Collaboration between Peering ISPs for Economic Management of Overlay Traffic

Eleni Agiatzidou and George D. Stamoulis

Athens University of Economics and Business, Department of Informatics

{agiatzidou, gstamoul}@aueb.gr

Abstract

The SIS (SmoothIT Information Service) entity has been introduced by FP7 project SmoothIT, to support Economic Traffic Management (ETM) for overlays. It is employed by ETM mechanisms in order to convey information between the overlay application and the underlay network. In this paper we present two different ways of cooperation between the SIS entities of different ISPs, in the case that these two ISPs maintain a peering agreement with each other.

1. Introduction

Overlay applications (e.g. BitTorrent) generate large volumes of Internet traffic. FP7 Project SmoothIT [1] has already developed ETM mechanisms, aiming to improve the performance of the overlay applications experienced by users and to reduce the high costs that are imposed to ISPs due to the traffic that is produced by these applications (Win-Win situation). Such traffic is forwarded via logical overlay connections without taking into account the physical network. BGP-based Locality Promotion is an ETM mechanism that employs the SIS entity, which provides the application end users a ranked list of peers according to BGP information. Thus, locality of traffic is promoted, and the aforementioned goals are often attained. Locality promoting approaches are also presented in other work e.g. [4] and references within. In this paper we provide an extension of the above ETM mechanism, introducing the cooperation between SISes that belong to different ISPs with a peering agreement. This allows the traffic exchange between them without introducing any charges, as long as the ratio of traffic in the two directions is close to 1.

In Figure 1 we present an abstract network topology with many ISPs. Two Tier 3 ISPs (3A and 3B) have a peering agreement and also both of them are connected by symmetric transit links to a Tier 2 ISP (2A). The traffic that is exchanged through these transit links imposes charges to Tier 3 ISPs by Tier 2 ISP, unlike the traffic that they exchange over their peering link. In order to reduce those charges but also to improve performance of the users in ISPs 3A and 3B, we propose an approach that aims to reduce the download of redundant (duplicate) content through the costly inter-domain links with ISP 2A. This approach referred to as "Splitting Chunks", works under the assumption that peers from 3A and 3B download the same content through the same transit link. Then each ISP downloads from peers outside the domains of 3A and 3B a different subset of chunks of the same content (swarm). These chunks are then exchanged between the two peering ISPs over the peering link. For example, a peer from 3A downloads from peers in 3B all the even chunks of the content, while he downloads all the odd chunks from the rest of the remote ASes. At the same time he uploads to peers in 3B the odd chunks that he has already acquired. Peers in 3B act similarly. Next, in section 2 we introduce "Splitting Chunks" and present certain experimental results in section 3. In section 4 we present conclusions and directions for further research.

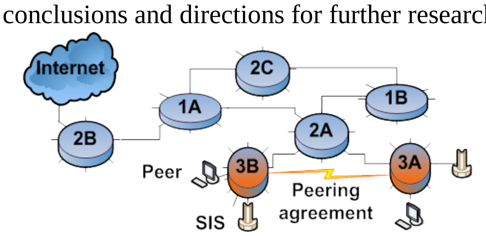

Figure 1. Abstract Network Topology

2. Splitting Chunks approach

As already explained, under the proposed approach, the peering ISPs (3A, 3B) can share transit costs, by deciding to split the downloading of different parts of content. The split can be done independently of the number of the peers that participate in the swarm from each ISP. Moreover ISPs have to decide which chunks of the content each of them will download from remote ASes. The chunks can be distinguished statically according to their ids (i.e. even – odds), which is a policy that can be maintained for all swarms regardless of the size of the content. Alternatively, the splitting of chunks can be decided dynamically by the SISes according to the progress of the collaborative downloading in the swarm. This scenario is not examined in this paper. The rest of the chunks that each ISP does not download from the remote ASes will be retrieved from the peering link. This restriction may lead to an increase in download times. Of course, local peers can exchange any chunks in order to promote locality. The SIS still maintains its basic functionality [2], i.e. assists the local peers by rating the list of peers using BGP locality information, thus offering peers the choice to connect to better (close by) neighbors. Under our approach, every peer also needs to be informed by the SIS about what chunk is authorized to download from remote peers, because the collaboration does not apply for each swarm. Thus, when the peer is informed about the new available chunks of a remote peer, then he declares his interest to him according to the chunks that he is authorized to download from that particular peer, depending on the domain where the latter belongs. Furthermore when this peer is unchoked, the chunks that he asks for are consistent with what he is authorized to download. As each peering ISP will download fewer chunks using the transit link, it will be charged less by the upper Tier. A preliminary evaluation of this approach using a special simulation tool is presented in the next section.

3. Experimental evaluation

In order to evaluate the aforementioned approach we used the simulator that was developed by the SmoothIT project and is based on the Protopeer simulator [3]. We implemented a network topology such as that in Figure 1, with thirteen ISPs. Eight of them are Tier 3 ISPs, with peers being evenly distributed among them. Two of these ISPs (namely, 3a and 3b) have a peering agreement. Initially, we ran the regular BitTorrent scenario and a simple scenario of BGP-Locality approach that uses the SIS for rating the peers in a swarm, in order to have a basis for comparison. Then, we ran the scenario of SIS collaboration with Splitting Chunks approach maintaining the same parameters as in the previous scenarios and modifying accordingly the way the peers in the peering ISPs 3a and 3b download chunks. During the various runs we focused on the inter-domain traffic between these ISPs of Tier 3 and the Tier 2 ISP 2A. Also we monitored the downloading times for the peers that belong to the peering ISPs. Our results are presented in Figures 2 and 3.

In particular, in Figure 2 we depict the 95th percentile of the inter-domain traffic volume that passes from ISP 2A to ISPs 3A and 3B. The charge to be incurred by each Tier 3 ISP in the corresponding transit link is proportional to this statistic. Clearly, the BGP Locality approach leads to very high reduction of the traffic in the transit link compared to the regular BitTorrent. Our approach runs together with BGP locality reduces this traffic even more as expected. In Figure 3, we depict the mean downloading times for the peers of the peering ISPs 3A and 3B. The mean download time for the Splitting Chunks approach is increased compared to regular BitTorrent. Nevertheless, the increase is less than 2%. Thus, in the cases examined the Splitting Chunks approach leads to a Win/Non-lose situation for the ISPs and the applications users.

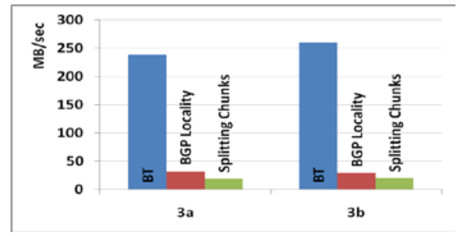

Figure 2. Inter-domain traffic between Tier 2 and Tier 3 ISPs

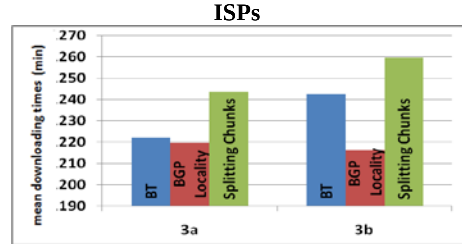

Figure 3. Downloading times for peers in Tier 3 ISPs

4. Conclusions and further work

Splitting chunks is an extension of the BGP Locality mechanism that aims to the mutually beneficial collaboration of the SISes between peering ISPs. The preliminary evaluation of the approach using a simulator has shown a reduction in the inter-domain traffic between the Tier 2 ISP and the Tier 3 peering ISPs. Thus, as the downloading times are slightly increased, the approach seems to achieve a Win/Non-lose situation.

Our future research will focus on further simulations for this approach, using also more dynamic splitting policies and more detailed study and evaluation of the incentive that the users have in order to adopt this approach. Furthermore we will evaluate other approaches for SIS collaboration, such as the splitting of the remote ASes that participate in the swarm.

5. Acknowledgments

This work has been performed partially in the framework of the EU ICT Project SmoothIT (FP7-2007-ICT-216259). Furthermore, the authors would like to thank all SmoothIT project members for providing insights and their discussions.

6. References

[1] http://www.smoothit.org/
[2] Racz, P et al. Economic Traffic Management for Overlay Networks. In: ICT-MobileSummit 2009 Conference Proceedings
[3] http://protopeer.epfl.ch/
[4] Haiyong Xie, et al. P4P: Provider Portal for Applications In: SIGCOMM 2008 351-362