Improvement of BitTorrent Performance and Inter-Domain Traffic by Inserting ISP-owned Peers

Ioanna Papafili, Sergios Soursos, George D. Stamoulis

Athens University of Economics and Business,
Department of Informatics,
76 Patission Str., Athens 10434, Greece
{iopapafi, sns, gstamoul}@aueb.gr

Abstract. Large volumes of Internet traffic are nowadays generated by BitTorrent. In this article, we propose the insertion of high-bandwidth ISP-owned peers as an optimization approach to improve end-users' performance and reduce inter-domain traffic. An ISP-owned peer participates in BitTorrent swarms in order to download chunks and subsequently serve regular peers. We have run simulations on the ns-2 platform showing that our approach results in considerable reduction of both inter-domain traffic and the downloading times of users. We also show that the insertion of an ISP-owned peer can complement effectively the use of locality awareness, and lead to further performance improvements.

Keywords: BitTorrent, insertion, interdomain, QoE, locality.

1 Introduction

File-sharing applications are used widely by Internet users to share content such as music tracks, movies or even software releases. Because of their high popularity and the large size of the files that are shared, file sharing applications generate huge volumes of traffic in the Internet. This in turn implies a change of traffic patterns and an increase of costs (in terms of both CAPEX and OPEX) for the ISPs. In addition, monetary penalties related to the existing interconnection agreements are incurred and the nature of such agreements might also change, e.g. a peering agreement may have to be converted to a transit one, due to the change of traffic ratio between the peering ISPs. Consequently, it is necessary for the underlay network to take the overlay applications and their generated traffic into serious account in order to achieve efficient traffic management and optimal resource utilization in the underlying network. However, the ISP should meet his objectives in a way that is incentive compatible for the overlay provider, i.e. improving (or at least retaining) the overlay application performance. This is actually the topic of FP7-ICT project SmoothIT [1].

BitTorrent being the most popular file sharing application is the source for up to 60% of the overall traffic in the Internet. The BitTorrent protocol [2] was originally designed and implemented with the objective to disseminate one large file or a composition of large files to a large number of users without the original distributor incurring entirely the costs of hardware, hosting and bandwidth resources. BitTorrent can be deployed either by using trackers, or by using structured lookup overlays without trackers, the so-called trackerless BitTorrent. The tracker is a centralized component which stores information about all peers that participate in a swarm. Its main role is helping peers to discover other peers.

An optimization approach that has been employed in several research works is the so-called locality awareness. This amounts to biased selection of peers based on locality criteria; e.g. being in the same autonomous system as the requesting peer. In this paper, we propose and investigate an innovative approach to achieve a more efficient operation of the underlying network and therefore a cost reduction for the ISPs together with performance improvements for the BitTorrent users; namely, the insertion of a high-bandwidth ISP-owned peer, which aims to reduce both inter-domain traffic and downloading-completion times. The motivation for this approach stems from BitTorrent's tit-for-tat mechanism; due to the high upload capacity of the ISP-owned peer, regular peers establish connections to the IoP with higher probability than to other peers, thus resulting in performance improvements. Moreover, we consider the combination of the ISP-owned peer insertion together with locality awareness. In this case, even further performance improvement is expected due to the fact that local peers are more likely to select the ISP-owned peer to download from. The performance improvements attained by these approaches are evaluated by means of simulation experiments, which verify the aforementioned arguments. We also explain that pure locality awareness without the ISP-owned peer may not always be a beneficial approach for the ISP, contrary to what is widely argued in the literature. It should be noted that the insertion of an ISP-owned peer coincides neither with the use of an intervening cache as proposed in [10], nor with the enforcement of biased selection of peers as studied in the various articles overviewed in Section 2.

The paper is organized as follows: in Section 2, we present research works related to optimization of BitTorrent performance. In Section 3, we propose the insertion of ISP-owned peer as an optimization approach and discuss implementation issues. In Section 4, we describe our simulation model. In Section 5, we present and discuss our simulation results. Finally, in Section 6, we further discuss conclusions and open issues to be studied in the future.

2 Related Work

Peer-assisted content distribution is a cost-effective and bandwidth-intensive solution for ISPs. Indeed, peer-assisted, decentralized and self-organized systems such as BitTorrent provide significant benefits to end-users and content providers. However, such systems create their own logical networks and perform their own routing based on performance metrics, without taking into account the underlying topology. In this sense, overlay paths might end up reusing unnecessarily physical links or even containing circles. On the other hand, network management may not take into account the requirements of overlay applications, leading to a tussle between ISPs and Overlay Providers. This is due to information asymmetry, and may cause an increase of traffic on ISPs both intra- and inter-domain links and, as a result, higher costs. In [3] and [4], it is shown that this objective misalignment of ISPs and P2P networks in combination with the information asymmetry lead to performance degradation both for the underlay and the overlay. An ISP-friendly peer-assisted content distribution protocol that would exploit topology information is expected in [5] to alleviate ISPs' induced costs and as well as to improve P2P users QoE.

It can be intuitively expected that using topological information in a P2P system would significantly improve network performance (both overlay and underlay), if a better selection of a "good" server or a close-by peer, in terms of latency, were performed. This approach has been undertaken in several research works. In particular, in order to achieve this, a distributed, scalable binning scheme that requires a small number of landmark machines spread across the Internet is proposed in [6]. Because of landmarking being not self-organizing though, a new model is proposed in [7]. The model uses Geographical Longest Prefix Matching (Geo-LPM) and RTT to organize nodes into clusters each of which is a group of nodes that share a common prefix and are close to each other. The fact that clustered Geo-LPM is combined with an appropriate RTT threshold ensures that a node entering the network will find other overlay nodes that belong to the same physical domain. However, in case of clusters that share a common prefix, a solution is given by means of Geo-Partitioning. Furthermore, in [8] a 2-D Euclidean space model of the connectivity among BitTorrent peers has been proposed with the objective to evaluate BitTorrent's topology. Parameters such as number of peers in the swarm, maximum number of unchokes, etc., have been also taken into account. Moreover, in [9], a lightweight approach to reduce inter-ISP costs is proposed that exploits network information derived at low cost from CDN queries.

In [10], biased neighbor selection is studied as an approach to enhance BitTorrent traffic locality, in which a peer is enforced by the tracker to select the majority of his neighbors from peers within the same ISP and only a few (namely, *k* neighbors) that are outside the ISP. Additionally, the peer is modified to request a new list of neighbors whenever its peer list has less local peers than a specific threshold. This locality-awareness scheme can be implemented either by modifying tracker and client, or by situating P2P shaping devices along-side the edge routers of the ISPs, so that deep packet inspection is used to identify P2P traffic and manipulate it accordingly by intercepting and modifying the exchanged messages. Instead of enforcing locality, in [11] an 'oracle' is proposed that ranks peers according to some metric, e.g. proximity, bandwidth, etc., and provides this underlay information to users so that they can choose appropriate neighbors. In order to reduce downloading times of BitTorrent networks, also alternative chunk selection policies have been proposed in [12], while in [13] a cost-aware model to reduce both ISPs' costs and distribution time is proposed. The latter approach employs also alternative peer selection policies based on chunk availability on each peer. Finally, in [14] alternative peer selection based on RTT and number of hops is considered, which seems to reduce ingress inter-domain traffic as well as downloading times.

3 Insertion of ISP-owned Peers

An **I**SP-**o**wned **p**eer (IoP) is an entity that aims at increasing the level of traffic locality within an ISP and at improving the performance enjoyed by the users of peer-to-peer applications. The IoP, either belongs to an ISP's infrastructure and is controlled by the ISP itself; or is a regular but **h**ighly **a**ctive **p**eer (HAP) that is granted by the ISP with extra resources, e.g. higher downlink/uplink bandwidth, at no extra cost. In principle, if dynamic adjustment of the end-user's bandwidth is possible, then the end-users might even not be aware of this enhancement. However, agreement between the ISP and the HAP is also meaningful in order for the approach to be more effective; e.g., in order to assure extended seeding time by the HAP. In any case, the most important issue and what differentiates IoP from other related approaches, is that IoP runs the standard overlay protocol, e.g. BitTorrent, like every other peer in the swarm; yet, there are introduced certain changes in some parameters of the protocol that serve IoP's purposes and that are beneficial for other peers as well. In particular, the IoP is capable of unchoking more peers than the regular ones, in order to exploit its extra uplink capacity. Since the IoP runs the overlay protocol, it is also assumed that is capable of storing the content that it downloads and of course offering it back to the network. In other words, until an IoP has a complete copy of a file, it is considered to be a leecher in that file's swarm; subsequently, it is considered to be a seed. Henceforth, we use only the term ISP-owned peer, and include the HAP in this as well. The term "HAP' is used only when certain subtleties of this approach are discussed. Below, we distinguish two approaches for deploying an ISP-owned peer:

A. Plain insertion of IoP in a BitTorrent network: All peers are assumed to run the original BitTorrent protocol. No other mechanism such as locality awareness is employed by the ISP, and no agreement with the overlay provider is considered. Thus, the overlay, e.g. the tracker, is *not* aware of the IoP's existence as a special entity but treats it as a regular peer. In this case, the IoP is expected to be preferred by other peers due to the tit-for-tat mechanism employed by BitTorrent's unchoking algorithm and because of its high uplink capacity. The IoP follows here the tit-for-tat rule exploiting the immediate incentives of the latter that are directly related to the underlay [14].

B. Combination of IoP with locality-awareness mechanisms: The use of locality-awareness mechanisms that affect the overlay network's structure is considered here as being imposed by the ISP. Furthermore, depending on their implementation, these mechanisms could be either: a) transparent to the peers, i.e., they run along with the original protocol, or b) non-transparent i.e., they are introduced along with a modified version of the protocol. Metrics that can be used as proximity criteria are RTT and number of hops associated with remote peers, peers' autonomous system identity, BGP information, etc. Due to these locality-awareness mechanisms, the IoP would be mostly preferred by peers that are 'closer' to it according to one or more of the proximity criteria.

Below important issues regarding implementation are addressed:

Dimensioning of the IoPs: Dimensioning is expressed in terms of downlink/uplink bandwidth and storage capacity the IoP should be equipped with. Recall that the aim of the approach is twofold: meet the objectives of the ISP while coping with the users' performance requirements from the overlay application, since it is highly important not to downgrade their completion times.

Number of IoPs: Increasing the number of IoPs up to a certain number implies improvement of performance but also increases the CAPEX of the ISP. Additionally, the more IoPs exist in an ISP's network, the more intra-domain traffic is generated. After this traffic exceeds a threshold, more congestion on intra-domain links may result, thus leading to deteriorated performance and increase of OPEX for the ISPs. Thus, the number of IoPs should be carefully selected.

Physical location of the IoPs: The ISP should decide, based on the overlay traffic patterns, the physical locations where the IoPs should be deployed, e.g. 1) one “large” IoP in a specific location (centralized approach), 2) multiple “smaller” IoPs in a specific location (moderately centralized) or 3) multiple “smaller” IoPs in different locations (decentralized). Terms such as “large” or “small” refer to resources capacity. Before the selection of such an approach many issues require to be addressed such as availability, content duplication, etc. The location of the IoPs is related but not identical to problems on cache dimensioning and placement. Related techniques from that field could be employed.

Generally, both the number and the location of the IoPs within an ISP have to be decided by the ISP itself taking under consideration traffic measurements on inter- and intra-domain links, as well as impact of the traffic on the interconnection costs.

Content Selection. The ISP has also to make certain decisions that are expected to have impact on the efficiency of the IoP. First, the ISP should decide on which content will the IoP be downloading, i.e., in which swarms to participate. The selection of the content can be deployed either in a centralized or in distributed way. In the centralized cases, it could be performed with or without human intervention. In the distributed case, it would probably be more efficient, if it were performed automatically. In particular, content selection approaches could be:

1) Trial-and-error: The IoP could join randomly selected swarms in popular trackers, monitor whether his intervention has the desired impact for the ISP and decide whether to maintain its position, and/or when to leave a swarm etc.

2) Swarm-size based: The selection of content to be downloaded would greatly benefit from information provided by the overlay, e.g. trackers keeping statistics about the number of peers that participate in each file’s swarm.

3) Popularity-based: The underlying idea is that the IoP should download a file that is expected to become popular before other peers start asking for it even of the swarm size is originally small but expected to become larger.

Content legality. In the case of IoPs (excluding HAPs) the content downloaded is stored in ISP’s equipment. Thus, only licensed or non-copyrighted content can be downloaded by the ISP. Additionally, the ISP could establish agreements with content providers, e.g. content distribution networks, software vendors, music industry, movies distributors, TV channels, etc. and they should consequently establish agreements with the overlay in order to download, store and serve licensed content. On the other hand, in the case of HAPs no licenses or agreements are required since the content is stored in the users’ premises.

4 Simulation Model

Our simulation experiments were performed on the ns-2 simulator [16] using the BitTorrent patch [17] implemented by K. Eger. This patch contains four classes that implement a simplified version of the BitTorrent protocol that was not originally implemented for the ns-2 platform. We have modified several methods of the

BitTorrent classes in order to deploy a locality-aware BitTorrent protocol, which is employed in half of the experiments.

As a base topology, we use the Dumbbell Topology, e.g. a complex topology that comes from the interconnection of two simple star topologies (Fig. 4.1). Each star topology represents an AS; the left network is AS0 and the right network is AS1. We considered both the symmetric case (the two ASs have the same number of peers), e.g. two Tier-2 ISPs, and the asymmetric case (one AS has many more peers than the other), e.g. corresponding to a Tier-2 ISP and a Tier-3 ISP. Each peer within an AS is considered to be a regular BitTorrent peer and is connected to the router of its AS via a duplex asymmetric intra-domain link. One of the peers belonging to AS0 is considered to be the unique seed of the swarm and appears in the system at time 0. Furthermore, each peer becomes a seed after finishing its download.

In Table 4.1, we present all parameters used in the simulation experiments and their respective values.

Table 4.1: Simulation parameters.

Description	Value	Description	Value
Number of regular BT peers	50, 52	Number of local peers replied by tracker	20, 15
Number of initial seed	1	Number of autonomous systems (AS)	2
File size	20M	Number of IoP(s)	1, 2
Upload capacity of reg. peers	512K	Download capacity of regular peers	8x512K
Unchoked connections permitted per reg. peer	4	Unchoked connection permitted per IoP	20
Number of open connection per peer	20	Up-/download capacity of IoP	20x512K
Total number of peers replied by tracker	25	Inter-AS capacity (each direction)	30x512K

The main performance metrics of interest are the users’ downloading time and the ingress inter-domain traffic to both ASs. Downloading time is defined as the difference between the time when the peer received the last chunk of the file and the time when the peer sent a request to the tracker to get a first list of neighbors. Ingress inter-domain traffic is measured in the inter-domain link, i.e. the link that interconnects the two star topologies. We monitor traffic to both directions in order to make conclusions for interconnection costs following specific charging schemes. Namely, we measure all TCP traffic moving towards AS1 and denote this as ingress inter-domain traffic to AS1. The traffic is calculated from the original trace file generated by ns-2. We considered different scenarios in a variety of cases:

- **Pure BitTorrent scenario (BT):** The tracker replies a random list of peers to each peer’s request.
- **BitTorrent with locality awareness:** We assume that the tracker replies a localized list according to *Biased Neighbor Selection* approach of [9].
- **Insertion of IoP in BitTorrent:** One ISP-owned peer (down/up-load capacity: 20x512K, 20 unchokes) is inserted in AS1 in pure BT (Fig. 4.1).

- **Insertion of IoP in BitTorrent with locality awareness:** Combination of the two aforementioned scenarios (Fig. 4.2).

Below, we present further variations of the above scenarios:

- **Symmetric:** AS0 and AS1 have 25 regular peers each. The seed is in AS0.
- **Asymmetric:** AS0 has 35 regular peers, while AS1 has only 15.
- **All-together:** All peers' starting times are selected according to the uniform distribution U(0,10); note that 10 coincides with the choking interval.
- **Split:** The 5 latest peers' starting times of each AS are selected according to U(150,300); the IoP always follows U(0,10).

  The experimental results, e.g. the downloading times and the traffic volumes on inter-domain links, are presented in trace figures in the section 5.

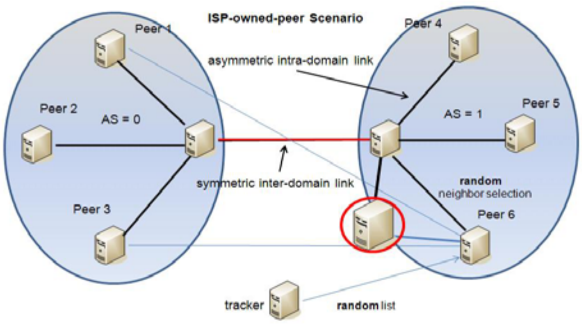

Fig. 4.1: Insertion of ISP-owned peer in pure BT.

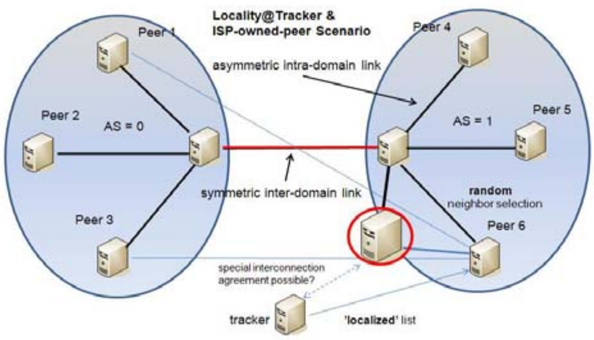

Fig. 4.2: Insertion of ISP-owned peer in BT with locality awareness.

5 Results

Reduction of inter-domain traffic: The insertion of ISP-owned peers achieves important reduction of the inter-domain traffic that enters the AS which deploys the IoP, namely in our scenario AS1. On the other hand, due to the fact that no constraints are posed, the IoP can serve peers outside its domain; thus, inter-domain traffic that exits AS1 towards AS0 will be increased. In Fig. 5.1 and Fig. 5.2, the ingress inter-domain traffic to AS1, in symmetric and asymmetric cases respectively, is shown. We note here that we consider only the case where peers start “all together”, but similar results hold also for the “split” case. In each case, we compare all four scenarios: pure BitTorrent, insertion of IoP in pure BitTorrent, BitTorrent and locality awareness and insertion of IoP in BitTorrent with locality. In the symmetric scenario, we observe that the insertion of the IoP achieves up to 35% inbound traffic reduction both in pure BitTorrent scenario (red-cyan lines) and in BitTorrent employing locality awareness (blue-green lines). Overall, we see that the combination of the IoP with locality awareness brings up to 53% improvement of the traffic compared to the pure BitTorrent scenario where no locality or IoP are considered. Note also that the IoP insertion on its own achieves similar results to pure locality awareness. On the other hand, in the asymmetric scenario, we observe up to 31% and 37% reduction of inbound traffic by the IoP insertion, in pure BitTorrent and BitTorrent with locality awareness respectively. In this case the IoP insertion on its own is *more* effective than pure locality. The gain reaches up to 44% when comparing the IoP combined with locality awareness and the pure BitTorrent scenarios. On the other hand, the traffic that enters AS0 increases up to 10-20% (not presented here due to space limitations) when compared with the respective non-IoP scenario, due to aforementioned reasons.

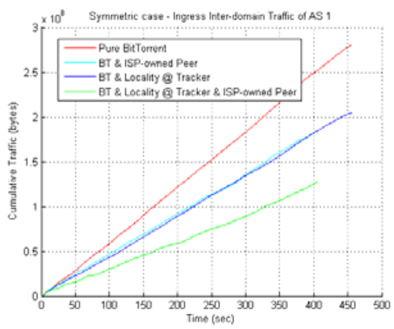

Fig. 5.1: Ingress inter-domain traffic to AS1 (symmetric case)

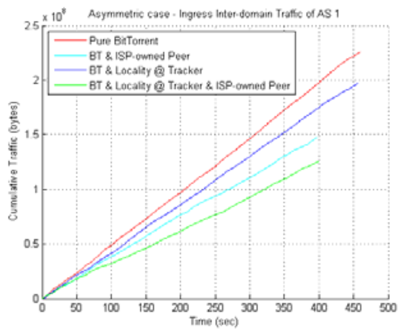

Fig. 5.2: Ingress inter-domain traffic to AS1 (asymmetric case)

Reduction of end-users’ completion times: In Fig. 5.3 and Fig. 5.4, we compare the end-users’ completion times for pure BitTorrent vs. insertion of IoP in pure BitTorrent, and BitTorrent with locality vs. insertion of IoP in BitTorrent with locality, respectively. In each figure are presented: (a) the simulation times and (b) the relative improvement of the completion times of the two scenarios compared. We have considered here the symmetric and “split” case, where the two ASs have same number of peers and some of the peers (peers with id 20 to 25 from AS0 and peers with id 45 to 50 from AS1) enter the swarm later. We observe that the insertion of IoP has important impact on the completion times. In particular, reduction of times up to 15% for peers starting along with the IoP (similar reduction is achieved also in ”all-together” case where all peers start along with the IoP), and reduction up to 35-40% for those starting later (spikes in Fig. 5.3 and Fig. 5.4).

Insertion of IoP vs. locality awareness; comparison w.r.t. reduction of charge for inter-domain traffic: As already noted, insertion of the IoP results in a higher reduction of inbound traffic, than locality awareness, particularly in the asymmetric scenario (see Fig. 5.2), which fits better to cases of transit agreements. Clearly, under a charging scheme for the inter-domain traffic that is based on statistics of the inbound traffic, the IoP insertion would lead to a higher reduction of charge than locality awareness. Furthermore, we turn attention to compare the effectiveness of the two approaches under charging models that are based on the *difference* between inbound and outbound traffic, e.g. using the 95th percentile rule[1]. Note that this 95th percentile charging scheme, like any difference-based scheme, is sensitive to asymmetric changes only, while symmetric changes have no direct impact on the costs.

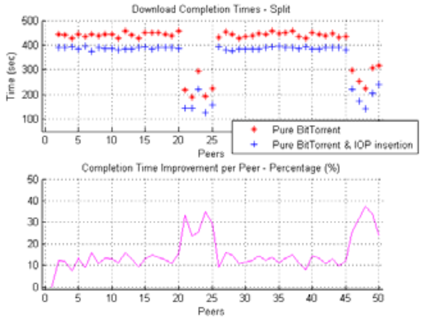

Fig. 5.3. End-users' completion times (scenarios without locality)

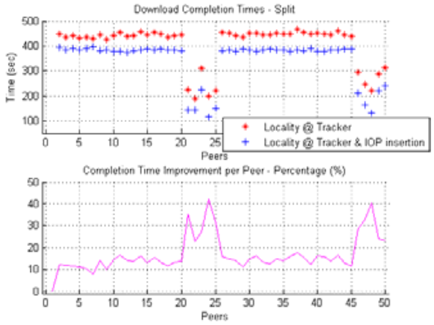

Fig. 5.4. End-user's completion times (scenarios with locality)

To this end, we present in Fig. 5.5 and Fig. 5.6 the instantaneous *difference* of inbound and outbound traffic to AS1, for pure BitTorrent and BitTorrent with locality, respectively. Again, we have considered here the asymmetric scenarios, i.e. AS0 has 35 peers and AS1 has 15 peers. We can observe that locality awareness on its own achieves more or less symmetric reduction of inbound and outbound inter-domain traffic. Indeed, the corresponding difference curve fluctuates around zero (Fig. 5.6 – top curve). This is due to the tit-for-tat mechanism that assures that the amount of traffic transferred to both directions is equivalent. On the contrary, the insertion of IoP in AS1 achieves asymmetric traffic reduction, regardless of whether locality is employed or not (Fig. 5.5 & Fig. 5.6 – bottom curves), due to the fact that the IoP quickly turns into a seed that serves peers regardless of the tit-for-tat mechanism which does not apply for seeds by the definition of the BitTorrent protocol. While locality does not affect this difference, the IoP clearly shifts the traffic difference in the favor of the AS deploying the IoP. When interpreting AS1 as a Tier 3-ISP and AS0 as a Tier 2-ISP the use of an IoP is beneficial for the Tier 3-ISP. Whether there is an actual monetary benefit depends on the OPEX and CAPEX for the IoP insertion, on the parameters of the charging model considered in each case and the achieved traffic reduction.

---

[1] The 95th percentile rule is applied every 5 seconds. The difference of inbound and outbound traffic is calculated and the upper 5% of that difference is cut away. The rest is what the ISP is charged for.

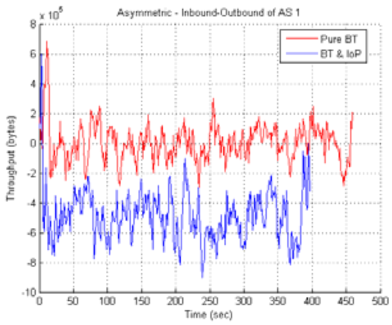

Fig. 5.5: Pure BitTorrent: Instantaneous difference between inbound-outbound to AS1

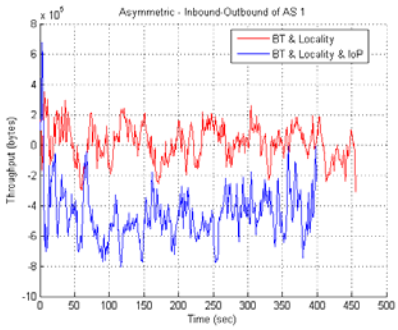

Fig. 5.6: BitTorrent with locality: Instantaneous difference between inbound-outbound to AS1

Impact of the IoP dimensioning on performance: Figures 5.7 and 5.8 present the ingress inter-domain traffic to AS1 and end-users' completion times, respectively, when an IoP is inserted in pure BitTorrent, for different values of capacity c assigned to the IoP. Recall that the number of unchokes of the IoP is equal to 20 and that its download and upload capacities are considered to be symmetric. In particular, in Fig. 5.7 traffic curves for $c = 10, 20, 30, 40$ x512 kbps are depicted. We observe that the traffic that enters the AS1 is generally decreasing when c increases. However, for $c = 40$ (cyan line), we see that the traffic is slightly higher than that for $c = 30$. This tradeoff is due to the fact that the IoP downloads more content from external peers before it becomes a seed. Furthermore, in Fig. 5.8 (top curve), the completion-time curves for $c = 10, 20, 30, 40$ and in Fig. 5.8 (bottom curve), the relative difference (%) of the completion times achieved for $c = 10, 30, 40$ compared to the times achieved for $c = 20$ are depicted. We can observe that, for $c = 10$, 5% worse times are achieved, whereas for $c = 20, 30, 40$, the completion times are similar. To summarize, providing more resources to the IoP is beneficial only up to a certain point.

6 Conclusions

In this paper, we have proposed and investigated the insertion of ISP-owned peers both in pure BitTorrent networks and in BitTorrent networks where locality awareness is also employed. The objective is to achieve both reduction of the inter-domain traffic caused by BitTorrent and reduction of downloading-completion times. Furthermore, we have conducted simulations for several scenarios in order to evaluate the performance implications of the IoP insertion and presented related results.

Simulations have shown that the insertion of the IoP achieves significant reduction of the inter-domain traffic that enters the AS where it is deployed. Further improvements are achieved when the IoP insertion is combined with locality-aware mechanisms. Moreover, the insertion of IoP in a pure BitTorrent network leads to higher inter-domain traffic reduction than just the use of locality awareness. On the other hand, the insertion of IoP achieves reduction of end-users' completion times in all cases that have been studied, whereas sole locality awareness implies slight performance degradation for end-users, as it was shown in [10] and was also observed in our simulation experiments. Furthermore, the symmetric reduction of inter-domain traffic achieved by locality awareness has no impact on interconnection costs when charging models based the difference of inbound-outbound traffic. On the contrary, the IoP achieves important asymmetric traffic reduction, which is expected to have also important impact on interconnection costs. Additionally, even when only inbound traffic is taken into account by the charging scheme, the IoP achieves further improvement and cost reduction than locality awareness.

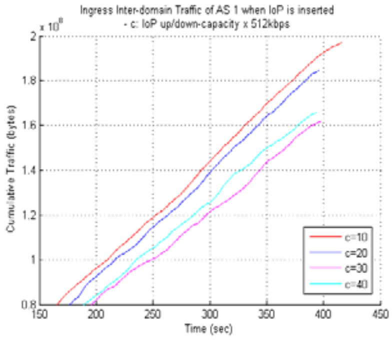

Fig. 5.7. Ingress traffic to AS1 for IoP capacity, c = {10, 20, 30, 40}x512kbps

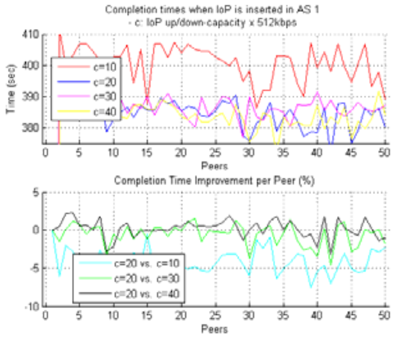

Fig. 5.8. End-user's completion times for IoP capacity, c = {10, 20, 30, 40}x512 kbps

The idea of the IoP insertion is related to the insertion of caches by the ISP, which store the content that is downloaded by peers, as considered in [5], [10]. However, the difference is that the solution of caches should be combined with interception of peers' messages whereas the IoP is part of the overlay itself. That is, it runs the overlay protocol, without requiring any enforcement. Therefore, communication between regular peers and the IoP is optional rather than being not enforced either at the application level or by means of special hardware. In this sense, the insertion of the IoP is an innovative idea. Of course, similarly to the case of caches, the IoP should only deal with legal content.

Furthermore, the insertion of IoP could be combined with bilateral agreements between the ISP and the overlay provider, or the ISP and the content provider (see also end of Section 3), which are also in line with the aforementioned legal issue. For instance, the overlay provider could favor the IoP when replying to peers' requests, e.g. by means of an IoP-aware overlay tracker. On the other hand, if the ISP has established some kind of agreement with a content provider, then its content can be stored directly in the IoPs and the torrent file generated would immediately contain as meta-info the IP addresses of the respective IoP. Essentially, the IoP acts as a seed, rather than as a cache that intercepts the requests. These kinds of agreements and related business models are under investigation.

Last, in this paper, we have restricted attention to the insertion of IoP in a BitTorrent file-sharing overlay. Investigation of the applicability of the IoP to the optimization of BitTorrent-like real-time or streaming applications is also currently in progress.

Acknowledgments This work has been accomplished in the framework of the EU ICT Project SmoothIT (FP7-2007-ICT-216259). The authors would like to thank all SmoothIT partners for useful discussions on the subject of the paper.

References

1. Simple Economic Management Approaches of Overlay Traffic in Heterogeneous Internet Topologies, European Seventh Framework Project FP7-2008-ICT-216259-STREP, http://www.smoothit.org/
2. Cohen B., Incentives build robustness in BitTorrent, In Proceedings of the First Workshop on the Economics of Peer-to-Peer Systems, Berkeley, CA, USA, June 2003.
3. Zhang H., Liu T., Gong W., Towsley D., Understanding the Interaction Between Overlay Routing and Traffic Engineering, 2004
4. Liu Y., Zhang H., Gong W., Towsley D., "On the Interaction Between Overlay and Underlay Routing," Proc. IEEE INFOCOM 2005
5. Karagiannis T.,Rodriguez P., Papagiannaki K., Should Internet Service Providers Fear Peer-Assisted Content Distribution? Proceedings of the Internet Measurement Conference 2005 on Internet Measurement Conference, 2005
6. Ratnasamy S., Handley M., Karp R., Shenker S., Topologically-Aware Overlay Construction and Server Selection, INFOCOM 2002
7. Le H., Hoang D., Simmonds A., A Self-Organizing Model for Topology-Aware Overlay Formation, Seoul, 2005
8. Jia J., Chen C., GMT: A Geometry-based Model on BitTorrent-like Topology, 8th International Conference on Signal Processing, 2006
9. Choffnes D.R., Bustamante F.E., Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems, SIGCOMM 2008
10. Bindal R., Cao P., Chan W., Medved J., Suwala G., Bates T., Zhang A., Improving Traffic Locality in BitTorrent via Biased Neighbor Selection, 26th IEEE International Conference on Distributed Computing Systems, p. 66, 2006
11. Aggrawal V., Feldmann A., Scheideler C., Can ISPs and P2P Users Cooperate for Improved Performance?, SIGCOMM 2007
12. Wu C.J., Li C.Y., Ho J.M., Improving the Download Time of BitTorrent-Like Systems, IEEE International Conference on Communications, p.1125-1129, 2007
13. Yamazaki S., Tode H., Murakami K., CAT: A Cost-Aware BitTorrent, 32nd IEEE Conference on Local Computer Networks, p. 226-227, 2007
14. Li W., Chen S., Yu T., UTAPS: An Underlying Topology-Aware Peer Selection Algorithm in BitTorrent, 22nd International Conference on Advanced Information Networking and Applications, p. 539-545, 2008
15. Oechsner S., Soursos S., Papafili I., Hossfeld T., Stamoulis G.D., Stiller B., Callejo M.A., Staelhe D., A framework of economic traffic management employing self-organization overlay mechanisms, International Workshop on Self-Organizing Systems, Vienna, December 2008
16. The Network Simulator – ns-2: http://www.isi.edu/nsnam/ns/
17. Eger K., Simulation of BitTorrent Peer-to-Peer (P2P) Networks in ns-2: http://www.tu-harburg.de/et6/research/bittorrentsim/