filesharing Networks", ATNAC 2004, Sydney, Australia, December 2004

[IS07] http://www.ipoque.com/media/internet_studies/internet_study_2007

[LCJ+07] J. Lv, X. Cheng, Q. Jiang, J. Ye, T. Zhang, S. Lin, L. Wang, „LiveBT: Providing Video-on-Demand Streaming Service over BitTorrent Systems", International Conference on Parallel and Distributed Computing, Applications and Technologies, PDCAT'07, Adelaide, South Australia

[LRL+08] J. Liu, S.G. Rao, B. Li, H. Zhang, "Opportunities and Challenges of Peer-to-Peer Internet Video Broadcast," Proc. IEEE, 2007.

[M07] C. T. Marsden, "Net Neutrality and Consumer Access to Content", SCRIPTe, Vol. 4, No. 4, available at SSRN: http://ssrn.com/abstract=1089063

[Ml08] MLDonkey homepage: mldonkey.sourceforge.net

[NOB] NOBEL phase 2 – Document, Deliverable 2.3. "Report on Resilience Mechanism for NOBEL Solutions in Medium and Long Term Scenarios". URL: http://www.ist-nobel.org/Nobel2/imatges/D2[1].3_v1.1am.pdf

[O08] A. Odlyzko, "Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets," available at http://www.dtc.umn.edu/~odlyzko/

[P06] J. M. Peha, "The Benefits and Risks of Mandating Network Neutrality and the Quest for a Balanced Policy," 34th Telecommunications Policy Research Conference, September 2006.

[PCE+07] A. Popescu, D. Constantinescu, D. Erman, D. Ilie, "A Survey of Reliable Multicast Communication", 3rd EuroNGI Conference on Next Generation Internet Networks (NGI2007), Trondheim, Norway, 2007

[PN08] P2P-Next integrated project, http://www.P2P-next.org

[SCK+06] A-J. Su, D. R. Choffnes, A. Kuzmanovic, F. E. Bustamante, "Drafting behind akamai (travelocity-based detouring)", SIGCOMM 2006

[Sk08] Skype homepage http://www.skype.com

[SMG+07] A. Sentinelli, G. Marfia, M. Gerla, L. Kleinrock, and S. Tewari, "Will IPTV ride the peer-to-peer stream?", IEEE Communications, vol. 45, pp. 86-92, June 2007.

[SP07] P. Shah, J.-F. Paris, "Peer-to-Peer Multimedia Streaming Using BitTorrent", 26th International Performance of Computers and Communication Conference (IPCCC 2007), New Orleans, LA, 2007

[Tr08] Tribler project homepage http://www.tribler.org

[VIF06] Aggelos Vlavianos, Marios Iliofotou, Michalis Faloutsos, "BiToS: Enhancing BitTorrent for Supporting Streaming Applications", INFOCOM 2006. 25th IEEE International Conference on Computer Communications, pp.1-6, April 2006.

[W07] D. Weitzner, "Twelve Billion Bargaining Chips, "The Web Side of the Net Neutrality Debate", Internet Computing, IEEE, vol. 11, pp. 78–81, January – February 2007.

[Wu08] Wuala homepage, http://wua.la

[ZDK08] S. Zoels, Z. Despotovic, and W. Kellerer, "On hierarchical dht systems – an analytical approach for optimal designs," Computer Communications (COMCOM), Vol 31, Issue 3 (February 2008), pp. 576–590, 2008.

[ZHD+08] Zoels, S.,Hofstaetter, Q., Despotovic, Z. and Kellerer, W., "Achieving and Maintaining Cost-Optimal Operation of a Hierarchical DHT System," Institute of Communication Networks, Technische Universitaet Muenchen, Munich, Germany, Technical Report, 2008.

[ZW2004] K. Zaraska, K. Wajda, "Identification of Internet services on the basis of statistical profiles of traffic streams", Secure 2004 conference, 21-22 October 2004, Warsaw, Poland.

10 Abbreviations

3G	3rd Generation
AGH	Akademia Gorniczo-Hutnicza im. Stanislawa Staszica W Krakowie
AS	Autonomous System
BGP	Border Gateway Protocol
BRAS	Broadband Remote access Aggregation Router
CAPEX	Capital Expenses
CDN	Content Delivery Network
CoS	Class-of-Service
CP	Content Provider
CU	Content User
DHT	Distributed Hash Tables
DiffServ	Differentiated Services
DRM	Digital Rights Management
DSL	Digital Subscriber Line
DSLAM	DSL Access Multiplexer
DWDM	Dense Wavelength Division Multiplexing
eBGP	external Border Gateway Protocol
ESM	End-system Multicast
ETM	Economic Traffic Management Mechanisms
Gbps	Gigabit per second
GPON	Gigabit Passive Optical Network
GPRS	General Packet Radio Service
HTTP	Hypertext Transfer Protocol
IBP	Internet Backbone Provider
ICOM	Intracom
IntServ	Integrated Services
IP	Internet Protocol
ISP	Internet Service Provider

ITU-T	International Telecommunications Union – Telecommunications Sector
IXP	Internet Exchange Point
Mbps	Megabit per second
MPLS	Multi Protocol Label Switching
MSD	Multiple Source Download
NAP	Neutral Access Point
NAT	Network Address Translator
NGN	Next Generation Networks
NN	Network Neutrality
OAM	Operation, Administration and Maintenance
OLT	Optical Line Termination
OPEX	Operational Expenses
P2P	Peer-to-peer
POP	Point of presence
PSTN	Public Switched Telephone Network
PoS	Packet-over-SONET
QoE	Quality of Experience
QoS	Quality of Service
RACF	Resource and Admission Control Function
RADIUS	Remote Access Dial-up Service
RTP	Real Time Protocol
SIS	SmoothIT Information System
SmoothIT	Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies
SSS	Service Structuring Stratum
STREP	Specific Targeted Research Project
TCP	Transmission Control Protocol
TE	Traffic Engineering
TID	Telefónica Investigacíon y Desarollo
TUD	Technische Universität Darmstadt
TV	Television
UDP	User Datagram Protocol
UniWue	Julius-Maximilians Universität Würzburg
UMTS	Universal Mobile Telecommunication System
UZH	University of Zürich
VoD	Video on Demand

VoIP	Voice-over-IP
VNC	Virtual Network Computing
VPN	Virtual Private Network

11 Acknowledgements

This deliverable was made possible due to the large and open help of the WP1 and WP2 team of the SmoothIT team within this STREP, which includes besides the deliverable authors as indicated in the document control, a number of additional persons as well. Many thanks to all of them.