Towards the Future Internet
G. Tselentis et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-539-6-1

ETMS: A System for Economic Management of Overlay Traffic

Sergios Soursos[a,b,1], María Ángeles Callejo Rodriguez[c],
Konstantin Pussep[d], Peter Racz[e], Spiros Spirou[b],
George D. Stamoulis[a], Burkhard Stiller[e]
[a]*Athens University of Economics and Business, {sns, gstamoul}@aueb.gr*
[b]*Intracom Telecom R&D, {souse, spis}@intracom.com*
[c]*Telefónica Investigacion y Desarrollo, macr@tid.es*
[d]*Technische Universität Darmstadt, pussep@kom.tu-darmstadt.de*
[e]*University of Zürich, {racz, stiller}@ifi.uzh.ch*

Abstract. The concept of Economic Traffic Management (ETM) encompasses various techniques for optimizing overlay networks considering both, underlay *and* overlay networks' performance requirements as well as the resulting economic implications for ISPs. This work presents several mechanisms through an overall ETM System (ETMS), identifying the possibility for synergies between mechanisms, both in the sense of complementarity of decision grounds and in the sense of functionality and components employed thereby. The paper describes the core ETMS architecture and how various mechanisms are instantiated. It continues with the discussion of the flexibility and modularity of this architecture, allowing for the accommodation of synergies. Finally, it presents selected results from the test-bed trials of the ETMS and a dedicated discussion on incentives behind these ETM mechanisms.

Keywords. Overlays, Peer-to-Peer (P2P), Economic Traffic Management (ETM), Incentives, Architecture, Modular Design

1. Introduction

Recent research focusing on the optimization of overlay networks has identified the tussle between Internet Service Providers (ISPs) and peer-to-peer (P2P) applications. This is due to different views of the former with respect to the overlay-optimized routing of traffic generated by P2P applications, which might result in increased interconnection costs for ISPs. Localization of overlay traffic appears as one of the most promising solutions to tackle this emerging problem. Despite the criticism on locality promotion [1], [2], there are significant efforts from the research community, industry, and standardization bodies to devise mechanisms for traffic localization of overlay-based applications [3], [4], [5], and [6]. One of the most popular of such application is BitTorrent. Those mechanisms, however, are partially similar, both in concept and in implementation, since in their core they provide the overlay with the information about network topology, even if they use different information sources and

---

[1] Corresponding Author.

granularity. Throughout this paper, we identify those approaches and spot the differences with the presented proposed solutions.

The objective of the SmoothIT project [7] is to define the framework for applying Economic Traffic Management (ETM) techniques, categorize them, implement the most promising ones, and evaluate their performance with respect to key benefits for all stakeholders. That is to achieve the TripleWin situation [8], where ISPs, overlay providers, and end users all benefit in terms of performance optimization and monetary gains. Thus, this work provides an overview of main ETM mechanisms (Section 2), presents the architecture of the ETM System (ETMS), which can be used to implement any of the aforementioned mechanisms (Section 3), describes how each of these ETM mechanisms is mapped to the generic architecture (Section 4), presents selected preliminary results from the deployment of the system in a controlled test-bed (Section 5), and finally, discusses economic implications of those mechanisms with respect to offered incentives for all stakeholders (Section 6), before concluding in Section 7.

2. Economic Traffic Management Mechanisms

Within the SmoothIT project, a multitude of approaches for enabling overlay traffic management techniques that lead to a TripleWin situation, have been identified. The most important of these can be classified into the following three main categories:

Locality Promotion involves peers of a domain having their overlay neighbors rated according to their underlay proximity. A peer sends a list of peers to the SIS (SmoothIT Information Service). The latter entity interfaces with the underlay network, obtains locality information, and rates each peer of the list based on certain location attributes, *e.g.*, on the number of hops in the Autonomous System (AS) path. The rated and accordingly sorted list is returned to the requesting peer where, in turn, overlay operations (*e.g.*, unchoking and neighbor selection in BitTorrent) are modified in order to favor peers that belong to the same domain with the requesting one. Thus, the interconnection costs for the ISP are reduced, while the performance for the peer is also improved in a variety of cases. Dynamic extensions of this mechanism can involve the promotion of locality only when necessary according to underlay conditions. The level of interference of SIS with overlay operations is reduced, also avoiding any negative impact on the quality of services offered.

Insertion of locality-promoting peers/resources: An indirect way of promoting locality for the ISP is to introduce to its domain special resources so that content is downloaded faster into the domain and distributed among the local peers. Two such ETM mechanisms developed are the ISP-owned Peer (IoP) and the Highly Active Peer (HAP). With the IoP, the ISP deploys and controls a peer with augmented bandwidth and storage capabilities that allow downloading content and making it available quickly to the domain's peers. To do so, the IoP initially participates in the swarm as a normal peer, but quickly acquires the entire content and starts serving the local peers. By the HAP this idea is implemented in a decentralized manner, with the functionality being passed to normal peers. In particular, based on overlay and underlay information, certain highly active peers in a domain are promoted by the ISP to enhanced peers and their Internet connection speed is increased dynamically by the ISP, so as to download and offer overlay content more efficiently.

Inter-domain collaboration addresses cases, where information available to a domain is not sufficient to reach some of the aforementioned optimization objectives.

This can be due to asymmetry of routing information between source and destination domains or it is due to the lack of connection details of communicating peers. In such cases, inter-domain collaboration can lead to fine-tuning and improvement of local decisions in all participating domains.

3. ETMS Architecture

Figure 1 depicts the ETMS architecture developed. Underlay and overlay networks and their components, which interact with the ETMS, are also included for completeness. As observed, the SIS encompasses the main functionality offered by the ETMS.

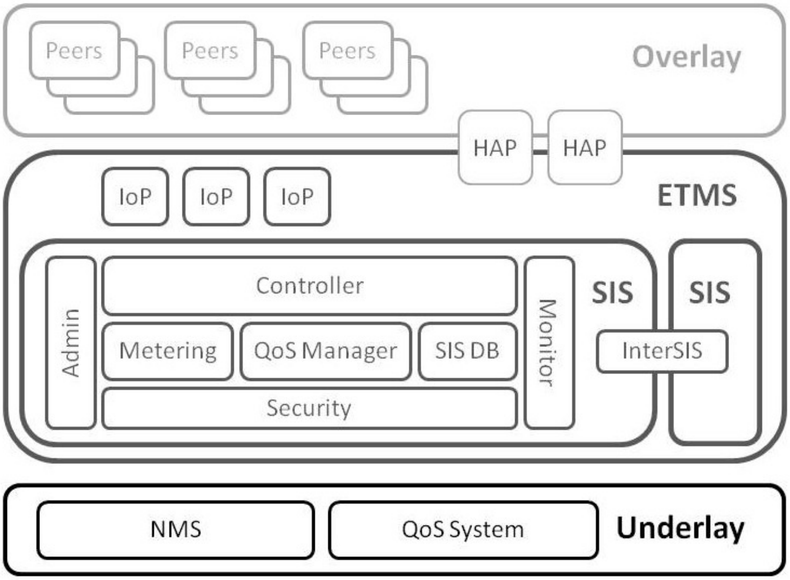

Figure 1. The ETMS Architecture

Components of the SIS, the entire ETMS architecture, and their functionality cover:

- The *Controller* determines the core of the system. Its main responsibility is to coordinate the ETM mechanism. Therefore, it receives requests from the overlay application (simple peer, IoP, or HAP), performs calculations based on the ETM mechanism deployed and according to several underlay metrics, such as metering and policy information, and returns to the overlay application the information required. Usually, the information returned is in the form of ranked lists, rating either peers as possible “good” neighbors, or swarms as “popular” ones.
- The *Metering* component collects network information from the underlying Network Management System (NMS) in order to support ETM mechanisms implemented by the SIS. This information can include, *e.g.*, BGP (Border Gateway Protocol) routing tables in order to support locality enforcement algorithms, network performance parameters and network usage by users that is necessary to support accounting and charging.

- The *QoS Manager* checks the availability of network resources, guarantees resources requested by the end user, and enforces QoS (Quality-of-Service) policies in the network. For example, it interfaces to the network by using the NGN (Next Generation Network) transport control functionalities, if available.
- The *InterSIS* component facilitates the exchange of information between two or more SISs to fine-tune local decisions.
- The *SIS DB* (Data Base) is a repository storing all information that may be useful for all modules, such as configuration parameters or business logic.
- The *Security* component provides security services to the SIS, including authentication, access control, and secure communication.
- The *Admin* component is used by the administrator of the SIS to access and configure the server as well as the deployed ETM mechanisms.
- The *Monitor* component is a supportive component used to gather certain overlay and underlay statistics for evaluation purposes.
- The *IoP* and the *HAP*, as already explained, denote two special cases of peers. In case of the IoP mechanism, this special peer belongs to the ETMS and is fully controlled by the ISP (through the SIS). In case of the HAP, the peer is a regular peer that is granted with special network capabilities by the ISP, for a given time, hence it belongs both to the overlay and to the ETMS.

The ETMS architecture designed has been implemented as a prototype. It follows a modular design and allows interfacing with external elements (e.g. NMS) without affecting the mechanism logic.

The overlay application is based on NextShare, a P2P streaming application developed by the P2P-Next project [9]. The application is written in Python and it has been extended to communicate with the SIS. The modified application also includes the biased neighbor selection and biased unchoking mechanisms to support locality promotion. Additionally, the application supports live monitoring by periodically sending XML (eXtended Markup Language) reports to the Monitor component. Reports include overlay statistics (*e.g.*, total download and upload, download and upload rate, or download progress) as well as video playback statistics (*e.g.*, like play time or stall time) and they are stored in the SIS DB.

All components of the SIS have been implemented in Java and are deployed on a JBoss application server. The Controller provides a Web Service interface over SOAP (Simple Object Access Protocol) to communicate with peers in the overlay application. The current version of the Controller supports the BGP-based locality promotion (*BGP-Loc*) and the *IoP* mechanisms, while the *HAP* mechanism will be implemented as a next step. The Metering component supports reading the BGP routing table over SNMP (Simple Network Management Protocol) from a router and provides information based on the routing table to the Controller. The SIS DB is implemented using the Java Hibernate persistence service, making it independent of the underlying database server. The prototype uses a MySQL database server, but it supports any relational database systems. The Admin component provides a Web interface to configure and manage the system. It is implemented using the JavaServer Pages (JSP) technology. The Security component provides the authentication and authorization to log in to the admin interface.

The final prototype will include all three mechanisms and will be deployed in a real environment, i.e. in the premises of an existing ISP, in order to evaluate the performance of these mechanisms and complement existing simulative and theoretical evaluations.

4. Mapping of the Architecture to ETMs

The ETMS provides a service to the overlay network to offer any ETM mechanism, which is supported by components involved and of the respective functionality. In order to support a specific ETM mechanism architectural components are extended with the desired functionality that implements the specific ETM algorithm. Depending on the algorithm, a subset of components presented in the previous section is extended. In any case the Controller must be extended with the specific ETM logic and the Admin interface is adjusted to enable parameter configuration. Furthermore, SIS DB is extended to store the required data per mechanism. Changes in further components depend on the mechanism, but typically only few components must be extended.

4.1. Support for Locality Promotion

The mechanism for promoting locality (*BGP-Loc*) requires two features to be supported by the architecture: (i) the ability to obtain locality information from the underlay and (ii) the provision of the pre-processed locality information to the overlay application. The first feature is supported by the Metering component, which acquires underlay information, like routing information from ISP's routers. This information is used by the Controller to characterize peers according to an underlay metric, *e.g.*, their AS, POP (Point-of-Presence), or any other relevant topology information like peering agreements among ISPs. One such metric is the BGP routing preference that combines several BGP attributes to provide the final rating value [10].

For the second feature, the Controller implements a generic rating algorithm that can provide ratings for a peer based on various metrics. This rating value is used by the SIS Client (residing in the peer) to bias the neighbor selection and unchoking algorithms [11] in a manner that promotes connecting to and exchanging data with peers of the same or of a nearby domain. For this mechanism, the extension of the client is crucial to take advantage of the locality promotion.

The extension of dynamic locality enforcement requires the Metering component to gather information related to the load level of interconnection links, as well as to the charging level of flowing traffic (*e.g.*, 95th percentile). By doing so, the Controller is able to know, when it is required to promote locality to both, while keeping charging levels below an ISP-defined threshold and limiting the intervention of the ISP to overlay-related procedures. This decision on whether to promote or not locality can be seen as another parameter in the Controller's generic rating function.

4.2 Support for ISP-owned Peer

The IoP [12] involves deploying special peers (actual IoPs) in the ISP's domain. These peers act as locality-aware ultra-peers and bias the overlay traffic for higher locality degree. The IoP support requires the EMTS architecture to acquire both underlay *and* overlay information to allow the IoP to make right decisions regarding the swarm, neighbor, and unchoke management. Therefore, the IoP and the Controller need to communicate, in order for the IoP to identify popular swarms to join, and decide on how much bandwidth to allocate to a swarm. The Controller provides estimates on the popularity of various swarms by collecting overlay statistics from local peers. Additionally, the Controller can advertise the IoP(s) by including it to the list of peers sent for rating by local peers, thus combining the IoP and locality promotion mechanisms.

4.3 Support for Highly Active Peers

In case of the HAP mechanism, the Controller decides, based on underlay and overlay metrics, which peers to promote to HAPs. By doing so, special policy profiles will be activated for those peers by the QoS Manager. These profiles provide additional upload and/or download capacity, for which purpose the QoS Manager communicates with the NMS or Traffic Shaping devices depending on the ISP's infrastructure and access type of the user.

The behavior of HAPs is monitored to decide whether their extra resources should be maintained and to what extent. Peers promoted to HAPs can be rewarded, *e.g.*, in terms of monetary discounts or for contributing to enhance the overlay's performance. But the main incentive to be a HAP is a better overlay performance (in terms of valuable upload capacity) resulting in better contribution rankings, *e.g.*, for BitTorrent networks, and higher download rates experienced by the users. In order to avoid the misuse of the HAP status by peers, per-user statistics are collected by the Metering equipment and provided to the Controller to adjust HAP promotion decisions. Alternatively, the overlay application can provide its own statistics to the Controller that allows offloading NMS equipment. Unlike the BGP-Loc and the IoP, this mechanism does not require any changes to or support of the overlay protocol.

4.4 Support for InterSIS Collaboration

For the refinement of local decisions, SIS instances must implement an InterSIS protocol and decide on which information to exchange and how to use it. This protocol connects the Controllers of different SIS instances (probably run by different ISPs or other entities). Again, this information is translated to another parameter to be considered by the generic peer rating algorithm that affects the overlay behavior of local peers.

Table 1: Overview of Architectural Adjustments

Supported ETM mechanism	Involved Components
BGP-Loc	Controller, Peer, Metering, (SIS DB, Admin)
IoP	Controller, IoP, Peer, (SIS DB, Admin)
HAP	Controller, QoS Manager, Peer or Metering, (SIS DB, Admin)
InterSIS	Controller, (SIS DB, Admin)

5. Test-bed Trials and Preliminary Trial/Prototype Results

The SmoothIT test-bed has been designed to validate the ETMS implementation in a medium scale scenario. The first mechanism validated is BGP-Loc. For this purpose, the ModelNet emulator [13] has been used which is a large-scale network emulator that allows users to evaluate distributed networked systems in realistic Internet-like environments. ModelNet enables the testing of unmodified prototypes running over unmodified operating systems across various networking scenarios.

In order to evaluate the impact of the BGP-Loc mechanism implementation, scenarios with different ISPs having different interconnection models, different access types, and distribution of peers have been generated. The topology that has been emulated in ModelNet is shown in Figure 2.

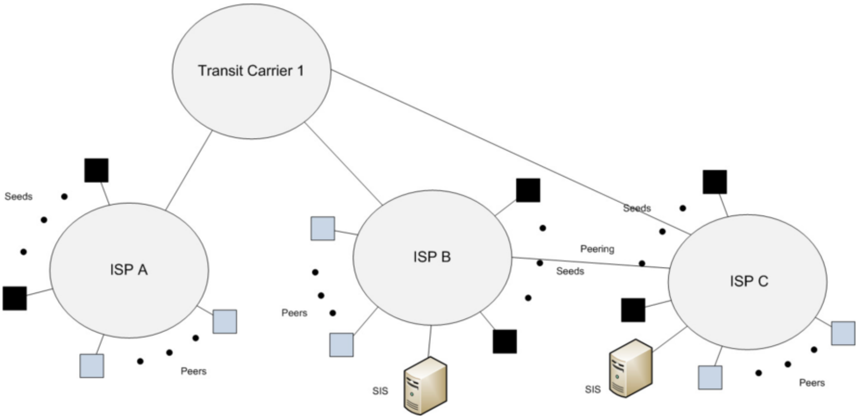

Figure 2. Internal Trial Topology

The main advantage of this environment is that it allows for the configuration of multiple network conditions in order to evaluate the performance of different mechanisms. In particular, while assuming a homogeneous distribution of the seeds (1 seed per domain) and leechers (8 leechers per domain) in different domains, the first scenario being evaluated shows the effectiveness of the BGP-Loc mechanism, especially when both ISPs B and C deployed the SIS with that ETM mechanism. These results are shown in Figures 3 and 4. Note that due to the limited scale of these experiments (small swarm size) these results are indicative.

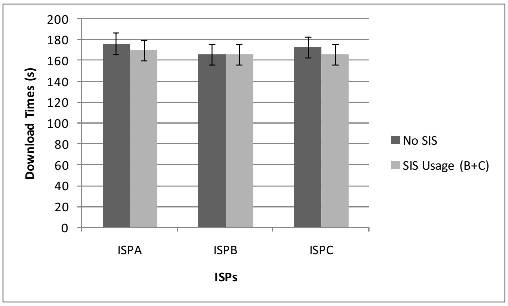

Figure 3. Download times

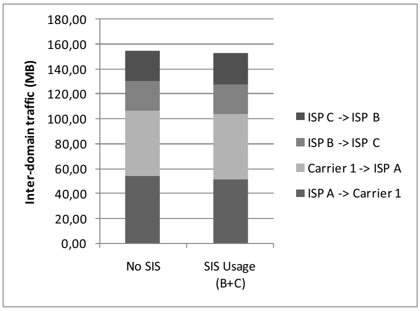

Figure 4. Inter-domain traffic

As shown, the usage of the BGP-Loc ETM can lead to a win/non-lose scenario, where clients download the content without any degradation in the service and the operator attains a reduction of the inter-domain traffic from those links that are more expensive (those ones related to transit interconnection agreement) in the short scenario, *e.g.*, the reduction of the inter-domain traffic in the transit link in the download direction (that can be represented in the traffic from the ISP A to Carrier 1) is around 3%.

Additional tests were conducted to evaluate the impact of this mechanism in other scenarios. *E.g.*, if the domain that implements locality has more seeds, then the download time is reduced significantly. Therefore, the BGP-Loc ETM can benefit from other mechanisms that: a) either provide incentives to seeds or b) deploy additional components that provide more upload capacity, such as the IoP or the HAP ETM mechanism. This is an observation that should be specifically considered if the ISP deploying the SIS offers low capacity access links to its clients. In this case, the usage of the pure BGP-Loc ETM may lead to degradation of the performance for end users, as they benefit by uploading content from external seeds with a higher upload bandwidth rather than from local ones that have slower connections.

These practical and experimental results are aligned with the simulation results presented in [14], where the BGP Loc mechanism was demonstrated as an effective procedure to reduce intra-AS traffic and peering traffic in scenarios with larger swarm sizes and varying peer distributions per domain, thus eliminating the inevitable limitations of the test-bed experiments.

6. Economic Implications

As introduced in Section 2, the presented ETM mechanisms aim at achieving a TripleWin situation. This is attained, with a simplifying discussion, by focusing mostly on incentives of the ISP and users. Note that these incentives of overlay providers are in general compatible with those of users, in the sense that an overlay provider can be reasonably assumed to benefit also whenever users of his application are better off. Also, the SmoothIT approach does not take for granted a collaboration of overlay providers (by accepting a modification of the tracker), as appears to be the case in [3] and [6].

In the case of the locality promotion mechanism (BGP-Loc), on one hand, the incentive of the ISP addressed is the reduction of the inter-domain traffic, and thus of the resulting interconnection costs. On the other hand, regarding the user, the incentive addressed is the improvement of performance. This pair of objectives also applies to other locality mechanisms, such as for [3], [4], [5], and [6]. Locality is in general a “win” for the ISP under the criterion of inter-domain traffic and costs. This is apparent from results of these aforementioned approaches, and confirmed in the overall assessment of locality in [1]. Moreover, locality is often a “win” for the user, or at least a “no lose”; (i.e. maintains performance, without improvement), but not always, as mentioned in [1]. This is due to the fact that certain local peers may have less upload bandwidth to offer than the remote ones that would be discovered by the original version of the application, particularly in cases where a bandwidth criterion is adopted for the selections made at the application level.

Therefore, if locality is not imposed in a compulsory way, it is likely that certain users can do better by not adopting it. This depends on the swarm distribution among ISPs and access resources of various users. Thus, the benefit of locality may vary in time, even for the same user within a certain swarm, since swarms are highly dynamic. However, with the SmoothIT approach, locality is offered to the user as an option. Hence, users have the possibility to verify that the mechanism is to their benefit or at least not harmful. That is, a user can ignore from time-to-time the recommendation given by the SIS and examine, whether such a bypassing of the mechanism leads to any noticeable improvement or deterioration of the QoE (Quality-of-Experience). To the best of our understanding, among the locality mechanisms of the aforementioned articles, only those of [4] and [5] can function in a contestable way as SmoothIT mechanism does. Moreover, to this end also, it is more meaningful for assessment of BGP-Loc to mostly focus on situations where it is employed only by those users that do benefit.

Under the other two ETM mechanisms, the ISP-owned Peer (IoP) and the Highly Active Peer (HAP), the ISP tries to be "by definition" compatible with selections performed by the overlay optimization mechanism: This is achieved by putting in place more resources in his own network, thus making the IoP(s) or the HAP(s) more attractive. With the right policy for the utilization of these additional resources, peers that reside in this network prefer uploading from IoP(s) or the HAP(s). Therefore, these approaches can attain both a "win" for the user regarding performance incentives and traffic localization for the ISP. Whether this situation is indeed beneficial for the ISP depends on the trade-off between the cost for these additional resources and the reduction of the inter-connection costs. Note also that an ISP employing such an approach may have a longer-term benefit, because it may improve its reputation and thus increases its customer basis and revenues. Such cost-benefit analysis is left for future work in the assessment of these approaches.

Regarding the "coordination" of traffic localization decisions between different ISPs, the inter-domain collaboration mechanism (InterSIS) provides a new communication "channel" for ISPs to resolve any information asymmetry. ISPs have an incentive to collaborate so that they can be sure that their decisions will have the desired positive impact on their inter-domain costs, or even be more effective than in the absence of this collaboration. For example, under BGP-Loc, the existence of asymmetric routes may result in unexpected increase of certain inter-domain costs, in the case of a mutli-homed ISP. The collaboration between source and destination ISPs may identify and resolve any such side-effects, leading to a "win" situation for both ISPs. Of course, the effectiveness of such a collaboration heavily depends on the truthfulness of the ISPs when reporting routing information. It is plausible that in certain cases an ISP may not want to fully reveal his routing paths or any other routing-sensitive information that depends on his business relationships with other ISPs. Compared to the scenario with BGP-Loc only in place, such cases may result in a "no-lose" situation for the ISPs, in the sense that they can collaborate by providing information up to the level that their business position permits..

7. Summary and Future Work

In this work, a set of ETM mechanisms have been outlined to show various alternatives that exist in order to promote localization of overlay traffic compared to existing work in the literature. More specifically, this paper described three different mechanisms that promote (directly or indirectly) the localization of overlay traffic, namely the *BGP-Loc*, the *IoP*, and the *HAP* mechanisms. Furthermore, the ETMS system developed has been presented, which can accommodate these mechanisms, while highlighting how different mechanisms can be combined under the generic architectural approach. This approach is backed by preliminary results of internal test-bed trials and referred to simulations undertaken, too. Finally, the work has analyzed incentives behind these mechanisms proposed and commented on the detailed difference of SmoothIT approaches with respect to relevant mechanisms proposed in the literature. The circumstances under which a TripleWin situation is achieved are highlighted, too.

Future work will include the implementation and testing of the HAP mechanism, the evaluation of the IoP in the internal trial test-bed, and the deployment of all three ETM mechanisms in a real environment to evaluate their performance under realistic circumstances and in operational networks. Moreover, SmoothIT as a project will continue simulative and theoretical evaluations of different ETM mechanisms and plans to identify the key conditions under which each of them is most effective, thus, resulting in operational guidelines for operators and end users.

Acknowledgements

This work has been performed in the framework of EU ICT Project SmoothIT (FP7-2007-ICT-216259). Furthermore, the authors would like to thank all SmoothIT project members for providing insights and their discussions, especially the University of Würzburg team for their simulation work backing the implementation results in a larger scale.

References

[1] E. Marocco, A. Fusco, I. Rimac, V. Gurbani: *Mythbustering Peer-to-peer Traffic Localization*, Internet Draft, July 2009.

[2] M. Piatek, H. Madhyastha, J. John, A. Krishnamurthy, T. Anderson: *Pitfalls for ISP-friendly P2P Design,* 8th ACM Workshop on Hot Topics in Networks (Hotnets), New York, U.S.A., October 2009.

[3] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, A. Zhang: *Improving Traffic Locality in BitTorrent via Biased Neighbor Selection,* 26th IEEE International Conference on Distributed Computing Systems, Lisboa, Portugal, July 2006.

[4] V. Aggarwal, O. Akonjang, A. Feldmann: *Improving User and ISP Experience through ISP-aided P2P Locality,* 11th IEEE Global Internet Symposium 2008 (GI '08), Phoenix, Arizona, U.S.A., 2008.

[5] D.R. Choffnes, F.E. Bustamante*: Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-peer Systems,* ACM SIGCOMM Computer Communications, Vol. 38. No. 4**,** October 2008, pp 363-374.

[6] H. Xie, Y.R. Yang, A. Krishnamurthy, Y. Liu, A. Silberschatz*: P4P: Provider Portal for Applications,* ACM SIGCOMM Computer Communications, Vol. 38, No. 4, October 2008, pp 351-362.

[7] *The SmoothIT Project,* http://www.smoothit.org/, January 2010.

[8] P. Racz, S. Soursos, M.A. Callejo, S. Spirou, F. Hecht, I. Papafili, G.D. Stamoulis, Hasan, B. Stiller: *Economic Traffic Management for Overlay Networks,* ICT-Mobile Summit 2009, Santander, Spain, June 2009.

[9] *The P2P-Next Project*: http://www.p2p-next.org/, January 2010.

[10] P. Racz, Z. Despotovic*: An ALTO Service based on BGP Routing Information,* Internet Draft, June 2009.

[11] S. Oechsner, F. Lehrieder, T. Hoßfeld, F. Metzger, K. Pussep, D. Staehle*: Pushing the Performance of Biased Neighbor Selection through Biased Unchoking,* IEEE P2P 2009, Seattle, Washington, U.S.A., September 2009.

[12] I. Papafili, S. Soursos, G.D. Stamoulis*: Improvement of BitTorrent Performance and Inter-Domain Traffic by Inserting ISP-owned Peers,* 6th International Workshop on Internet Charging and QoS Technologies (ICQT'09), Aachen, Germany, May 2009.

[13] UCSD Computer Science: *ModelNet, Systems and Networking,* https://modelnet.sysnet.ucsd.edu, January 2010.

[14] The SmoothIT Consortium: *ETM Models and Components and Theoretical Foundations (Final),* Deliverable D2.3, http://www.smoothit.org/, October 2009.