

Flexible Peer Selection Mechanisms for Future Internet Applications

Pedro Sousa

Informatics Department, University of Minho, Braga, Portugal

Email: pns@di.uminho.pt

Abstract—This paper focuses on the development of solutions able to enhance the peer selection semantics supported by P2P based applications. The proposed framework relies on the concept of a flexible P2P tracker able to support multiple versatile peer selection procedures that might be activated and configured during the application lifetime. The devised flexible tracker may also resort to additional cross-layer informations provided by the network level or other external entities. Using a BitTorrent-like P2P approach as a case study, the proposed solution fosters the capabilities of P2P applications to differentiate, benefit or penalize specific peers, also giving ground for the development of advanced applications based on the P2P paradigm.

I. INTRODUCTION

Internet usage patterns have greatly evolved in the last years. One example of that is the increasing usage of P2P overlay networks [1], where peers form self-organized network infrastructures. Among many distinct P2P protocols, BitTorrent [2] is an example of one of the most popular solutions [3] and is now responsible for more than one third of the Internet Traffic [4]. Thus, the massive use of P2P applications is an undeniable reality and is dramatically changing the traffic profile and opening new opportunities for application development.

As consequence, ISPs are now facing serious problems such as: high traffic variability and distortion, unpredictable loads in crucial links, generation of unnecessary inter-domain traffic and difficulties to use traditional traffic engineering techniques for network optimization [5][6]. All this leads to possible disruptions in ISPs economics resulting in serious coexistence problems with new Internet applications [7]. To overcome such problems, several solutions have been adopted in order to improve network and applications performance. As brief examples, ISPs often resort to the adoption of caching devices [9] to reduce bandwidth consumption and to several mechanisms to detect and control P2P traffic [10]. In addition, also the application level is fostering the battle for efficient storage and dissemination of several contents. In fact, there are a wide variety of P2P solutions [1] each one having distinct behaviors, objectives, adaptation strategies, routing and peering solutions. The proposed solution will shift away from these selfish and independent optimization perspectives and will be centered on the development of novel alternative paradigms. In this way, using a BitTorrent-like P2P approach as a case study, this work proposes a framework able to enrich the decisions adopted by P2P applications, namely in configuration tasks such as peer selection procedures. This will allow the development of collaborative models able to simultaneously improve both

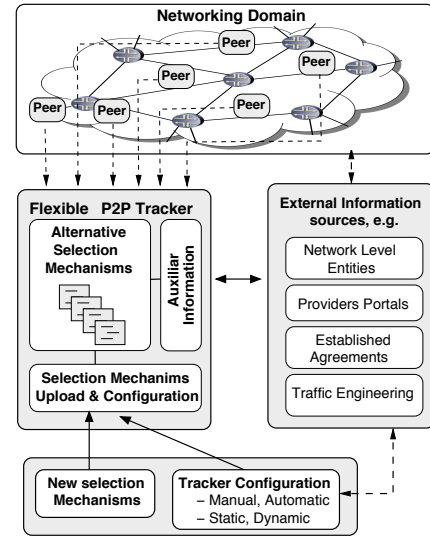


Fig. 1. An illustrative framework for flexible peer selection decisions.

applications and ISP performances. The proposed framework also intends to aid the development of collaborative solutions to foster the quality and differentiation capabilities of the P2P services supported by the network provider. Such capabilities are essential to support off-line agreements between service providers and network providers involving some type of preferential treatment for specific peers in the network, or for the development of mechanisms able to improve (or degrade) the quality of some peers according to specific P2P application objectives or other possible collaborative agreements.

II. FLEXIBLE PEER SELECTION FRAMEWORK

The presented framework assumes the specific case of P2P applications following a BitTorrent-like approach with a centralized tracker. Similar concepts to the ones presented by this research work could be devised and adapted to other P2P approaches (including other BitTorrent variants) taking into account their specificities and operation modes. Figure 1 presents a general framework allowing flexible peer selection decisions. The proposed framework integrates the following elements: (i) a flexible P2P tracker; (ii) the networking domain; (iii) a set of external information sources and (iv) a configuration module. The Flexible P2P Tracker is the core of the

proposed system. As proposed by BitTorrent-like approaches this element tracks a specific swarm and is contacted by newly arrived peers wishing to participate in a given session. The tracker should return a peer sample to the client which will be then used for establishing P2P connections with other peers of the swarm. In the proposed approach this element integrates a programmable area where several peer selection mechanisms will reside. These mechanisms are programmed and uploaded by the tracker administrator in order to induce specific peer differentiation behaviors in the network. The administrator (or other external entity) will configure the selection mechanism to be used by the tracker within a specific P2P session. As mentioned, the framework also allows that such configuration procedures could be performed by other external entities. In such way, the tracker configuration procedures may assume a manual or automatic perspective and the selected configurations may be static or modified on-the-fly during the swarm lifetime. Each one of the selection mechanisms programmed in the tracker may also require specific information from external information sources. Figure 1 includes examples of possible information sources, such as: *i)* network level entities able to provide privileged network level information; *ii)* Provider Portals (e.g. as defined in [8]); *iii)* general information related with established agreements with other network or service providers requiring particular differentiation semantics at the P2P level and *iv)* traffic engineering related information. Next, three illustrative selection mechanisms that were programmed in the flexible P2P tracker are presented.

A. Topology aware peer selection

In this configuration mode the P2P tracker will try to reduce the inter-domain traffic generated by a given P2P session. In order to prevent the generation of unnecessary inter-domain traffic in this configuration mode the network level (e.g. entities such as Provider Portals suggested in [8]) will inform the P2P level about the administrative location of current peers in a specific swarm along with the location of newly arrived peers requesting the tracker services. Such location informations will influence the tracker peer selection behavior. In this configuration mode, when receiving a request from a new peer, the tracker will return a random sample of peers in the swarm taking into account two distinct phases: 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 random sample of existing peers to the newly arrived peer, the current number of peers in the swarm (or other P2P level information) and the definition of a specific threshold might be used to assess the state of the swarm; otherwise, if the swarm is not considered to be in a initial state then the returned sample will be mainly composed by peers belonging to the same networking domain of the requesting peer.

B. Restricted peer selection

Another possible family of selection mechanisms is based on restricting the number of the peers returned in a given sample. This simple differentiation technique is expected to

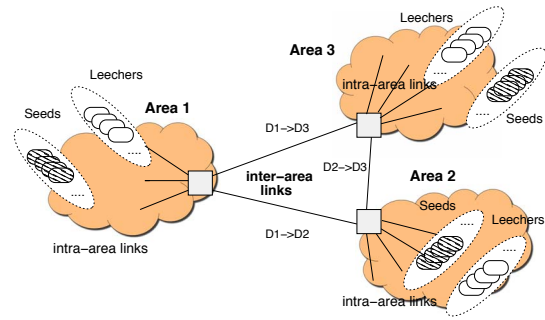


Fig. 2. Network level simulation scenario used for results presentation.

originate distinct levels of service quality as now low priority peers will have a reduce opportunity of discovering and connecting to other peers in the swarm. This selection mode might be used as a pure penalizing mechanism where specific peers contact the tracker and receive a random sample with a reduced number of peers. As consequence, and comparatively with peer samples having an higher dimension, such swarm elements are expected to experience low quality service levels. This penalizing mechanism may be useful to punish non-conforming peers with some pre-defined P2P level rules or, due to agreements with the network level, punish peers which behavior is degrading the overall performance of the system. In order to prevent service starvation, these strict restrictions affecting low priority peers might be gradually relieved during the swarm operation. However, it is important to note that such type of selection mechanisms may also be used with other objectives, such as controlling the traffic generated by a set of peers, protecting specific paths of the network from excessive P2P traffic, avoiding the congestion of critical links, force P2P connections among a specific set of peers, etc..

C. Incentive based peer selection

This type of selection mechanisms aims to provide incentives to specific peers in a given swarm. In the same perspective of the previously explained type of mechanisms, it is possible to provide incentives through a careful selection of the peers included in the samples returned by the tracker. Once again, incentive based selection mechanisms 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; allow the creation of enhanced sub-swarms where a restricted set of peers have access to high upload capacity seeds, among many others possibilities. As an example of a selection mechanism of this type Section III will show the results of a tracker based 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 peers.

III. EXPERIMENTS AND RESULTS

In order to test the peer selection framework the ns-2 [13] simulator was used to develop a prototype of the architecture

explained in Section II. The framework and the illustrative peer selection mechanisms were tested following a packet-level simulation approach. As discussed in [11], although packet-level simulations are more complex and require more computational power than flow-level approaches 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 protocol was used for the development of the framework [12]. This patch was extended in order to allow the definition of distinct peer selection techniques to be adopted by the tracker. The internal structure of the tracker was also modified in order to store additional state information supporting the peer selection decisions. Moreover, additional interfaces were developed allowing the exchange of information between the flexible tracker and other external entities. Several debugging and log functionalities were also integrated in the tracker. Figure 2 illustrates one of the network topologies used to test the proposed peer selection framework. 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 the internal network structure. In Figure 2, and in the context of the simulation results presented here, the concept of an area may have two distinct interpretations. As an example, when testing selection 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 a single networking domain. In such cases intra and inter area links will be viewed in fact as internal links of a domain. As regards to the P2P applicational level, most of the parameters controlling the BitTorrent-like protocol might also be configured by the user. This includes 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 protocol, among many others. Due to space constraints only a set of results obtained from a specific configuration will be presented in the following sections in order to illustrate the differentiation semantics obtained by the proposed framework. In the selected examples most of the results were taken from a simulation scenario assuming nearly 50 leechers per area, resulting in a total number of 150 peers. The file size is 100 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. Whenever possible super-seeding 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 times higher than this value. 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




Fig. 3. Default vs Location aware peer selection - Download times.

support a share of 10 Mbps for P2P traffic and their propagation delays are at least two times higher than the maximum value considered for intra-area links. The peers performance is measured taking into account the download time needed for a complete file transfer and for results visualization each peer is assigned with an identification number in the interval $[1, 150]$.

A. Simulation Results

The next sections present the results of the tracker configured with the illustrative mechanisms discussed in Section II.

1) *Topology aware peer selection*: Figures 3 and 4 present the results of the location aware peer selection mechanism compared with the ones obtained using the default tracker peer selection procedures. Figure 3 plots the file download time experienced by each peer and Figure 4 show the total amount of inter-domain traffic exchanged during the swarm lifetime. For the selected scenario, and as observed in Figure 3, both methods achieve downloading times very similar without noticeable differences. This is explained by the rationale supporting the location aware selection mechanism that was previously described. After an initial stabilization period, peers adopting the location aware strategy will be provided with peer samples preferentially containing elements from the same domain. If this, at least at a first sight, may restrict the piece selection opportunities for those peers, it is also true that the option of choosing near peer locations leads to a high probability of establishing lower RTT connections which benefits the throughput of TCP based BitTorrent connections. A detailed analysis of Figure 4 shows that the use of a flexible tracker running the location aware selection strategy will effectively reduce the amount of inter-domain traffic. As observed in Figure 4, using the proposed mechanism the inter-domain traffic is now five (link $D1-D3$) to ten (link $D2-D3$) times lower than the observed in the default configuration.

2) *Restricted peer selection*: Figures 5 a) and b) show the results for a BitTorrent scenario having a flexible tracker configured with the restricted peer selection mechanism explained in Section II. In this case the tracker decides to penalize a specific set of peers by restricting the number of candidate peers included in the returned sample. In this case the number of returned peers in the samples halves the maximum number of active connection allowed in the swarm. By this way, those peers should experience low service quality levels as they are




Fig. 4. Default vs Location aware selection - Total amount of Inter-domain traffic for: a) $D1 - D2$ link; b) $D1 - D3$ link and c) $D2 - D3$ link.




Fig. 5. a) Restricted peer selection - penalized peers = $\{10, 20, 30, 40, 50, 60, 70, 80, 90, 100, 110, 120, 130, 140, 150\}$; b) Restricted peer selection - penalized peers in $[10, 20]$ and $[110, 120]$; c) Incentive based peer selection - benefitted peers is a subset from area two, i.e. peers in $[70, 80]$.

limited in establishing several P2P connections. In the case of Figure 5 a) distinct peers from distinct network areas were penalized by the tracker (to make easy the perception of the results the applied rule was: $peer_{id}$ is multiple of ten). As observed, it is clear the service degradation affecting those peers and, as consequence of a discriminative penalization, their download times are significantly higher than the obtained by other peers. Another example of this selection mode is presented in Figure 5 b) also presenting clearly visible differentiation results, with tracker receiving information regarding two sets of nonconforming peers (sets $[10, 20]$ and $[110, 120]$).

3) *Incentive based peer selection*: The example selected in this section assumes that the tracker will benefit a specific set of peers in networking area two, in this case $peers_{id}$ in the interval $[70, 80]$. For that purpose, in the selected example, the tacker knows the location of two high upload capacity seeds. Those seeds are hidden from normal peers and will be only included in the returned peer sample when the requesting peer belongs to the previously mentioned interval. In addition, the tracker manipulates the returned samples in order to form a kind of high priority sub-swarm only integrating nodes in the interval $[70, 80]$ and the two hidden seeds. As consequence, and as observed in Figure 5 c), the overall service experienced by those peers is better than the obtained by normal peers, which means in this case lower downloading times.

IV. SUMMARY

This paper proposes the concept of flexible trackers to be used in the context of P2P applications. Their ability to support multiple versatile peer selection strategies turns possible that a wide range of differentiation semantics be introduced at the

P2P application level. The use of the devised framework allows that P2P level and ISPs be able to develop collaborative efforts and that such enhanced differentiation capabilities could be used as a starting point for devising new services and business models based on the P2P paradigm for the Future Internet.

REFERENCES

- [1] Keong Lua et al. *A survey and comparison of peer-to-peer overlay network schemes*. IEEE Comm. 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] T. Karagiannis et al. *Is p2p dying or just hiding?* In Globecom, Dallas, TX, USA, November 2004.
- [4] 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.
- [5] R. Keralapura et al. *Can ISPs take the heat from overlay networks?*, In Proc. of HotNets-III, San Diego, CA, Nov. 2004.
- [6] L. Qiu et al. *Selfish routing in Internet-like environments*. In Proc. of SIGCOMM, Karlsruhe, Germany, Aug. 2003.
- [7] H. Xie, A. Krishnamurthy, A. Silberschatz, and Y. R. Yang, *P4P: explicit communications for cooperative control between P2P and network providers*, http://www.dcia.info/documents/P4P_Overview.pdf
- [8] H. Xie, et al. *P4P: Provider Portal for Applications*, SIGCOMM 2008, August 1722, Seattle, Washington, USA, 2008.
- [9] G. Shen et al. *HPTP: Relieving the tension between ISPs and P2P*. In Proc of IPTPS, Feb. 2007.
- [10] 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.
- [11] 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, 2007.
- [12] Simulation of BitTorrent Peer-to-Peer (P2P) Networks in ns-2. <http://www.tu-harburg.de/et6/research/bittorrentsim/index.html>
- [13] ns-2 (The Network Simulator). Sources and Documentation from <http://www.isi.edu/nsnam/ns/>.