Context aware Programmable Trackers for the Next Generation Internet

Pedro Sousa

Department of Informatics
University of Minho, Braga, Portugal
pns@di.uminho.pt

Abstract. This work introduces and proposes the concept of context aware programmable trackers for the next generation Internet. The proposed solution gives ground for the development of advanced applications based on the P2P paradigm and will foster collaborative efforts among several network entities (e.g. P2P applications and ISPs). The proposed concept of context aware programmable trackers allows that several peer selection strategies might be supported by a P2P tracker entity able to improve the peer selection decisions according with pre-defined objectives and external inputs provided by specific services. The flexible, adaptive and enhanced peer selection semantics that might be achieved by the proposed solution will contribute for devising novel P2P based services and business models for the future Internet.

1 Introduction

The massive use of P2P applications in the Internet is changing the traffic profile and is introducing additional problems to Internet Service Providers (ISPs). In fact, Internet usage patterns have greatly evolved in the last years mainly due to the usage of P2P overlay networks [1], where peers form self-organized network infrastructures. ISPs are now facing serious problems such as high traffic variability and distortion and excessive and unpredictable loads in critical links. Moreover, P2P applications often generate unnecessary inter-domain traffic also making difficult the use of traditional traffic engineering techniques for network optimization [7][8]. All this leads to possible disruptions in ISPs economics and foster serious coexistence problems with new Internet applications [9]. In such context, several solutions have been adopted by ISPs in order to improve network performance such as the adoption of caching devices [11] to reduce bandwidth consumption and several mechanisms to detect and control P2P traffic [12].

Internet applications following the P2P paradigm adopt a wide range of selfish strategies to improve their performance irrespectively of the side effects induced at the network level. In this context, there is a wide variety of P2P solutions [1] having distinct behaviors, objectives, adaptation strategies, routing decisions and peering solutions [13][14]. Among many distinct P2P protocols, BitTorrent [2][4] is a good example of one of the most popular solutions [5] and is now responsible for more than one third of the Internet Traffic [6]. ISPs should

take into account the possibility that in the Future Internet applications such as P2P-based ones (e.g. BitTorrent like solutions) be mainstream and common users use them to access a wide range of services. Furthermore, P2P paradigm has several advantages for fast distribution of large content when compared with the traditional client-server model also opening new opportunities in areas as content distribution, distributed file systems, games, virtual reality, software updates, etc. As consequence, ISPs should rethink how P2P traffic will be handled in the future due to the fact that ISPs offering a poor service for P2P traffic might have to face user dissatisfaction and possible economic penalties.

In this perspective, future internet applications should no longer continuously assume selfish behaviors, as ISPs will be forced to use more restrictive solutions to control (or even deny) such traffic. This leads us to the need of devising flexible solutions guiding P2P-based applications in order to: *(i)* foster the development of collaborative optimization models able to simultaneously improve both applications and ISPs performances; *(ii)* improve the differentiation capabilities of P2P services in order to support off-line agreements between service providers and network providers involving some type of preferential treatment for specific peers and *(iii)* improve (or degrade) the quality of specific peers according with their conforming or non-conforming behaviors. This proposal assumes a BitTorrent-like P2P approach and focuses on a specific component, the context aware tracker, able to achieve all the previous mentioned objectives.

The paper is organized as follows: Section 2 presents the rationale underpinning the concept of context aware programmable trackers also providing examples of distinct selection mechanisms; Section 3 illustrates a set of experiments and corresponding results; finally, Section 4 summarizes the presented work.

2 Context aware Programmable Trackers

This section introduces the concept of context aware trackers and presents some peer selection configurations which results will be verified in Section 3.

2.1 Supporting Rationale

In order to illustrate the concept of context aware trackers lets assume the specific case of BitTorrent-like applications. Here, new peers wishing to join a specific swarm contact a tracker providing the clients with a random sample of peers. This sample is used by the peers for establishing new P2P connections with other peers in order to download a given resource[1]. In this context, the concept of context aware programmable trackers is presented by the framework of Figure 1. The several modules integrating the *Programmable P2P Tracker* internal architecture are now described in detail:

- *Peer Selection Module:* This module holds the current peer selection strategy that is being used by the context aware tracker for peer selection purposes.

[1] Additional details of the BitTorrent protocol can be found in [1–3].

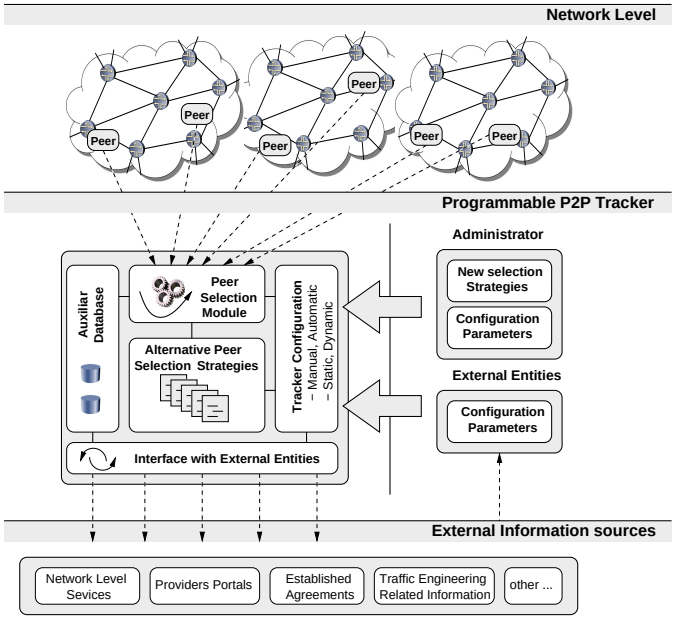


Fig. 1. Illustrative architecture for the use of Context aware Programmable Trackers.

- *Alternative Peer Selection Strategies Module:* A repository with alternative selection mechanisms to be used by the tracker. These mechanisms might be programmed, uploaded and activated using appropriate configuration commands.

- *Tracker Configuration Module:* The interface for all the allowed programming and configuration procedures. This module receives commands responsible for tasks such as: selecting the active selection strategy; uploading alternative selection strategies; defining which external entities should be contacted by the tracker to gather additional information; etc. The configuration procedures may be manual or follow an automatic approach. Moreover, tracker configurations may be static or dynamic, with the later allowing the change of the peer selection strategy during the swarm lifetime.

- *Interface with External Entities:* This module is responsible for controlling all interactions with the external information sources consulted by the tracker.

- *Auxiliar Database:* This module stores several auxiliary data characterizing active peers in a given swarm. In this way, this module is responsible for storing a large part of the information required by the tracker to support peer selection procedures. This module is also responsible to temporarily hold data retrieved from external sources before new interactions occur.

Figure 1 presents some examples of external information sources that may be used in the context of the proposed solution, e.g.: network level services able to provide privileged information regarding network state data; Provider Portals (e.g. as defined in [10]); information related with established agreements with

other providers; traffic engineering information, among other possibilities. All this information might be useful to support the tracker peer selection decisions.

2.2 Illustrative Configuration Examples

This section describes three illustrative peer selection strategies that were programmed in the context aware tracker. The first example describes a possible collaborative mechanism between the P2P and ISP network levels. The second and third examples describe tracker configurations able to differentiate the quality of the service provided to distinct peers of a specific P2P swarm.

Collaborative Optimization - This first example assumes a collaborative context between the application and the ISP levels. This particular selection mechanism is mainly devised for traffic engineering purposes, with the P2P tracker trying to reduce the inter-domain traffic generated by a given swarm. In this selection mechanism the tracker was programmed to gather information about the location of current peers in a specific swarm along with the location of newly arrived peers requesting the tracker services. Such information may be provided by network level entities cooperating with the P2P level. When receiving a request from a new peer, the tracker was programmed to return a random sample of peers in the swarm taking into account two distinct phases. First, if the swarm is in an initial state (or with a limited number of peers) then the default behavior is assumed, i.e. the return of a random sample of the existing peers to the newly arrived peer. The current number of peers in the swarm (or other P2P level information) might be used to assess the state of the swarm. Otherwise, if the swarm is not considered to be in an initial state then the returned sample will be mainly composed by peers belonging to the same networking domain of the requesting peer. This strategy intends to drastically reduce the intrer-domain traffic generated by P2P applications without noticeable degradation of the service quality. The first phase of this mechanism allows that diverse peering relations occur independently of peers locations. From that point on, newly arrived peers will mainly use local peers to download the network resource.

Service Differentiation *i) Penalizing Peers in a Swarm* - This example illustrates a peer selection mode that might be used as a pure penalizing mechanism able to punish non-conforming peers with some pre-defined P2P application level rules or, due to specific agreements with ISPs, punish peers which behavior is degrading the overall performance of the system. However, other scenarios may also benefit from these differentiation capabilities, such as: the need of controlling the traffic generated by a set of peers; protecting specific paths of the network from excessive P2P traffic to avoid the congestion of critical links; the need of forcing P2P connections only among a specific set of peers, among many others. This illustrative selection mechanism resorts to a tracker which is programmed in order to restrict the number of peers returned to specific clients[2]. This simple differentiation technique is expected to originate distinct levels of service quality as now low priority peers will have a reduced opportunity of

[2] More details of this illustrative peer selection mechanism are provided in Section 3.

discovering and connecting to other peers in the swarm. As consequence, and comparatively with peer samples having a higher dimension, such swarm elements are expected to experience lower quality service levels. Moreover, in order to prevent service starvation, these strict peer restrictions affecting low priority peers might be gradually relieved by the tracker during the swarm lifetime.

Service Differentiation *ii) Benefiting Peers in a Swarm* - A tracker operating under this programming mode is able to provide incentives to specific peers in a given swarm. In this case, such incentives are provided through a careful selection of the peers included in the samples returned by the tracker. To exemplify a selection mechanism of this type Section 3 will show the results of a tracker selection mechanism that benefits a set of peers in a given swarm by providing them privileged information regarding high upload capacity seeds that are hidden from other nodes. As consequence, such set of peers will form a kind of high priority sub-swarm that is expected to receive a better overall service from the P2P application level. This incentive based selection mechanism can be used simply to benefit specific peers in the network or with other side-effects in mind such as: divert traffic from specific links or paths of the infra-structure; avoid the generation of inter-domain traffic by providing high quality local peers in the samples; to allow the creation of enhanced sub-swarms where a restricted set of peers has access to high upload capacity seeds, among many others possibilities.

3 Experiments and Results

The ns-2 [17] simulator was used to develop and test the proposed architecture presented in Section 2, following a packet-level simulation approach. Although P2P packet-level simulations are more complex and require more computational power than flow-level approaches [15], they present more accurate results also taking into account specific cross-layer interactions which are crucial in the context of this research work. A simulation patch implementing a BitTorrent-like protocol was used for the development of the proposed solution [16]. This patch was extended in order to allow the definition of distinct peer selection techniques to be adopted by the tracker. The tracker internal structure was also modified according with the architecture proposed by Figure 1. Additional state information storage for peer selection decisions and tracker communication interfaces were also developed allowing the interactions with external entities. Several debugging and log functionalities were also integrated in the tracker.

3.1 Simulation Scenarios

Figure 2 illustrates one of the network topologies used to present some illustrative results of the proposed context aware programmable tracker concept. At the top level the network is divided in three distinct areas interconnected by inter-area links. Each area is then composed by a second level of links which configurations allow the definition of each area internal structure. In Figure 2 the concept of an area may have two distinct interpretations. For instance, when testing selection

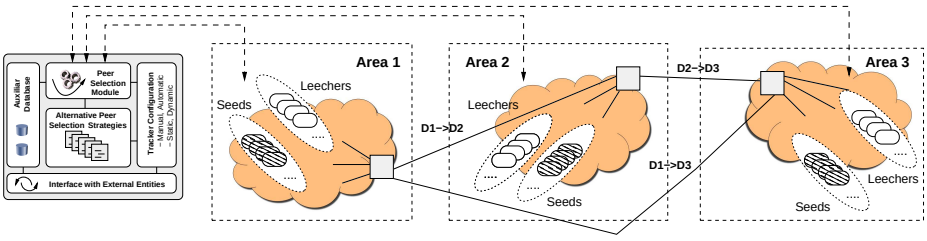


Fig. 2. Network level simulated topology for testing the programmable tracker.

mechanisms having the objective of reducing the inter-domain traffic an area will be assumed in fact as a networking domain. So, links $D1 \rightarrow D2, D1 \rightarrow D3$ and $D2 \rightarrow D3$ will be viewed as interconnections between distinct networking domains. Otherwise, for simulations disregarding domain related issues, the three areas will be interpreted as integrating an unique domain. In such cases, intra and inter area links will be viewed in fact as internal links of a domain, and their distinct capacities and propagation delays will be used to increase the heterogeneity of the domain topology. Most of the parameters controlling the BitTorrent-like protocol may be configured, including parameters such as the number of seeds and leechers per domain and their arrival processes, tracker related configurations, the use (or not) of superseeding, chunk size, file size, several timers and intervals guiding the P2P protocol, among many others. The proposed context aware P2P tracker was tested resorting to a high number of simulation experiments and each one of the individual scenarios was tested several times using distinct seeds controlling parameters such as links propagation delays, leechers arrival times to the swarm, peers distribution patterns, etc.

Due to space constraints only a set of illustrative results obtained from a specific configuration will be presented. In the selected examples most of the results were taken from a simulation scenario assuming nearly 100 leechers per area, resulting in a total number of 300 peers. The file size is 50 MB and the chunk size 256 KB. The maximum number of peer addresses requested from the tracker is 25, however depending on the selected mechanism the tracker may manipulate this value for specific peers. Most of the selected results assume the worst case scenario for file dissemination, i.e. initially only one seed and a number of leechers in the network (i.e. the flash crowd effect)[3]. Whenever possible superseeding mode was used in the simulations. At the network level the peers have, on average, an upload capacity of 1 Mbps and a download capacity which is considered to be eight timer higher than this value (i.e to simulate asymmetric access links, such as ADSL for home users). In order to improve the heterogeneity of each area, the propagation delays of the access links were randomly generated in the interval of 1-50 ms. The inter-area links were considered to be able to support a share of 10 Mbps for P2P traffic and their propagations delays are at least two times higher than the maximum value considered for intra-area

[3] To assure fairness among all peers some scenarios also assume one seed per area.

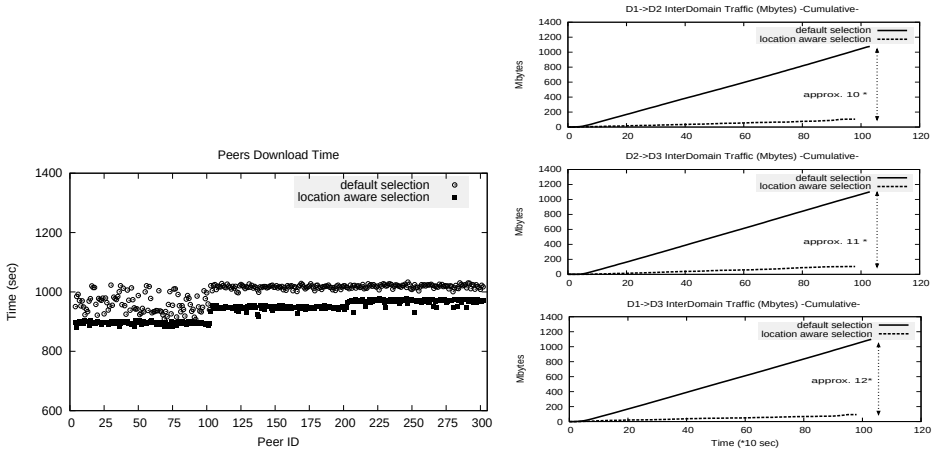


Fig. 3. Collaborative optimization: download times and inter-domain traffic (1 seed).

links. The peers performance is measured taking into account the download time needed for a complete file transfer. To simplify results visualization each peer is assigned with a $peer_{id}$ identification, in this case in the interval $[1, 300]$.

3.2 Simulation Results

This section presents simulation results of the tracker configured with illustrative selection mechanisms explained and discussed in Section 2.

Collaborative Optimization - As explained, in this example the tracker was programmed to behave in a collaborative perspective, receiving peer location information from the network level with the objective of reducing the interdomain traffic generated by the P2P level. Figure 3 shows comparative results of the tracker configured in the default selection mode and when programmed with the location aware peer selection mechanism, in this case only one initial seed is considered to exist in network area one. As observed, when the tracker is programmed to perform a location aware peer selection strategy the inter-domain traffic generated is at least ten times lower than the observed in the default selection mode (see the three graphs plotting the generated inter-domain traffic in the right side of Figure 3). Moreover, and even taking into account that peer selection decisions are now constrained and local peers have a higher probability of being selected, the average download times of the peers are also improved (see the peers download times in the left side of Figure 3). This means that it is possible to develop collaborative approaches effectively attending both P2P and ISP particular objectives. A second example is given by Figure 4 where one seed is considered to be present in each network area. Similar conclusions might be drawn, with an even higher decrease in the inter-domain traffic (e.g. $D1 \rightarrow D2$ inter-domain traffic is eighteen times lower comparatively with the default selection mode). In contrast with the previous case presented in Figure

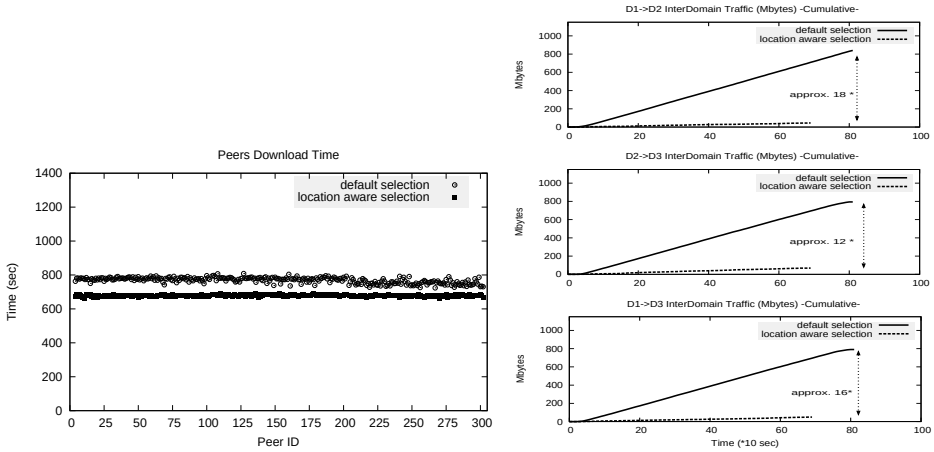


Fig. 4. Collaborative optimization: download times and inter-domain traffic (3 seeds).

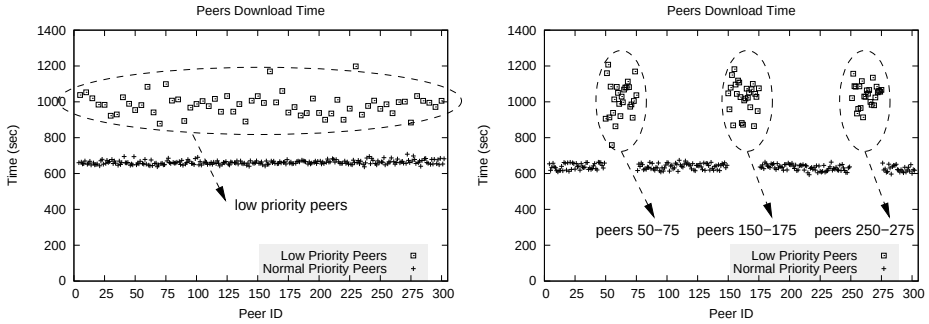


Fig. 5. Penalized peers with a) peers IDs multiple of 5 b) peers IDs in the intervals $[50, 75]$, $[150, 175]$ and $[250, 275]$.

3, where peers inside area one have slightly lower download times, now all the peers experience similar performances due to the fairness of seeds distribution.

Penalizing Peers in a Swarm - The results presented in Figure 5 a) and b) show two distinct scenarios with the tracker programmed to penalized specific peers in a swarm by restricting the number of peers included in the samples, as explained in Section 2.2. In this specific case, for penalized peers, the number of peers returned in the samples halves the maximum number of active connections allowed in the swarm. By this way, those low priority peers should experience a lower service quality level as they are constrained in the way they are able to establish P2P connections to get all the pieces of the original shared file. In the case of Figure 5 a) several peers in all the networking areas were penalized and, in this specific scenario, to make easier the results visualization penalized peers are those having a peer ID which is multiple of 5. In the second

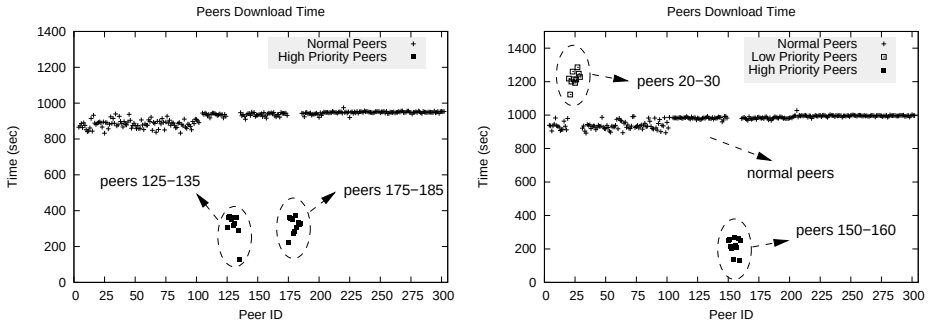


Fig. 6. a) Benefited peers in $[125, 135]$ and $[175, 185]$ b) Mixed configurations: penalized peers in $[20, 30]$ and benefited peers in $[150, 160]$.

scenario, which results are plotted in Figure 5 b), specific peer groups within each network area were selected to be penalized. In this case peer IDs in the intervals $[50, 75]$, $[150, 175]$ and $[250, 275]$ were penalized. As observed, in both scenarios the tracker was able to induce the proposed differentiation semantics.

Benefiting Peers in a Swarm - Figure 6 a) shows the results obtained using a programmable tracker configured to benefit two groups of peers, in this case belonging to the intervals $[125, 135]$ and $[175, 185]$. In this case, the strategy adopted by the tracker is to include in the returned samples two high upload capacity seeds that are unaccessible to other peers in the swarm. Additionally, the tracker assures that peers in the mentioned intervals and those seeds form a kind of high priority sub-swarm, i.e. they exchange data apart from the other peers of the swarm. As consequence, and as plotted by Figure 6 a), service differentiation is effectively achieved and high priority peers effectively obtain a better service quality from the P2P level (i.e. lower download times).

Mixed Configurations - The last selected example uses the tracker configured in a hybrid differentiation mode. The results of Figure 6 b) were obtained with the tracker programmed to benefit a specific group of peers in the network area two, in this case peers in the interval $[150, 160]$, and to penalize a group of peers in the network area one, in this case peers in the interval $[20, 30]$. As observed in Figure 6 b), the results clearly show the correctness of the devised hybrid mode, showing that mixed and enhanced configurations are possible to be achieved using the proposed context aware programmable tracker architecture.

4 Summary

This paper introduced and explained the concept of context aware programmable trackers. A detailed description of an architecture devised for that purpose was firstly presented and discussed, along with illustrative examples of possible peer selection strategies which may use additional information provided by external entities. The proposed context aware tracker solution was implement resorting to

simulation and, as corroborated by the presented results, enhanced collaborative behaviors and differentiation semantics are possible to be achieved at the P2P level using the proposed solution. In this way, the proposed approach will benefit the development of advanced P2P-based applications in the future internet, also underpinning the development of intelligent collaborative approaches between ISPs and the P2P applicational level. Moreover, due to the enhanced differentiation semantics that could be obtained using the proposed context-aware tracker, novel Internet services and business models based on the P2P paradigm could also take advantage of the proposed solution.

References

1. Keong Lua et al. *A survey and comparison of peer-to-peer overlay network schemes.* IEEE Communications Surveys & Tutorials, pp. 72-93, 2005.
2. B. Choen, *Incentives build robustness in BitTorrent*, In Proc. 1st Workshop on Economics of Peer-to-Peer Systems, Berkeley, June 2003.
3. A. Legout et al. *Clustering and Sharing Incentives in BitTorrent Systems.* In Proceedings of ACM SIGMETRICS'2007, June 12-16, 2007, San Diego, CA, USA.
4. Ashwin R. Bharambe et al. *Analyzing and Improving a BitTorrent Networks Performance Mechanisms*, IEEE INFOCOM, 2006.
5. T. Karagiannis et al. *Is p2p dying or just hiding?* Globecom, Dallas USA, Nov. 2004.
6. Hendrik Schulze and Klaus Mochalski, *Internet Study 2007: The Impact of P2P File Sharing, Voice over IP, Skype, Joost, Instant Messaging, One-Click Hosting and Media Streaming such as YouTube on the Internet*, Tech. report, 2007.
7. R. Keralapura et al. *Can ISPs take the heat from overlay networks?*, In Proc. of HotNets-III, San Diego, CA, Nov. 2004.
8. L. Qiu et al. *SelFIsh routing in Internet-like environments.* In Proc of SIGCOMM 2003, Karlsruhe, Germany, Aug. 2003.
9. H. Xie et al. *P4P: explicit communications for cooperative control between P2P and network providers*, http://www.dcia.info/documents/P4P_Overview.pdf
10. H. Xie, et al. *P4P: Provider Portal for Applications*, SIGCOMM 2008, August 1722, Seattle, Washington, USA, 2008.
11. G. Shen, et al. *HPTP: Relieving the tension between ISPs and P2P.* In Proc of IPTPS, Bellevue, WA, Feb. 2007.
12. A. Spognardi et al. *A Methodology for P2P File-Sharing Traffic Detection, Proc. Second International Workshop on Hot Topics in Peer-to-Peer Systems 2005 (HOT-P2P 2005)*, pp. 52- 61, July 2005.
13. T. Karagiannis et al. *Should internet service providers fear peer-assisted content distribution?*, Proc. of the Internet Measurement Conf., Berkeley CA, Oct. 2005.
14. H. V. Madhyastha et al. *iPlane: An information plane for distributed services.* In Proc. of OSDI, Seattle, WA, 2006.
15. Kolja Eger et al. *Efficient Simulation of Large-Scale P2P Networks: Packet-level vs. Flow-level Simulations*, 2nd Workshop on the Use of P2P, GRID and Agents for the Development of Content Networks, Monterey Bay, USA, 2007.
16. Simulation of BitTorrent Peer-to-Peer (P2P) Networks in ns-2. http://www.tu-harburg.de/et6/research/bittorrentsim/index.html
17. ns-2 (Network Simulator). Documentation from http://www.isi.edu/nsnam/ns/.
18. A. Odlyzko. *Data networks are lightly utilized, and will stay that way.* Review of Network Economics, 2(3):210−237, September 2003.