Estimating and Controlling the Traffic Impact of a Collaborative P2P System

Pedro Sousa

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

Abstract. Nowadays, P2P applications are commonly used in the Internet being an important paradigm for the development of distinct services. However, the dissemination of P2P applications also entails some important challenges that should be carefully addressed. In particular, some of the important coexistence problems existing between P2P applications and Internet Service Providers (ISPs) are mainly motivated by the inherent P2P dynamics which cause traffic to scatter across the network links in an unforeseeable way.
In this context, this work proposes a collaborative framework of a BitTorrent like system. Using the proposed framework and based on the exchange of valuable information between the application and network levels, some novel techniques are proposed allowing to estimate and control the traffic impact that the P2P system will have on the links of the underlying network infrastructure. Both the framework and the presented techniques were tested resorting to simulation. The results clearly corroborate the viability and effectiveness of the formulated methods.

Keywords: P2P; Traffic Engineering; BitTorrent; Collaborative systems

1 Introduction

P2P overlays [1] can be considered as self-organized systems operating on top of a given network infrastructure. Such systems usually adopt specific protocols and peering strategies which may significantly change the traffic profiles observed in the network, thus also posing new problems to the Internet Service Providers (ISPs). BitTorrent [2] is just an example of a widely used P2P protocol over which many applications rely to exchange considerable large resources among a significant number of users, being also responsible by a considerable amount of the Internet traffic [3,4]. However, several coexistence problems between ISPs and P2P applications emerged in the last years, being this motivated by several factors. In fact, P2P dynamics cause traffic to scatter across the network links in an unforeseeable way. As consequence, P2P approaches are not always consistent with ISP economic models, as specific links from the underlying network might be under excessive and unpredictable traffic loads and some unnecessary interdomain traffic could also be generated [5,6]. Another important issue is that ISPs

many times use several Traffic Engineering (TE) techniques for tasks such as capacity planning, resilience improvements, routing optimization, among others. One of the critical inputs for TE tasks is the estimation of the traffic matrix of the network infrastructures. In this perspective, P2P overlay networks make complex the demand matrix estimation, and such estimation errors will also affect all the other TE related tasks [7,8].

Given all the above mentioned, this work presents a contribution to attain a BitTorrent-like P2P system architecture with the ability to collaborate and help network level entities to better deal with the P2P traffic generated by the application level. For that purpose, the devised framework includes methods allowing to estimate the traffic impact that the P2P system will have on the underlying network links. Such qualitative impact estimation values are able to provide a preliminary view about the traffic patterns that will traverse the network infrastructure, thus being an important asset from the ISP point of view. Additionally, within the proposed framework, the P2P tracker ruling the P2P swarm behavior is able to be dynamically configured in order to make an effort to protect specific network links from the generated P2P traffic. The presented framework corroborates the advantages of pursuing collaborative efforts in this area, as also highlighted by other works (e.g. [9,10]). In this case, the proposed solution raises the P2P application level with mechanisms allowing to estimate and control the P2P traffic impact, making such enhanced methods available to network level entities through specific configuration interfaces. As compensation, ISPs are expected to provide such collaborative P2P systems with a privileged traffic treatment, in contrast with more aggressive techniques used to punish other nonconforming P2P approaches.

The paper is organized as follows: Section 2 presents the proposed framework rationale, also explaining the devised mechanisms for estimating the network impact of P2P swarms and for protecting specific links from P2P traffic. Section 3 describes the implemented simulation platform and presents illustrative results corroborating the effectiveness of the devised mechanisms. Finally, Section 4 concludes the presented work.

2 P2P System Architecture and Devised Methods

This work proposal focus on a BitTorrent-like framework having some enhanced features. In the devised system the network level is able to obtain estimations about the traffic impact that the P2P system will have on the network infrastructure and, if required, control how such traffic traverses the network domain. For this end, the framework depicted in Figure 1 assumes a collaborative perspective between the application level (e.g. Service Providers) and network level entities (e.g. ISPs) with the exchange of valuable information. The framework might be used in distinct contexts. For instance, it could allow ISPs to offer Internet users a friendly P2P system behaving in a collaborative way with the network level. In a different perspective, it can be also used by Service Providers to develop specific services involving their clients, such as the upload of large

resources (e.g. generic data files, media, software packages, etc.) to all of their customers in a P2P fashion, also benefiting from specific agreements made with the network provider. This exchange might occur in previously scheduled time periods, having end users to notify the Service Provider that they intend to integrate the corresponding P2P swarm on such allocated time slots.

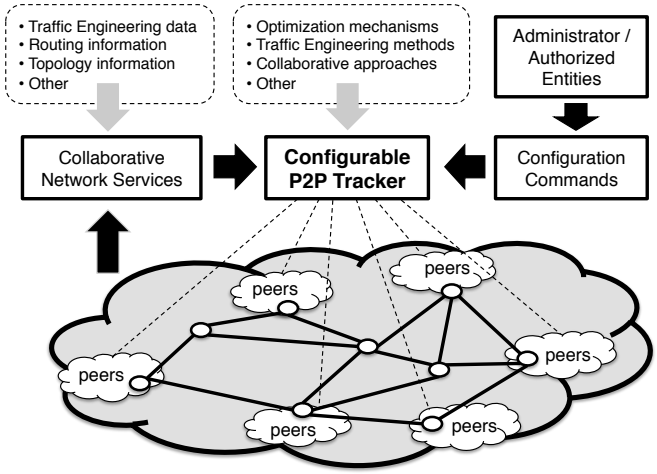

Fig. 1. General view of the envisioned P2P framework and participating entities.

The operation of a BitTorrent like system usually implies that interested peers establish a contact with the P2P tracker controlling a specific P2P swarm. As consequence, the P2P tracker sends a random sample of peers already present in the swarm to the contacting peer. With that, the newly arrived peer attempt to establish network connections with other peers in order to exchange the pieces of the file. From that point on, the BitTorrent defines several rules that affect the data transfer and the choke/unchoke processes among the swarm peers [1,2,11]. In addition, periodically, peers are allowed to contact the tracker to obtain a renewed sample of peers. The framework depicted in Figure 1 assumes that the P2P tracker fully controls the peering information provided to clients, which means that clients are not allowed to exchange peering information between them. Most of the classical BitTorrent based P2P systems can also behave in this way by the use of specific options conveniently defined in the *.torrent* file.

As also visible in Figure 1, the framework adopts the use of configurable P2P trackers making possible that distinct configurations could be made on the tracker, also allowing the tracker interaction with other external entities. The tracker internal architecture is not the focus of this work, but could follow similar directives as the presented in [10,12]. Moreover, if required, the tracker may also resort to several optimization mechanisms, including mechanisms from the field of computational intelligence (e.g. [16], [15]). The configurable P2P tracker

integrated in Figure 1 is then able to receive valuable network level information from collaborative network services, such as topological, routing and other traffic engineering related inputs. The tracker can also be programmed with internal methods that might be activated through specific configuration commands. For that purpose the tracker receives configuration commands from administrators, or other authorized entities, instructing it to adopt a specific behavior for a particular P2P swarm. As a reward for the use of the devised P2P collaborative approach, network level providers are expected to give a better traffic treatment to this P2P system, in counterpoint to other P2P approaches that will suffer from the restrictions usually imposed by ISPs (e.g. bandwidth throttling).

The following Section 2.1 describes a method allowing to estimate the traffic impact of a P2P swarm in the network links of the underlying infrastructure. Following that, Section 2.2 explains how is possible to protect specific network links from excessive P2P traffic.

2.1 P2P Link Impact Values

This section describes a method allowing to attain an estimation about the impact that traffic generated by a given P2P swarm will have of the network links, when involving a considerable number of peers. Thus, for a given swarm composition and assuming the tracker behaving in the classical mode, the objective is that such qualitative link impact information could be provided to the ISP.

Lets assume a classical mathematical representation of a network, with the graph $G = (N, L)$ expressing a network domain (e.g. an ISP network), were N is a set of the network nodes/routers and L a set of the interconnecting network links, for which routing link weights are also considered for shortest path computation. Part of the network nodes/routers might also be viewed as Points of Presence (PoP) to end-users areas having peers interested to participate in a given P2P swarm. For convenience, the location of such end-users areas is denoted by the corresponding ISP network router, a, with $a \in A$ and $A \subseteq N$.

Within the scope of the proposed mechanism, several graph measures (e.g. [13,14]) could constitute valuable inputs, in particular the concept of betweenness centrality in a graph, here adapted and extended to provide estimations of the P2P traffic link impact. The devised impact estimation metric combines distinct factors that could present a preliminary snapshot of the traffic patterns exchanged within a large P2P swarm. For a specific ISP link, l, and a pair of end-users areas, $i, j \in A$, we consider the ratio between the number of shortest paths from i to j, $sp_{i,j}$, and the number of such paths that effectively pass through link l, $sp_{i,j}(l)$. By this way, each link l is assigned with a partial impact value of $\frac{sp_{i,j}(l)}{sp_{i,j}}$ for the case of peering adjacencies between areas i, j. When accounting all possible area adjacencies this metric will present higher values for links which integrate a higher number of shortest paths among the areas, thus having such links higher probabilities of being traversed by the P2P swarm traffic. A second weighting factor, $w_{i,j}$, is also considered for case of P2P swarms where end-user areas have an unbalanced distribution of peers. This factor considers the ratio

between the number of peers involved in the peering adjacencies of areas i, j over the total number of peers involved in all possible adjacencies, favoring the importance of shortest paths connecting areas involving higher number of peers.

The above mentioned rationale can be further enhanced taking into account some characteristics of the TCP protocol that is used in the data transfers among BitTorrent peers. In fact, in such protocolar approach, peers often have a higher probability to establish peering connections with nearest peers in the network, taking advantage of lower network round-trip times (RTT). Thus, for shortest paths between areas i and j a preference value[1] ($p_{i \leftarrow j} \in [0, 1]$ with $\sum_{j \in A, j \neq i} p_{i \leftarrow j} = 1$) is assigned to such adjacencies, implicitly denoting how close are areas j and i. Considering all the above mentioned reasoning, and for the case of a tracker returning random samples to contacting peers, Equation 1 presents the devised normalized P2P link impact value ($P2P_{LIV}$) value for link l, within the interval $[0, 1]$. The tracker may announce these estimations to network services or administrators which in turn are able to instruct the tracker to protect specific links from the infrastructure.

$$P2P_{LIV}(l) = \sum_{\substack{i,\, j \,\in\, A \\ i \neq j}} [(|A| - 1) \cdot p_{i \leftarrow j}] \cdot \frac{sp_{i,j}(l)}{sp_{i,j}} \cdot w_{i,j} \quad l \in L \tag{1}$$

The metric presented by Equation 1 has the major objective of gathering a preliminary snapshot of which links are expected to be traversed by higher amounts of P2P traffic. The objective is that the comparison between the $P2P_{LIV}$ values of two links can be used to foresee which one will be traversed by higher amounts of P2P traffic, i.e. that the order relations between $P2P_{LIV}$ values could also somehow express the order relation between the P2P traffic that will flow over such links.

In order to validate the correctness of such impact estimations, the function $f(l, z)$ (presented in Equation 2) is defined for two distinct links $l, z \in L$. As observed in Equation 2, the function $f(l, z)$ might return two alternative values $\{0, 1\}$ according with the estimated $P2P_{LIV}$ metrics and the traffic that effectively traverses such links (function $T(l)$) after running a real/simulated experiment of the framework. If the $P2P_{LIV}$ order relations also express the $T(l)$ order relations the value returned by $f(l, z)$ is 1, otherwise 0. For the particular case of links having exactly equal $P2P_{LIV}$ values a small deviation (controlled by the γ variable) is accepted when comparing the observed traffic on each link.

[1] This value is then multiplied by the total number of distinct external areas adjacencies that could be made by peers in a given area, i.e. $|A| - 1$, for normalization purposes.

$$f(l,z) = \begin{cases} 1 \text{ if } (P2P_{LIV}(l) > P2P_{LIV}(z)) \& (T(l) > T(z)) \\ 1 \text{ if } \big(P2P_{LIV}(l) < P2P_{LIV}(z)) \& (T(l) < T(z)) \\ 1 \text{ if } \big(P2P_{LIV}(l) = P2P_{LIV}(z)) \& (T(l) \in [T(z) \cdot (1-\gamma), T(z) \cdot (1+\gamma)]) \\ 0 \text{ otherwise} \end{cases} \tag{2}$$

Based on the $f(l,z)$ function, Equation 3 defines now the function $\psi(l)$ expressing the order conformity of the $P2P_{LIV}$ impact value of link l. Thus, $\psi(l)$ represents the average $f(l,z)$ values obtained when directly comparing link l with all the other links of a given network topology. Therefore, $\psi(l)$ values will vary within the interval $[0,1]$, with values close to 1 expressing that most of the order relations among $P2P_{LIV}$ values also express the order relations between the P2P traffic that effectively traverses the links. Thus, function $\psi(l)$ will be used to assess the quality of the $P2P_{LIV}$ results obtained in the experimental part of this work.

$$\psi(l) = \frac{\sum_{z \in L \setminus \{l\}} f(l,z)}{|L| - 1} \quad l \in L \tag{3}$$

2.2 Protecting Links from P2P Traffic

As previously explained, links having higher $P2P_{LIV}$ values are expected to be traversed by larger amounts of traffic. In this perspective, we now explore a possible method allowing the tracker to control the P2P swarm traffic distribution in the network domain, namely by protecting specific links of the network from P2P traffic. The devised method allows the tracker to reduce the link impact values of specific network links by conveniently manipulating the peer samples returned to the contacting peers.

Algorithm 1 presents the pseudo-code of the proposed method. As inputs it receives the swarm identification, an ordered set of the protected links and collaborative information provided by the network level. The methods starts by considering a set with all the area pairs combinations of the network (X_s, line 2), where each pair (a_i, a_j) means that when contacted by a peer from area a_i the tracker is able to include in the random sample peers from the area a_j. After that, and for each protected link $link_l$, the algorithm uses the topology and routing information provided by collaborative network entities to construct a subset Y containing the (a_i, a_j) pairs for which the shortest paths connecting such areas traverse $link_l$ (line 4). In the next step the algorithm verifies if is possible to remove a specific (a_i, a_j) entry from X_s in order to reduce the impact of the P2P swarm traffic on such link. The function *swarm_totally_connected*() (in line 6) verifies if the swarm is still totally connected when considering that the tracker will not include peers from area a_j in peer samples sent to peers

Algorithm 1 *protecting_links_from_P2P_Traffic (s, K, data)*

1: {**Comment:** ***s***- *a swarm identification;* ***K***- *a decreasingly ordered set with all* $link_l \in L$ *protected links (ordered by priority);* ***data***- *auxiliary information provided by collaborative services (topology, routing, etc.)*}
2: $X_s \leftarrow$ decreasingly ordered set with all (a_i, a_j) area pairs having peers from swarm s, $a_i, a_j \in A$ {**Comment:** X_s is a $w_{i,j} * p_{i \leftarrow j}$ ordered set}
3: **for all** $link_l \in K$ **do**
4: $Y \leftarrow$ decreasingly ordered subset of X_s with (a_i, a_j) pairs which shortest paths include $link_l$ {**Comment:** Y is a $w_{i,j} * p_{i \leftarrow j}$ ordered set}
5: **for all** $(a_i, a_j) \in Y$ **do**
6: **if** $swarm\_totally\_connected(s, X_s \setminus \{(a_i, a_j)\}) = TRUE$ **then**
7: $X_s \leftarrow X_s \setminus \{(a_i, a_j)\}$
8: **end if**
9: **end for**
10: **end for**
11: $update\_tracker(s, X_s)$

from area a_i. The swarm is assumed to be totally connected if all peers have the opportunity to contact one of the swarm seeds, or contact other peers that directly or indirectly have access to the pieces sent by a seed. Otherwise the swarm is considered to be partitioned and some peers will never receive all the pieces of the shared file. In the case that the swarm would not become partitioned the (a_i, a_j) pair is effectively removed from X_s (line 7).

Algorithm 2 *get_peer_sample(peer p, swarm s)*

1: *peer_area* $\leftarrow$ get_peer_location(p)
2: **if** swarm_in_initial_state(s) **then**
3: *peer_sample* $\leftarrow$ random_sample(s)
4: **else**
5: *peer_sample* $\leftarrow$ random_sample_from_X_s (s, *peer_area*, X_s) {**Comment:** X_s *was previously computed by the tracker using Algorithm 1*}
6: **end if**
7: update_swarm_info(p, s)
8: return(*peer_sample*)

As result, at the end of Algorithm 1, the set X_s will contain all area pairs that the tracker should consider to build the random peers samples. The considered peering adjacencies are sufficient to build a totally connected swarm, having also the minimum possible traffic impact on the considered protected links. After the computation of the final X_s set, the tracker will adopt Algorithm 2 to return a peers sample whenever contacted by any peer. As illustrated in Algorithm 2 in the initial state of the swarm, i.e. during a short initial period over which only

few peers have contacted the tracker, the tracker behaves in the classical mode, i.e. a random peer sample is build considering all the available peers[2]. After that period, the tracker takes into account the area of the contacting peers and builds random peer samples constrained by the allowed peering adjacencies expressed in the X_s set returned by Algorithm 1 (line 5 of Algorithm 2).

3 Simulation Testbed and Results

Figure 2 presents the modules that were implemented in a simulation platform (ns-2 [17]) and the selected network topology to present illustrative results. A patch implementing the dynamics of the BitTorrent protocol [18] was used as the baseline over which other components from the devised framework were added. Specific modules were built for the implementation of the configurable tracker as well as the collaborative network services module providing network level information to the tracker. A specific interface was devised, allowing to interact with the tracker and provide configuration commands to activate some implemented methods. The link impact estimation and the link protection methods (explained in Sections 2.1 and 2.2) were programmed in the tracker internal logic, being able to be activated whenever triggered by specific configuration commands.

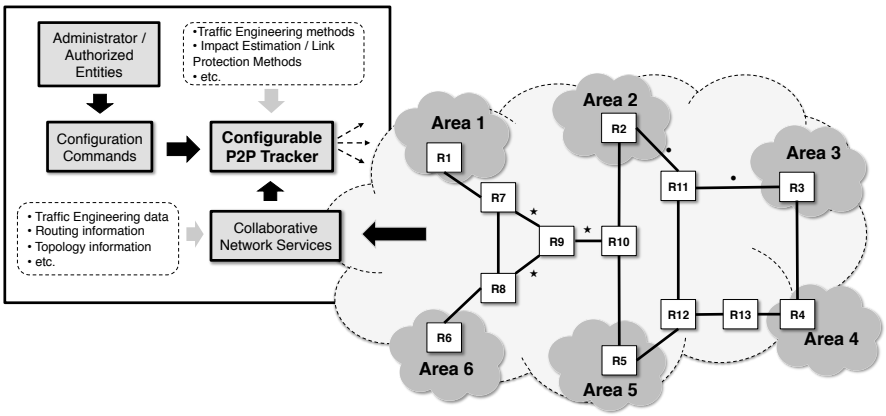

Fig. 2. Implemented framework and illustrative topology used in the experiments.

To present illustrative simulation results the network topology depicted in Figure 2 was used. It illustrates a network topology consisting of several end-users areas interconnected by several links and core routers (some of which can also be viewed as possible Points of Presence (PoPs) of the ISP). The depicted

[2] This limited initial period will generate an almost negligible volume of traffic on the protected links, but contributes for assuring an improved initial distribution of some file pieces among all the areas.

topology intentionally includes several topological characteristics making more challenging the test of the devised mechanisms, such as: areas connected by paths with distinct distances, equal cost paths between some areas, critical links which failure will originate a partition on the network, single and multi-homed areas, etc. For routing purposes it is assumed that the ISP shortest paths are the ones having a small number of hops between a given source/destination pair (i.e. routing link weights of 1 for all links). The scenario assumes that peers participating in the P2P swarm are distributed along six areas, being each area composed by a second level of routers/links. In the developed simulation platform, several parameters can be configured, including the number of peers and seeds per area, the file size, the chunk size, among others.

The examples presented in the following sections assume a total number of 300 peers in the swarm, exchanging a file of 50 MB and operating with a chunk size of 256 KB. The parameters P_D and S_L might be used to control the distribution of the peers and seeds in the distinct network areas, respectively. By default, the peer sample returned by the tracker includes 25 peer contacts. At each area the peers have an upload capacity of 1 Mbps and a download capacity of 8 Mbps, thus simulating common residential scenarios where users have higher download capacities. To force some heterogeneity within each area, propagation delays of the users access links randomly vary within the interval $[1, 50]$ ms. Due to the collaborative nature of the devised P2P system, the scenario also assumes that the ISP allows on each link a share of 50 Mbps exclusive for P2P traffic generated by the proposed P2P system, and the propagation delays of such links are at least two times higher than the end users access links. In the following sections, for each one of the described experiments, five simulations were made and the corresponding mean values were taken for analysis.

3.1 P2P Traffic Impact - Link Impact Values ($P2P_{LIV}$)

Based on the scenario depicted in Figure 2 several results are now presented regarding the tracker method to estimate the P2P impact on the network links. In the provided examples several scenarios were considered for distinct combinations of peers distribution in the network, P_D, and seed locations, S_L, and the obtained results are shown in Figure 3. The scenarios vary from an uniform distribution of peers in the network areas (first row of Figure 3 with all areas having 50 peers, i.e. P_D=$(50, 50, 50, 50, 50, 50)$) to other scenarios where a higher density of peers is considered to exist in specific parts of the network. The results of such additional peer distributions are presented in the other rows of Figure 3, assuming that the left, right, upper and bottom sides of the topology of Figure 2 have a higher density of peers, respectively. In addition, for each of the mentioned P_D distributions, three distinct seed positioning scenarios are considered: *i)* all areas having one seed; *ii)* a single seed positioned in area 1 and *iii)* a single seed positioned in area 4 (first, second and third columns of Figure 3).

Each graph of Figure 3 presents the results obtained on each particular scenarios (five independent simulations runs were made for each one and the plotted results are averaged values). For comparative analysis, on each graph, the

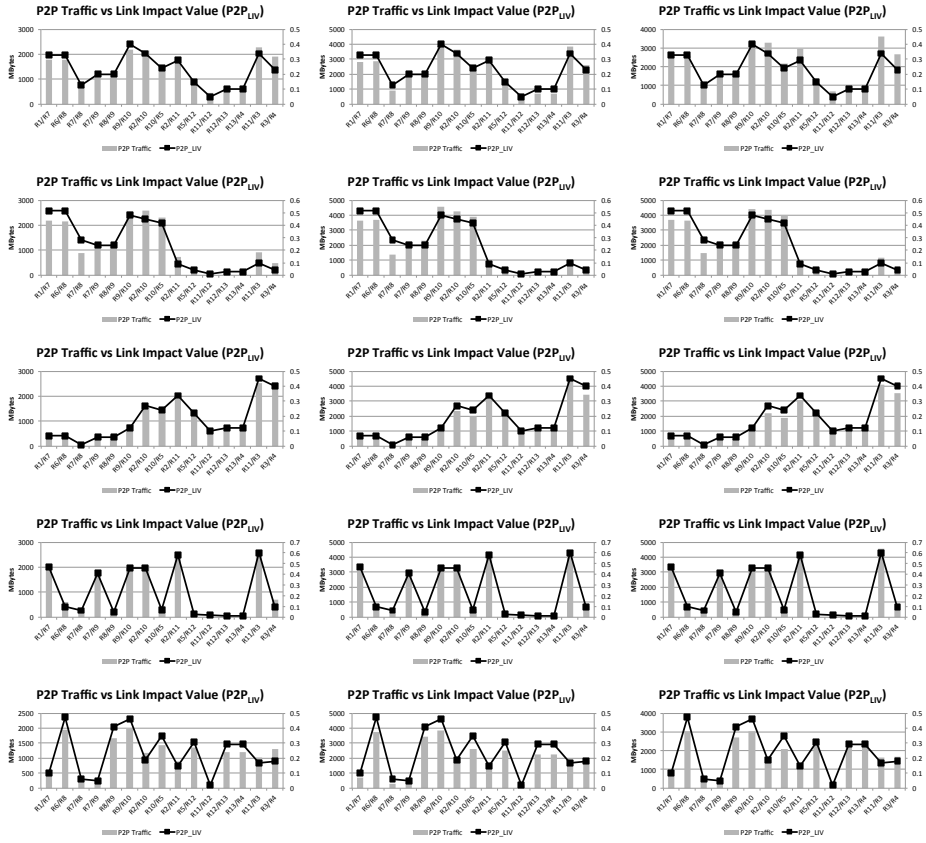

Fig. 3. P2P traffic on links vs $P2P_{LIV}$ values for distinct P_D and S_L values. Row1: P_D=(50,50,50,50,50,50); Row2: P_D=(70,70,10,10,70,70); Row3: P_D=(10,70,70,70,70,10); Row4: P_D=(90,90,90,10,10,10); Row5: P_D=(10,10,10,90,90,90); Column1: S_L=all ; Column2: S_L=A_1; Column3: S_L=A_4.

cumulative P2P traffic which traversed each link during the swarm lifetime is represented by gray filled columns (in MBytes), being the previoulsy computed $P2P_{LIV}$ impact metrics[3] for each link (Equation 1) represented by a black line-plot representation (normalized values within $[0, 1]$). A detailed analysis of Figure 3 allows to verify that in all of the considered scenarios both the $P2P_{LIV}$ link values and the overall P2P traffic on each link follow a similar trend. This constitutes a preliminary indication that $P2P_{LIV}$ metric could in fact denote the relations between the P2P traffic traversing each link during the swarm lifetime.

In order to verify the correctness of the $P2P_{LIV}$ metrics, the link impact order conformity metric (function $\psi(l)$ in Equation 3) was evaluated for each one

[3] As in real scenarios the tune of $p_{i \leftarrow j}$ values is difficult, in the experiments only nearest areas are differentiated ($p_{i \leftarrow j}$=0.4), the remaining areas have values of 0.15.

of the topology links within each one of the simulated scenarios. The obtained $\psi(l)$ values are summarized in Table 1[4]. As observed the link impact metrics obtained high order conformity values. In fact, in most of the presented scenarios and independently of the peers distribution and seed locations the $\psi(l)$ averaged values fall within the interval $[0.89, 97]$. This means that, for an expressive majority of the cases, the $P2P_{LIV}$ link impact values computed by the tracker also denote the foreseeable order relations between the P2P traffic traversing each link. In that way, $P2P_{LIV}$ values can effectively be used to have a preliminary view about which links will suffer higher impact from the P2P swarm traffic, thus being this information a valuable asset for ISPs and network administrators.

Table 1. Link Impact Value Order Conformity $\psi(l)$ on the Simulated Scenarios (for each simulated instance of Figure 3)

Scenar.		**Link Impact Value Order Conformity $\psi(l)$**															
P_D	S_L	R_1 R_7	R_6 R_8	R_7 R_8	R_7 R_9	R_8 R_9	R_9 R_{10}	R_2 R_{11}	R_{10} R_5	R_2 R_{11}	R_5 R_{12}	R_{11} R_{12}	R_{12} R_{13}	R_{13} R_4	R_{11} R_3	R_3 R_4	**Avg** $\overline{\psi}(l)$
50,50,	*all*	0.86	0.86	1.00	1.00	1.00	0.93	0.93	0.93	0.79	1.00	1.00	1.00	1.00	0.86	0.71	**0.92**
50,50,	A_1	0.93	0.93	1.00	1.00	1.00	1.00	0.93	0.93	0.86	1.00	1.00	1.00	1.00	0.93	0.93	**0.96**
50,50	A_4	0.86	0.86	1.00	1.00	1.00	0.93	0.93	0.93	0.86	1.00	1.00	1.00	1.00	0.86	0.79	**0.93**
70,70,	*all*	0.79	0.79	0.79	0.93	0.93	0.79	0.79	0.86	1.00	0.93	1.00	1.00	1.00	0.93	0.93	**0.90**
10,10,	A_1	0.79	0.79	0.86	0.93	0.93	0.86	0.86	0.86	1.00	0.93	1.00	1.00	1.00	1.00	0.93	**0.91**
70,70	A_4	0.79	0.79	0.86	0.86	0.86	0.86	0.86	0.86	1.00	0.93	0.86	0.93	0.93	1.00	0.93	**0.89**
10,70,	*all*	0.93	0.93	1.00	0.93	0.93	0.86	1.00	1.00	1.00	1.00	1.00	0.93	0.93	0.93	0.93	**0.95**
70,70,	A_1	0.93	0.93	1.00	0.93	0.93	0.79	0.93	0.93	1.00	0.86	0.93	0.93	0.93	1.00	1.00	**0.93**
70,10	A_4	0.93	0.93	1.00	0.93	0.93	0.79	1.00	0.93	1.00	0.93	0.93	0.93	0.93	1.00	1.00	**0.94**
90,90,	*all*	0.86	0.93	0.93	1.00	0.93	0.86	0.86	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.93	**0.95**
90,10,	A_1	0.86	0.93	0.86	1.00	0.93	0.86	0.86	1.00	1.00	0.93	1.00	1.00	1.00	1.00	0.93	**0.94**
10,10	A_4	0.86	0.93	0.93	1.00	0.93	0.86	0.86	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.93	**0.95**
10,10,	*all*	1.00	0.93	0.93	0.93	1.00	0.93	0.93	1.00	1.00	1.00	1.00	0.93	0.93	1.00	0.79	**0.95**
10,90,	A_1	1.00	0.93	0.93	0.93	1.00	0.93	1.00	1.00	1.00	1.00	1.00	1.00	1.00	0.93	0.93	**0.97**
90,90	A_4	1.00	0.93	0.93	0.93	1.00	0.93	1.00	0.79	1.00	0.93	1.00	0.93	0.93	1.00	1.00	**0.95**

$\mathbf{P_D}$ - Peers distribution in the network (A_1,...,A_6), $\mathbf{S_L}$ - Seeds location in the network

3.2 Protecting Network Links from P2P Traffic

This section illustrates the tracker configuration mode explained in Section 2.2, namely in Algorithm 1, where some specific network link(s) are protected from the traffic generated by the P2P swarm.

[4] For $\psi(l)$ computation (Eq. 2) variable γ was assigned with a value of 0.025, i.e. only allowing a traffic deviation of 2.5% when comparing links with equal $P2P_{LIV}$ values.

In the first example the tracker was instructed to protect the links $R_7 \rightarrow R_9$, $R_8 \rightarrow R_9$ and $R_9 \rightarrow R_{10}$ from the network topology (links identified with a $\star$ mark in Figure 2) considering the scenario with a balanced distribution of peers in the network and one seed in all the network areas. The resulting traffic behavior is plotted in Figure 4, which compares the traffic observed in the network when the tracker behaves in the classical mode, Figure 4 a), and when configured with Algorithm 1 to protect the mentioned links, Figure 4 b). As observed in Figure 4 b) with the devised mechanism the cumulative P2P traffic traversing the selected links is almost imperceptible[5], comparatively with the scenario where the tracker assumes the classical behavior and a significant amount of traffic is observed in links $R_7 \rightarrow R_9$, $R_8 \rightarrow R_9$ and $R_9 \rightarrow R_{10}$ (plotted in Figure 4 a)), i.e. 2175, 1087 and 1087 MBytes, respectively. In this example, the protection of the links is obtained as Algorithm 1 computes a X_s set that only maintains area adjacencies pairs in two independent groups. In the first group, peers from areas A_1 and A_6 are not allowed to receive peers samples involving peers from other areas (i.e. A_2, A_3, A_4 and A_5), and in the second group peers from areas A_2, A_3, A_4 and A_5 are not able to receive peers samples integrating peers from areas A_1 and A_6. In this way all the P2P swarm traffic that would intersect the protected links is avoided by the tracker computed peering constraints. Thus, in the example of Figure 4 b) the tracker has computed the following allowed adjacencies:

$$X_s = \{(A_1, A_1), (A_1, A_6), (A_6, A_1), (A_6, A_6), (A_2, A_2), (A_2, A_3), (A_2, A_4), (A_2, A_5), (A_3, A_2), (A_3, A_3), (A_3, A_4), (A_3, A_5), (A_4, A_2), (A_4, A_3), (A_4, A_4), (A_4, A_5), (A_5, A_2), (A_5, A_3), (A_5, A_4), (A_5, A_5)\}$$

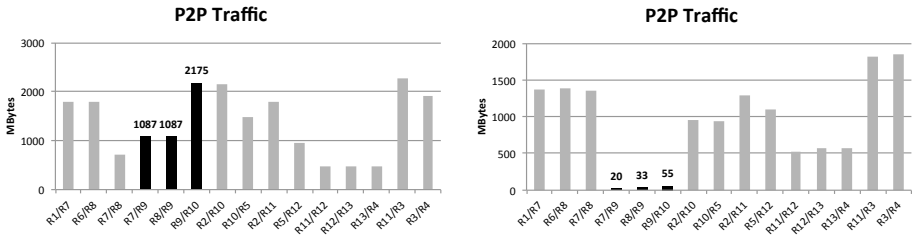

Fig. 4. Scenario with P_D=(50, 50, 50, 50, 50, 50), S_L=*all* a) classical tracker configuration; b) tacker configured to protect $R_7 \rightarrow R_9$,$R_8 \rightarrow R_9$,$R_9 \rightarrow R_{10}$ using Algorithm 1

The example presented in Figure 4 could be considered as having lower complexity due to the fact that distinct seeds were considered to exist on each network area. Thus, the behavior of Algorithm 1 could be considered has somehow foreseeable, not having to deal with possible swarm partitioning problems that could occur in more complex scenarios. In this perspective, a second example is now presented with a more challenging task. This case assumes the same peer

[5] The residual values observed are due to the first phase of Algorithm 2 were no peering adjacencies constraints are considered.

distribution as in Figure 4, but considering now that only a single seed in area A_1 exist, for the same set of links to be protected. As consequence, Algorithm 1 returns in this case a slightly distinct solution to the tracker, also integrating the (A_6, A_4) areas pair in the X_s set of the previous example. Otherwise, without such pair, a partition will occur in the swarm[6]. Figure 5 b) plots the results for this new scenario. As observed, this time the links $R_8 \rightarrow R_9$,$R_9 \rightarrow R_{10}$ have been traversed by some traffic from the P2P swarm, which is required to preserve the swarm totally connect (traffic exchanged between areas A_6 and A_4). Nevertheless, as Algorithm 1 tries to minimize traffic on protected links, there is still a significant traffic reduction even in such links, as observed when comparing Figures 5 a) and b). In fact, traffic on link $R_9 \rightarrow R_{10}$ is now five times lower than in the classical configuration, traffic on link $R_8 \rightarrow R_9$ is nearly two and a half times lower and traffic on link $R_7 \rightarrow R_9$ only presents residual values.

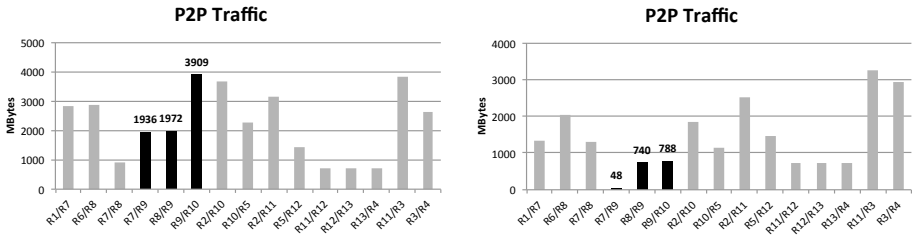

Fig. 5. Scenario with P_D=(50, 50, 50, 50, 50, 50), S_L=A_1 a) classical tracker configuration; b) tacker configured to protect $R_7 \rightarrow R_9$,$R_8 \rightarrow R_9$,$R_9 \rightarrow R_{10}$ using Algorithm 1

The last example presented assumes that the tracker was instructed to protect the links $R_2 \rightarrow R_{11}$ and $R_{11} \rightarrow R_3$ (identified with a • mark in Figure 2), for the same scenario as in Figure 5. The results presented in Figure 6 a) and b) corroborate again the effectiveness of the proposed link protection approach, as only residual traffic values are observed in the protected links. For this specific example the tracker has computed the following allowed adjacencies:

$X_s = \{(A_1, A_1), (A_1, A_2), (A_1, A_5), (A_1, A_6), (A_2, A_1), (A_2, A_2), (A_2, A_5), (A_2, A_6), (A_3, A_3), (A_3, A_4), (A_4, A_3), (A_4, A_4), (A_4, A_5), (A_5, A_1), (A_5, A_2), (A_5, A_4), (A_5, A_5), (A_5, A_6), (A_6, A_1), (A_6, A_2), (A_6, A_5), (A_6, A_6)\}$

A more depth analysis of the X_s computed by the tracker allows to verify that, in this example, peers from area A_3 are very constrained in peering opportunities, only being allowed to contact peer in the same area or in area A_4. However, peers in area A_4 are allowed to contact peers in area A_5 which, in turn, have directly or indirectly access to the pieces sent by the seed in area A_1. In

[6] Note that area A_6 is able to contact peers from Area A_1 where the seed is located. Thus, the (A_6, A_4) peering adjacency now added to X_s indirectly allows that areas A_2, A_3, A_4 and A_5 have also access to all pieces of the files exchanged in the swarm.

this perspective, once again the computed X_s solution ensures the integrity of the P2P swarm and the protection of the considered links.

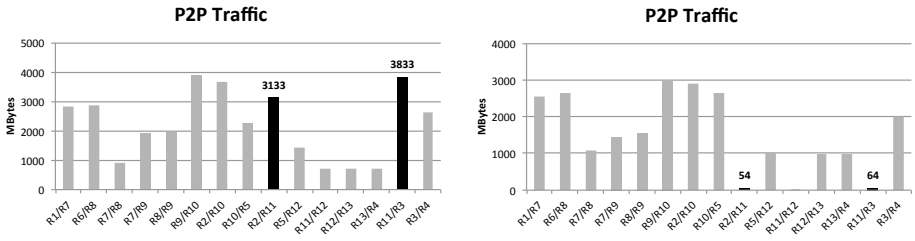

Fig. 6. Scenario with P_D=(50, 50, 50, 50, 50, 50), S_L=A_1 a) classical tracker configuration; b) tacker configured to protect $R_2 \rightarrow R_{11}$,$R_{11} \rightarrow R_3$ using Algorithm 1

4 Conclusions

This paper described a framework for a collaborative BitTorrent-like system involving network level (e.g. ISPs) and application level (e.g. Service providers) entities. In particular, this work focused on a system with the ability of providing link impact estimations about the traffic generated by P2P BitTorrent swarms. This allows to foresee how the network level links will be affected by the P2P traffic, thus being an important asset for ISP administrators. Complementary, a method was presented allowing to manipulate in an intelligent manner the peering information sent by the trackers. As consequence, the P2P tracker can be informed about which link(s) it should protect from the P2P swarm, generating for that purpose an optimized set of the allowed peering adjacencies, still ensuring the full connectivity of the swarm.

As a proof of concept, both the framework modules as well the devised methods were implemented in a simulation platform. The preliminary results obtained clearly corroborate that the mechanisms for P2P link impact estimations and for the protection of links from P2P traffic presented acceptable behavior. As future work, we intend to pursue the study on the effectiveness of the proposed mechanisms, analyzing additional complementary scenarios and configuration parameters. In a similar way, it is also intended to further enrich the proposed framework with other intelligent mechanisms that could benefit the integration of collaborative P2P applications in current networking environments.

Acknowledgments: This work has been supported by FCT - Fundação para a Ciência e Tecnologia within the Project Scope: PEst-OE/EEI/UI0319/2014.

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*, vol 7, Issue 2, pp. 72-93 (2005).
2. Choen, B.: Incentives build robustness in BitTorrent. *In Proceedings 1st Workshop on Economics of Peer-to-Peer Systems*, Berkeley (Jun. 2003).
3. Karagiannis, T., et al.: Is p2p dying or just hiding?. *In Proceedings of GLOBECOM*, Dallas, USA, (Nov. 2004).
4. 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. *Technical Report* (2007).
5. 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* (2008).
6. Seetharaman, S., Ammar, M.: Characterizing and mitigating inter-domain policy violations in overlay routes. *In Proceedings of IEEE International Conference on Network Protocols (ICNP)* (2006).
7. Keralapura, R., Taft, N., Chuah, C., Iannaccone, G.: Can ISPs take the heat from overlay networks?. *in Proceedings of HotNets-III*, San Diego, CA (Nov. 2004).
8. Qiu, L., Yang, Y. R., Zhang, Y., Shenker, S.: On selfish routing in Internet-like environments. *In Proceedings of SIGCOMM*, Karlsruhe, Germany (Aug. 2003).
9. Xie, H. et al: P4P: Provider Portal for Applications. *In Proceedings of ACM SIGCOMM 2008*, August 17-22, Seattle, Washington, USA (2008).
10. Sousa, P: Context Aware Programmable Trackers for the Next Generation Internet. *EUNICE 2009 - The Internet of the Future*, Barcelona, Spain, Springer, LNCS 5733, pp. 78-87, Barcelona, Spain (2009).
11. Legout, A., et al: Clustering and Sharing Incentives in BitTorrent Systems. *In Proceedings of ACM SIGMETRICS'2007*, June 12-16, San Diego, USA (2007).
12. Sousa, P.: Flexible Peer Selection Mechanisms for Future Internet Applications. *In Proceedings of BROADNETS 2009 - Sixth International ICST Conference on Broadband Communications, Networks and Systems*, Madrid, Spain (2009).
13. Opsahl, T., Agneessens, F., Skvoretz, J.: Node centrality in weighted networks: Generalizing degree and shortest paths. *Social Networks*, vol. 32, Number 3, pp. 245-251 (2010).
14. Narayanan, S.: The betweenness centrality of biological networks. *MSc Thesis*, Faculty of the Virginia Polytechnic Institute and State University (2005).
15. Rocha, M., Sousa, P., Rio, M., Cortez, P.: QoS constrained internet routing with evolutionary algorithms. *In Proceedings of IEEE Congress on Evolutionary Computation*, pp. 2720-2727 (2006).
16. Sousa, P., Rocha, M., Rio, M., Cortez, P.: Efficient OSPF Weight Allocation for Intra-domain QoS Optimization. In: Parr, G., Malone, D., O Foghlu, M. (eds.), IPOM 2006. LNCS, Vol. 4268, pp. 37-48. Springer, Heidelberg (2006).
17. ns-2 (The Network Simulator). Sources and Documentation from http://www.isi.edu/nsnam/ns/.
18. Eger, K., Hofeld, T., Binzenhofer, A., Kunzmann, G.: Efficient Simulation of Large-Scale P2P Networks: Packet-level vs. Flow-level Simulations. *In Proceedings of 2nd Workshop on the Use of P2P, GRID and Agents for the Development of Content Networks* (2007).