

Traffic Engineering Approaches in P2P Environments

Pedro Sousa

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

Abstract. Nowadays, P2P applications proliferate in the Internet with distinct utilization contexts, being also an attractive model for the deployment of advanced Internet services. However, there are several undesirable effects that are caused by such applications, raising coexistence problems with Internet Service Providers (ISPs). In this context, using as case study BitTorrent like applications, this work explores a collaborative framework allowing that advanced efforts could be carried on between P2P applications and network level entities. In order to illustrate such framework, several Traffic Engineering (TE) mechanisms are devised in order to align some P2P dynamics with particular objectives pursued by network administrators. The simulation results show that both the proposed framework and corresponding illustrative mechanisms are viable and can effectively foster future research efforts within this field.

1 Introduction

The nature of Internet applications has greatly evolved in the last years and there is an increasing usage of P2P overlay networks [1], where peers form self-organized network infrastructures. Within this class of applications, BitTorrent [2][3] is a common example of one of the most popular solutions [7], being responsible for a large amount of the total Internet traffic [8].

Thus, the massive use of P2P applications and their inherent operating models opened new application opportunities in areas as content distribution, distributed file systems, games, virtual reality, software updates, etc. However, it is also true that ISPs are facing serious coexistence problems with the P2P operational paradigm. In fact, P2P usually generates high variability and distortion in traffic patterns, along with excessive and unpredictable loads in crucial links. Moreover, P2P behaviors many times preclude the use of classical Traffic Engineering (TE) techniques for network optimization [9][10]. As consequence, this results in several coexistence problems between network providers and P2P based applications [12]. In order to deal with that, ISPs often resort to caching devices [15][16] to reduce bandwidth consumption, or inspection tools to detect and control P2P traffic [17]. Nevertheless, P2P applications are permanently fostering the battle to surpass some of these mechanisms and there is a wide range of P2P approaches with distinct selfish behaviors, adaptation strategies

and peering solutions [18][19]. In this perspective, this work assumes the inherent advantages of devising collaborative approaches between P2P and network level entities (e.g. ISPs). For that purpose, using a BitTorrent-like P2P approach as case study, this work proposes a framework able to enrich the decisions adopted by P2P applications, also taking into account specific requirements imposed by the underlying network level. Within this context, and resorting to a highly re-configurable P2P tracker, several illustrative TE mechanisms are explored by considering some mathematical foundations from the graph theory field. The explored models constitute preliminary approaches to deal with P2P swarms involving a high number of peers, aiming to raise the network level with some control and estimation capabilities to better accommodate such P2P traffic aggregates in the underlying infra-structure.

The paper is organized as follows: Section 2 describes the rationale of the proposed framework along with illustrative tracker configurations; Section 3 presents the experimental platform and obtained simulation results; finally, Section 4 draws the main conclusions related to the proposed solution.

2 Proposed Framework and Illustrative Configurations

In order to illustrate the proposed framework (see Figure 1) we assume the specific case of BitTorrent-like applications, considering that such system principles could be used to develop proprietary applications offered by content/service providers to their end-users. Additionally, the application scenario adopted within the proposed collaborative system assumes that the tracker is the only entity able to provide peering information. Thus, client side software provided to end-users is not able to exchange peer identities with other peers, meaning that the tracker fully controls the peering informations provided to the clients. In BitTorrent classical systems, new peers wishing to join a specific swarm contact a tracker which then provides the clients with a random sample of peers. This sample is used by the peers for establishing new P2P connections with other peers in the swarm to obtain a given shared resource. After this stage, several BitTorrent rules will drive the data transfer processes among the peers. These additional details about the BitTorrent protocol regarding pieces selection algorithms and choking strategies to determine which peers to choke/unchoke can be found in [1, 2]. In this perspective, the P2P tracker main role is to keep track of the current peers participating in a given swarm and dynamically provide random peer samples to newly arrived peers in the swarm.

The proposed architecture, depicted in Figure 1, assumes an ISP networking domain, consisting of several core routers (which in this work context may also express possible Points of Presence (PoPs) of the ISP), interconnected by several links. At the P2P application level, peers are distributed among several end-users areas, which access ISP infrastructure through the corresponding PoP. The P2P tracker is able to use alternative configurations, which could be defined by the administrator (or other external entities) and might be programmed and activated using appropriate configuration commands. Distinct configuration strate-

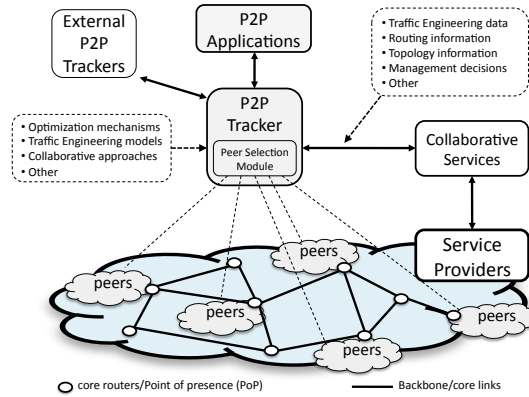


Fig. 1. Illustrative description of the proposed operational scenario

gies adopted by the tracker may require distinct types of information, depending on their objectives and operation modes. Moreover, if required, the tracker may also integrate additional intelligent optimization (e.g. [11, 14]) or forecasting (e.g. [4]) mechanisms to deal with high complex optimization problems. The external collaborative services, mainly controlled by the network providers, are expected to provide useful information for the P2P system. Examples of possible external information sources may include: network level entities/services able to provide privileged network level information (e.g. topology inputs, routing information, peers location, or other TE data); Provider Portals (e.g. as defined in [13]); administrative services providing specific management policies to the tracker; or any other entity able to interact with the tracker providing valuable information. These interoperation possibilities between the P2P applicational level and network level entities present incentives attracting both sides. Service level entities (e.g. content providers) can benefit from gathering underlying network information and network providers might also try to influence application level peering decisions in order to improve their network resources usage. In this perspective, and depending on the defined objectives, distinct advanced TE mechanisms could be implemented at the P2P tracker.

Some examples of possible configurations useful for TE purposes will be presented in the following sections. Within the context of P2P swarms involving a considerable number of peers, the objective at this stage is to explore how some simple mathematical foundations from the graph theory field can be adapted and used as underpinning models to gather preliminary snapshots of the global P2P application and corresponding traffic aggregates dynamics. If such estimation efforts could be effectively accomplished, the studied models will be then able to be enriched and improved in the future with additional modeling capabilities. The first two examples of the presented illustrative mechanisms focus on providing some preliminary estimation approaches able to achieve qualitative measures

about the impact of P2P traffic in network links, also allowing that ISPs may divert traffic from specific links of their infrastructures. The third example could be useful for other collaborating scenarios, with the tracker providing useful information regarding seeds placement within the network.

2.1 P2P Link Impact Measure

This section proposes a heuristic to be used by the tracker to estimate the impact of the traffic generated by a P2P swarm in the network links of the ISP. Such P2P link impact measures calculated by the tracker can be used by external network services or administrators to better manage their networking resources. For that purpose, we assume a collaborative scenario where the tracker can contact authorized collaborative network services in order to collect topology and routing information, along with peers location data. Based on that, there are some aspects from the graph theory field that can be adapted within this context. As example, several graph measures [5, 6] could constitute valuable inputs to be adapted in order to devise estimation techniques of P2P link impacts values.

To illustrate such concepts we assume that the tracker may resort to a network representation using a mathematical model which represents ISP nodes (routers) and transmission links by a set of nodes (N) and links (L), respectively, in a simple graph $G = (N, L)$. For simplicity, we consider network scenarios with symmetric links, which can be modeled by an undirected graph. Each pair of nodes in the graph ($x, y \in N$) is connected by a given path comprising one or more links, according to a given routing strategy adopted at the network domain (e.g. as shortest-path based mechanisms). Also, each link ($l \in L$) has specific attributes, such as an assigned weight for routing purposes, used by the ISP to compute shortest-paths among the nodes. Moreover, we also assume that the location (area) of end-users peers participating in the swarm is denoted by the corresponding ISP PoP/core router, a , with $a \in A$ and $A \subseteq N$. In order to compute the P2P link impact values, the tracker will evaluate a P2P betweenness centrality measure for each one of the links, taking into account the locations of the swarm peers (identified by the corresponding ISP router/PoP). For a particular link, l , and a specific pair of end-users areas, $i, j \in A$, the metric takes into account the ratio between the number of shortest paths from i to j , $nsp_{i,j}$, and the number of shortest paths from i to j that pass through link l , $nsp_{i,j}(l)$, resulting that link l will be assigned with a partial impact value of $\frac{nsp_{i,j}(l)}{nsp_{i,j}}$ for the particular case of i, j peering adjacencies. It is then possible to sum all the partial impact values involving link l , and obtain a reference value within the interval $[0, 1]$ by considering all the possible area peering adjacencies, i.e. $|A| \cdot (|A| - 1)$. In the case of P2P swarms where end-user areas show a considerable unbalanced distribution of the number of peers (also reflected in the number of peers from each area that are included in the random samples returned by the tracker) an additional weighting factor could be introduced, $w_{i,j}$, for each specific i, j ¹ case.

¹ $w_{i,j}$ factor considers the ratio between the number of peers involved in the area peering adjacency i, j over the total number of peers involved in all possible area peering

This will increase the importance of shortest paths connecting areas having a higher number of peers. Thus, links presenting higher betweenness centrality values have a higher probability of being traversed by traffic of the corresponding BitTorrent P2P swarm. For the case of a tracker returning random samples to the contacting peers, Equation 1 presents the normalized P2P betweenness centrality value for link l , within the interval $[0, 1]$, which is from this point on designated as P2P link Impact Measure I (IM_I).

$$IM_I(l) = \frac{\sum_{i,j \in A, i \neq j} \frac{nsp_{i,j}(l)}{nsp_{i,j}} \cdot w_{i,j}}{|A| \cdot (|A| - 1)}, l \in L \quad (1)$$

The devised P2P link impact measure can be further enhanced taking into account some common application level dynamics assumed by the BitTorrent protocol. In fact, due to the inherent characteristics of the transport protocols used by BitTorrent (e.g. TCP), peers usually have a considerable probability of establishing peering connections with nearest peers in the network, taking advantage of lower RTTs. In this perspective, in Equation 1, when considering a given shortest path between areas i and j (assuming the context of peers in area i trying get data from peers in area j) it is also possible to assign a preference value² ($p_{i \leftarrow j} \in [0, 1]$ with $\sum_{j \in A, j \neq i} p_{i \leftarrow j} = 1$) to such shortest paths, which implicitly expresses how close is area j from area i . This value is then multiplied by the number of possible external peering adjacencies that could be made by peers in area i , i.e. $|A| - 1$. The resulting value is then used as a weighting factor when accounting the shortest path between areas i and j . Equation 2 expresses an alternative P2P betweenness centrality value for link l , from this point on designated as P2P link Impact Measure II (IM_{II}). The IM_{II} or IM_I impact measures could be then announced to network services or administrators which may require the tracker to change its behavior according to a given objective.

$$IM_{II}(l) = \frac{\sum_{i,j \in A, i \neq j} [(|A| - 1) \cdot p_{i \leftarrow j}] \cdot \frac{nsp_{i,j}(l)}{nsp_{i,j}} \cdot w_{i,j}}{|A| \cdot (|A| - 1)}, l \in L \quad (2)$$

2.2 Protecting Links from P2P Traffic

This configuration mode allows that the tracker could be configured in order to protect specific network links from excessive levels of P2P traffic. In a context of TE efforts, external network level services or administrators are now able to inform the tracker about the link(s) that it should protect, i.e. requiring that

adjacencies. In order to preserve the original form of the betweenness measure, this ratio is multiplied by $|A| \cdot (|A| - 1)$ for normalization purposes.

² If required, in highly heterogeneous scenarios, the estimation model could be enriched by also reflecting in this parameter the relative quality of the average upload capacities of area j peers, when compared with other peers in the domain.

the tracker reduces their P2P impact values. For that purpose, and taken the example of a given set of protected links, $K \subseteq L$, the tracker objective is to minimize the P2P impact values of links $k \in K$, e.g. induce peering adjacencies constrained by one of the illustrative objective functions expressed in Equation 3, depending on the adopted P2P link impact measure. To achieve this objective, the tracker should change its random behavior and carefully select which peers should integrate the peer samples returned to requesting clients.

$$\min \left(\sum_{k \in K, K \subseteq L} IM_I(k) \right) \text{ OR } \min \left(\sum_{k \in K, K \subseteq L} IM_{II}(k) \right) \quad (3)$$

The underpinning optimization concept is that the P2P tracker be able to induce peering adjacencies that should now avoid traversing network paths including the protected links. It is possible that under some peering configurations achieved by the tracker the previously presented P2P link impact equations need to be adapted in consonance with the new conditions³.

2.3 Seeds Placement Strategies

This tracker configuration example could be used when network level services, in collaboration with content providers, are intended to have an active participation in the definition of the swarm structure, namely as regards to seeds placement. In this context, and as will be illustrated in Section 3.3, this mechanism can be used with distinct objectives, such as benefit some peers areas in the swarm, or to achieve a more efficient usage of networking resources. For that purpose, the tracker should be able to provide valuable information to network administrators about the correct positioning of the seed(s) given a pre-defined criteria. Assuming that the tracker has the objective of finding the more appropriate seed locations for a given set of end-users areas Z (with $Z \subseteq A$), it is possible to resort to the notion of closeness centrality to compute the mean length of the shortest paths (lsp) between the candidate seed locations (any network node/PoP within the ISP infrastructure) and the considered areas. As before, for unbalanced distributions of the number of peers in the areas, an weighting factor could be introduced, w_i , for each specific area⁴. Equation 4 expresses then a closeness centrality measure for location n , from this point on designated as P2P Closeness Measure (CM), and candidate locations with lower CM values are expected to better serve the considered areas.

$$CM(n) = \frac{\sum_{i \in Z, Z \subseteq A} lsp_{n,i} \cdot w_i}{|Z|}, n \in N \quad (4)$$

³ As an example, if peers in some area are not able to contact peers in other specific areas, then the number of all possible area adjacencies will be no longer $|A| \cdot (|A| - 1)$ as assumed in Equations 1 and 2, for normalization purposes.

⁴ The w_i factor considers the ratio between the number of peers in area i over the total number of peers in all areas. Taking into account the original form of the closeness measure this ratio is multiplied by $|Z|$ for normalization purposes.

3 Experiments and Results

For testing purposes, the ns-2 [21] simulator was used to develop the proposed architecture and test some of the devised tracker configurations. A packet-level simulation approach was adopted for that purpose, using a simulation patch [20] implementing a BitTorrent-like protocol. This patch was extended to integrate a prototype with the major components of the framework presented in Figure 1, also including the illustrative tracker configurations previously described.




Fig. 2. Network topology used to collect illustrative results

Figure 2 illustrates the network topology adopted to present some illustrative results. At the top level the Internet Service Provider consists of several core routers (for this work context they can also be viewed as possible Points of Presence (PoPs)), interconnected by several links. For P2P application level simulation, six end-users areas with participating peers are assumed. Each area is composed by a second level of nodes/access links. Most of the parameters controlling the BitTorrent-like protocol could be configured, such as the number of seeds and leechers per domain, their arrival processes into the swarm group, tracker related configurations, the use (or not) of superseeding, chunk size, file size, among others. In the selected examples the results were taken from a simulation scenario assuming nearly 50 leechers per area, i.e. a total number of 300 peers. The file size is 50 MB and the chunk size 256 KB. The maximum size of the peer sample returned by the tracker is 25. At the end-users areas 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. 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. In this illustrative scenario, the ISP links were considered to be able to support a maximum share of 50 Mbps for P2P traffic and their propagations delays are at least two times higher than the values considered for access links. In the following sections, and for each particular tracker configuration example, five simulations (s_1, \dots, s_5) were made and the corresponding mean values taken for analysis.

3.1 P2P Link Impact Measures




Fig. 3. P2P traffic traversing each link a) on each simulation b) mean values




Fig. 4. a) P2P link impact measure IM_I b) IM_I vs observed traffic values

This section illustrates a scenario where the P2P tracker evaluates qualitative P2P impact measures for the network links of the domain. In this specific example a single seed is considered to exist on each end-user network area. The values presented in Figure 3 a) report the cumulative values of P2P traffic traversing each link, for each one of the simulation instances, with the corresponding mean values presented in Figure 3 b). As observed, the P2P traffic resulting from the swarm behavior has a major impact in some specific links of the network domain. The estimated P2P link impacts, using the IM_I technique⁵, are presented in Figure 4 a), where it is visible an acceptable match when the proportions over such link impact values are compared with the proportions among the measured traffic values. In order to provide a more straightforward comparative perception between impact values and traffic measures, in Figure 4 b) the scale of the impact values was converted to the same order of magnitude as the traffic measures. In this perspective, Figure 4 b) shows a similar trend among the link traffic values and the forecasted link impact values. This means that, even considering

⁵ For this scenario, the estimation model considers all i, j paths with $w_{i,j} = 1$.




Fig. 5. a) P2P link impact measure IM_{II} b) IM_{II} vs observed traffic values

that some distortions exist among the link impact values when compared with measured traffic (given by the plotted lines differences in Figure 4 b)), external entities or administrators can rely on trackers that use the IM_I technique to nearly forecast the qualitative impact of P2P traffic in the network domain. The estimation of the P2P link impact measures can be further enhanced using the IM_{II} method. In this context, Figure 5 a) presents the P2P link impact estimations using the IM_{II} model⁶, with the comparative values presented in Figure 5 b). As shown, a more accurate match between the link impact values and the proportions among real traffic measures is now obtained.

3.2 Protecting Links from P2P Traffic




Fig. 6. P2P traffic traversing each link a) on each simulation b) mean values

The results included in this section illustrate a tracker configuration protecting specific links of the network from excessive P2P traffic, using a similar simulation scenario. In this specific case, the tracker was informed (e.g. by the network administrator) that it should protect the following links: $R7 \leftrightarrow R9$, $R8 \leftrightarrow R9$ and $R9 \leftrightarrow R10$ (see Figure 2). For that purpose, the tracker will try to reduce the betweenness centrality values associated with such links in order to decrease the corresponding P2P traffic. After this optimization process the

⁶ $p_{i \leftarrow j}$ was set to 0.4 for the nearest area and 0.15 for the other areas.

tracker will verify which are the most appropriate peering adjacencies to follow, and will apply such knowledge when returning peer samples to the requesting peers. In this specific case the tracker will find that the best way to protect the mentioned links is to define two independent peering groups⁷, one with peers from areas 1 and 6, and another one with peers from areas 2, 3, 4 and 5. Figures 7 a) and 8 a) show the estimated P2P link impact values evaluated by the tracker, after the optimization process, for methods IM_I and IM_{II} , respectively. The real traffic measures obtained for this scenario are presented in Figures 6 a) and b). As plotted, it is clearly visible that under this configuration links $R7 \leftrightarrow R9$, $R8 \leftrightarrow R9$ and $R9 \leftrightarrow R10$ are effectively protected from the P2P swarm behavior, only presenting almost imperceptible values of P2P traffic⁸. As before, the P2P link impact measures show an acceptable match with the relative values of the traffic gathered in simulation, with the IM_{II} method providing more accurate estimations, as depicted by Figures 7 b) and 8 b).




Fig. 7. a) P2P link impact measure IM_I b) IM_I vs observed traffic values




Fig. 8. a) P2P link impact measure IM_{II} b) IM_{II} vs observed traffic values

⁷ This solution will completely avoid traffic from the P2P swarm to traverse the defined links, i.e. impact measures equal to zero. However, other not so severe solutions could also be defined by the tracker.

⁸ The residual values are due to the implemented algorithm at the tracker, with an initial phase where no constraints are applied to the peering adjacencies.

When protecting specific links of the network domain, the tracker changes its default behavior selecting now which peers samples should be returned to specific clients. It would also be interesting to analyze the consequence of such behavior when compared with the results of Section 3.1. In this perspective, Figure 9 a) presents the peers download times obtained in the scenario of Section 3.1. The downloading times differences obtained under this new tracker configuration are plotted in Figure 9 b). As observed, some peers from areas 1 and 6 obtained slightly higher download times (roughly an 8% increase), while peers from the other areas experience download times which may increase or decrease in the same order of magnitude. Overall, such values do not significantly affect the overall service quality meaning that the objective of protecting specific network links was accomplished, in this case, with lower costs from the end-users P2P service quality perception.




Fig. 9. a) Peers download time in Scenario 1 b) variations observed in Scenario 2

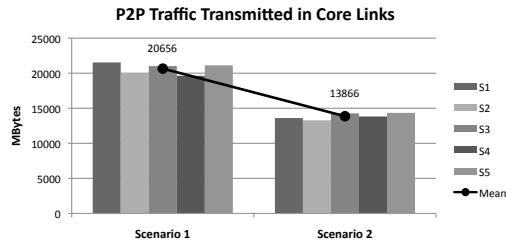


Fig. 10. Overall P2P traffic transmitted in the network domain links

Figure 10 also compares both scenarios as regards to the overall P2P traffic transmitted in the network. As observed, when the tracker was configured to protect some specific links (scenario 2) there is a clear decrease in the overall P2P traffic traversing the network domain (around a 32% decrease), representing a significant advantage from the operator perspective. This is justified by a side-effect resulting from the peering adjacencies induced by the tracker in Scenario

2, which forces nearing peers to participate in the data exchanges, thus avoiding unnecessary connections among distant peers. This example proves that intelligent tracker decisions can effectively improve network resource usage without significantly degrading end-users service quality.

3.3 Seed Placement Influence

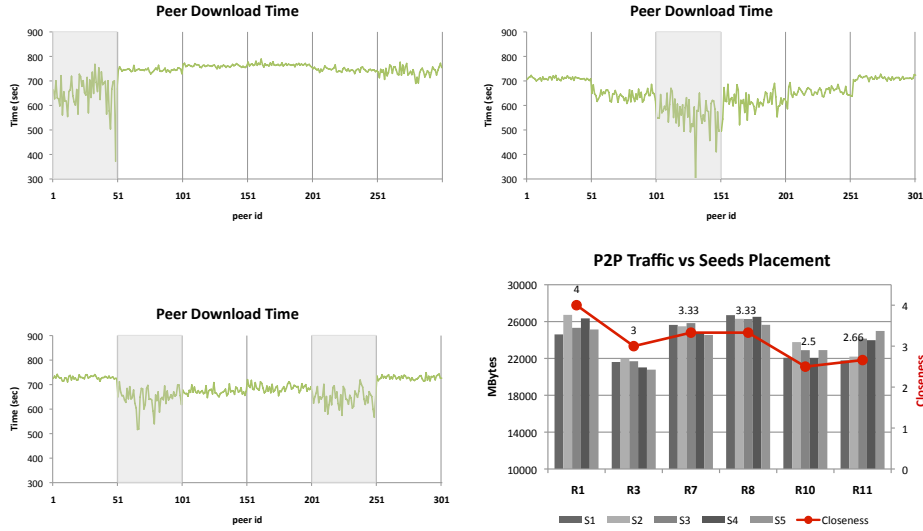


Fig. 11. Peers download times with seeds placed at a) R1 b) R3 c) R10; d) P2P traffic traversing the domain for distinct seeds placements

The results included in this section assume that the network level is able to participate in the P2P swarm configuration, namely being responsible for the placement of P2P seeds in the network. As explained, within the devised method, the tracker resorts to $CM(n)$ measures to devise appropriate placements for a given seed(s) and provide such information to the network level administrators responsible to place them in the pointed positions. In the next examples the tracker was required to devise the placement of two high upload capacity seeds in a given topology position, according to an administrative pre-defined criteria. The results presented by Figures 11 a) and 11 b) show the peers download times when the seeds are positioned in order to benefit end-users of areas 1 and 3, respectively. The results of Figure 11 c) were obtained when the closeness centrality measure was used by the tracker to benefit peers from areas 2 and 5. As observed, for each one of the cases, peers in the priority areas (denoted by gray filled areas in each one of the figures) achieve better service quality, i.e. a qualitative differentiation with peers within higher priority areas having

lower downloading times. In each one the three illustrative cases presented before the tracker indication was to place the seeds in network positions R1, R3 and R10, respectively (see Figure 2). In another distinct perspective, the tracker can also resort to the computation of the $CM(n)$ values not to induce service quality differentiation, but to provide feedback about the more appropriate seed(s) placement to avoid unnecessary P2P traffic in the domain. For that purpose, assuming a scenario with the tracker behaving in the classical mode, i.e. returning random samples to all peers, network positions having lower global closeness centrality values, $CM(n)$, are expected to be better candidates for seed positions. Figure 11 d) illustrates this by showing the centrality values of distinct candidate seed positions and the overall P2P traffic traversing the domain when seeds are placed in such positions. As observed, network locations having lower $CM(n)$ values show a tendency to originate lower amounts of P2P traffic.

4 Conclusions

This work explored the concept of a collaborative framework involving P2P applications and network level entities to underpin the development of advanced TE mechanisms. Taken the example of BitTorrent applications, several illustrative tracker configurations were explored being able to provide useful auxiliary information to network administrators, and better accommodate P2P traffic within the network infrastructure, e.g. by protecting specific links from excessive P2P traffic. Resorting to simulation, both the framework and the devised mechanisms were tested successfully. Even considering the inherent difficulties of controlling application level P2P dynamics and obtain precise impact estimations, it has been demonstrated that there is a wide range of possible fruitful collaboration efforts that could be made between the P2P and network levels. As future work, there are still many other TE related mechanisms that could be developed using the proposed framework. Moreover, some of the preliminary TE mechanisms proposed here can be further enhanced by also considering other network level specificities. In such cases, the tracker may gather additional information from the network level (e.g. congestion levels, packet loss, etc.) to further improve the modeling capabilities and the effectiveness of the devised TE mechanisms.

Acknowledgments

This work is partially funded by FEDER Funds through the Programa Operacional Fatores de Competitividade COMPETE and by National Funds through the FCT - Fundação para a Ciência e a Tecnologia (Portuguese Foundation for Science and Technology) within project FCOMP-01-0124-FEDER-022674.

References

1. Lua, K., Crowcroft, J., Pias, M., Sharma, R., Lim, S.: A survey and Comparison of Peer-to-peer Overlay Network Schemes. IEEE Communications Surveys & Tutorials, pp. 72-93 (2005)

2. Choen, B.: Incentives Build Robustness in BitTorrent. In: Proc. 1st Workshop on Economics of Peer-to-Peer Systems, Berkeley (June 2003)
3. Bharambe, A., Herley, C., Padmanabhan, V.: Analyzing and Improving a BitTorrent Networks Performance Mechanisms. In: Proc. IEEE INFOCOM (2006)
4. Cortez, P., Rio, M., Rocha, M., Sousa, P.: Multi-scale Internet traffic forecasting using neural networks and time series methods. *Expert Systems: The Journal of Knowledge Engineering* 29(2), 143-155, Wiley-Blackwell (2012)
5. Opsahl, T., Agneessens, F., Skvoretz, J.: Node Centrality in Weighted Networks: Generalizing degree and shortest paths. *Social Networks* 32(3), pp. 245-251 (2010)
6. Narayanan, S.: The Betweenness Centrality of Biological Networks. MSc Thesis, Faculty of the Virginia Polytechnic Inst. and State University (2005)
7. Karagiannis, T., et al.: Is P2P Dying or Just Hiding? In: Proc. Globecom, Dallas, TX, USA, November (2004)
8. Schulze, H., Mochalski, K.: 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)
9. Keralapura, R., Taft, N., Chuah, C., Iannaccone, G.: Can ISPs Take the Heat from Overlay Networks? In: Proc. HotNets-III, San Diego, CA (November 2004)
10. Qiu, L., Yang, Y. R., Zhang, Y., Shenker, S.: Selfish Routing in Internet-like Environments. In: Proc. of SIGCOMM, Karlsruhe, Germany (August 2003)
11. Sousa, P., Cortez, P., Rio, M., Rocha, M.: Traffic Engineering Approaches using Multicriteria Optimization Techniques. In: Proc. WWIC 2011 - 9th Int. Conf. on Wired/Wireless Internet Communications. LNCS, vol. 6649, pp. 104-115. Springer, Heidelberg (2011)
12. Xie, H., Krishnamurthy, A., Silberschatz, A., Yang, Y. R.: P4P: Explicit Communications for Cooperative Control between P2P and Network Providers. http://www.dcia.info/documents/P4P_Overview.pdf
13. Xie, H., et al.: P4P: Provider Portal for Applications. In: Proc. SIGCOMM 2008 Conference, August 17-22, Seattle, Washington, USA (2008)
14. Sousa, P., Rocha, M., Rio, M., Cortez, P.C.: Efficient OSPF Weight Allocation for Intra-domain QoS Optimization. In: Parr, G., Malone, D., O Foghlu, M. (eds.) IPOM 2006. LNCS, vol. 4268, pp. 3748. Springer, Heidelberg (2006)
15. Shen, G., Wang, Y., Xiong, Y., Zhao, B., Zhang, Z.: HPTP: Relieving the Tension between ISPs and P2P. In: Proc. of IPTPS, Bellevue, WA (February 2007).
16. Wierzbicki, A., Leibowitz, N., Ripeanu, M., Wozniak, R.: Cache Replacement Policies Revisited: The case of P2P traffic. In: Proc. of GP2P (2004)
17. Spognardi, A., Lucarelli, A., DiPietro, R.: A Methodology for P2P File-Sharing Traffic Detection. In: Proc. Second International Workshop on Hot Topics in Peer-to-Peer Systems 2005 (HOT-P2P 2005), pp. 52-61 (July 2005)
18. Karagiannis, T., Rodriguez, P., Papagiannaki, K.: Should Internet Service Providers fear Peer-assisted Content Distribution? In: Proc. Proceedings of the Internet Measurement Conference, Berkeley, CA (October 2005)
19. Madhyastha, H. et al. iPlane: An Information Plane for Distributed Services. In: Proc. of OSDI Conference, Seattle, WA, (2006)
20. Eger, K. et al.: Efficient Simulation of Large-Scale P2P Networks: Packet-level vs. Flow-level Simulations. In: Proc. 2nd Workshop on the Use of P2P, GRID and Agents for the Development of Content Networks (2007)
21. ns-2 The Network Simulator, <http://www.isi.edu/nsnam/ns/>