ISP-driven Practical P2P Traffic Control Technique

HyunYong Lee
The University of Tokyo
Tokyo, Japan
ifjesus7@gmail.com

Akihiro Nakao
The University of Tokyo
Tokyo, Japan
nakao@iii.u-tokyo.ac.jp

***Abstract*—We propose one practical approach to control BitTorrent traffic by exploiting existing features of BitTorrent: peer exchange (PEX) and tit-for-tat (TFT). For the traffic localization within target network domain, our approach injects a set of local peers (i.e., peers in the target network domain) to each local peer through PEX that allows BitTorrent clients to exchange their neighboring peer list with each other. In addition, we try to redirect unavoidable inter-domain traffic from transit link to peering link by affecting TFT strategy. We add an artificial delay to the inter-domain traffic of the transit link going to the local peers to affect TFT strategy. Our simulation shows that our approach achieves the traffic localization and redirects the inter-domain traffic while improving the content download performance.**

I. Introduction

There have been various P2P traffic control techniques to manage inter-domain P2P traffic. ISP has applied various approaches such as rate limiting [1]. These ISP-driven unilateral approaches are not so practically satisfactory due to the reactions of users (e.g., dynamic port change, data encryption) to avoid being identified [3], since these approaches usually degrade content download performance. There also have been P2P-driven approaches [2] relying on inferring network information such as topology and congestion status. By utilizing the network information, P2P applications try to localize P2P traffic. However, there are fundamental limits on what P2P can achieve alone. Reverse engineering of the network information, in particular cost and ISP's traffic engineering policy, is challenging if not impossible. As an alternative way to satisfy both ISP and P2P users by reflecting their major considerations, bilateral approaches based on the explicit collaboration between users and ISP [4], [5] have been proposed. In this approach, ISP generates a guidance reflecting its traffic engineering policy as well as accurate network information and provides the guidance to P2P applications. Then, P2P applications utilize the guidance to select communication partners. This approach is regarded as a win-win approach, since users can enjoy better content download performance with accurate network information and ISP can reflect its traffic control intention flexibly. However, there are practical limitations for deployment, since it requires modification of existing P2P systems.

In this paper, we propose practical win-win approaches for BitTorrent [6] that still accounts much portion of Internet traffic (e.g., in Asia-Pacific area during 2012 1H, BitTorrent accounts around 27% of total Internet traffic [10]). Instead of modifying existing BitTorrent system, our approach exploits two features of BitTorrent: PEX and TFT strategy. With PEX, BitTorrent clients can exchange a list of their neighboring peers with each other. Our approach injects a list of local peers to each local peer through PEX to increase a chance of local communications within the target network domain. In addition, to control unavoidable inter-domain P2P traffic (e.g., download of content that is not available in the target network domain), our approach adaptively adds artificial delay to inter-domain traffic going to local peers. With TFT strategy, BitTorrent client uploads the content to the peers who have uploaded content to itself in the past at high bandwidth. Therefore, we may be able to turn inter-domain traffic of transit link into inter-domain traffic of peering link if we can degrade the download throughput through transit link with the delay insertion.

Our simulation results show that the manipulation of neighboring peer set based on PEX increases the intra-domain traffic volume. This leads to the improved content download performance. The adaptive delay insertion increases the inter-domain traffic volume of peering links while decreasing the inter-domain traffic volume of transit links. This result shows that our approach can be a win-win approach that satisfies both users and ISP. In addition, our approach is easily deployable, since we do not need to modify BitTorrent.

The rest of this paper is organized as follows. Section II introduces our approach and brief comparison with existing approaches. Section III evaluates our approach and Section IV concludes this paper.

II. Practical P2P Traffic Control Technique

Our approach supports two main functions for traffic control of BitTorrent: *topology-aware PEX* for the traffic localization within target network domain and *adaptive delay insertion* for redirection inter-domain traffic from non-preferred inter-domain links to preferred inter-domain links. Two functions depend on crawling scheme in common, since they need peer information for their work (Fig. 1). For the sake of simplicity, in this paper, we assume that target network domain and target swarm are given. A swarm that has more than certain number of peers in the target network domain can be the target swarm. The crawler collects IP address and port number of the local peers by contacting BitTorrent tracker. The crawler repeats its work periodically to adapt peer dynamics. After crawling, if there exist any new local peers, the crawler provides IP address and port number of local peers to the topology-aware PEX component and the adaptive delay insertion component. We skip detailed description about the crawler due to the space limit, even though we have developed our own efficient crawler. Please refer [7], [8] for our crawler.

978-3-901882-50-0

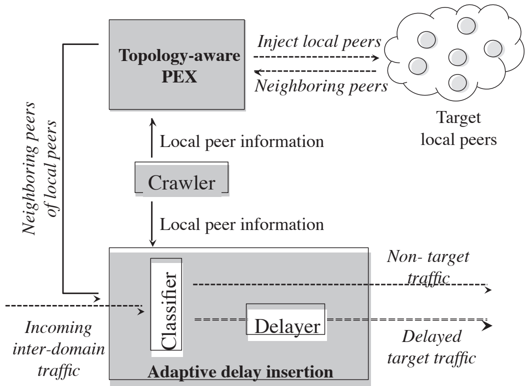

Fig. 1. Conceptual illustration of proposed system.

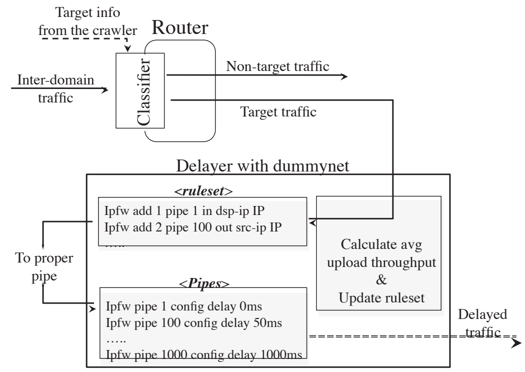

Fig. 2. The adaptive delay insertion component.

A. Topology-aware PEX

To achieve the traffic localization within the target network domain, we utilize PEX. PEX is one popular extension to BitTorrent and most existing BitTorrent clients support PEX (i.e., around 95% [9]). In the original design of the BitTorrent, the centralized server called tracker returns certain number of existing peers to (newly joining) peers. On the other hand, PEX allows each peer to exchange neighboring peers with each other directly.

Based on the peer information given by the crawler, the topology-aware PEX component contacts each local peer and injects a set of local peers randomly chosen through PEX. Only difference between the topology-aware PEX and existing normal PEX is that neighboring peer set to be exchanged only includes local peers in the topology-aware PEX component. Our intention of the topology-aware PEX is to increase a chance of local communications within the target network domain. Due to TFT strategy, BitTorrent clients prefer to communicate with other BitTorrent clients in the same network domain, since performance of local communication is usually better than performance across network domains. The topology-aware PEX component also receives neighboring peer information from the local peers through PEX. Then, it forwards the neighboring peer information to the adaptive delay insertion component.

B. Adaptive Delay Insertion

The adaptive delay insertion component utilizes TFT strategy to control the inter-domain traffic. Every 10 seconds, with TFT strategy, BitTorrent client selects (i.e., *unchokes*) peers (up to certain number, e.g., 4 is default) to upload a content for next 10 seconds. For the peer selection, BitTorrent client cares about download performance from its neighboring peers during previous 10 seconds. BitTorrent client selects the peers who have uploaded content to itself in the past at high bandwidth. Therefore, delay insertion to inter-domain traffic may affect the peer selection procedure of BitTorrent. For example, if we add the artificial delay to inter-domain traffic of transit link enough to degrade download throughput worse than download throughput through peering link, local BitTorrent client may select BitTorrent clients beyond peering link. To redirect inter-domain traffic from non-preferred links to preferred links, the artificial delay is added to inter-domain traffic of non-preferred links when the inter-domain traffic goes to local peers who know peers beyond both the preferred and the non-preferred links. The adaptive delay insertion component can identify those target local peers based on the peer information given by the crawler and the topology-aware component.

Even though the topology-aware PEX achieves the traffic localization by injecting a set of local peers, there exits unavoidable inter-domain traffic. The topology-aware PEX just injects a set of local peers while not removing existing remote neighboring peers from the local peers. In BitTorrent, peer communication is affected by a content availability. The local peers sometimes need to download chunk (i.e., fixed-sized segments of content being shared) from remote peers when the chunk is not available from local neighboring peers. In addition, all neighboring peers are not candidate peers to be selected based on TFT strategy every 10 seconds. Only peers who sent INTERESTED message are considered. A peer sends INTERESTED message to its neighboring peer that has the chunk that it does not have.

The adaptive delay insertion component can be divided into two subcomponents: traffic classifier and delayer (Fig. 2).

Traffic classifier: The traffic classifier classifies inter-domain traffic into target traffic (going to the local peers) and non-target traffic. Then, the traffic classifier forwards the target traffic to the delayer. For the traffic classifier, we can utilize existing DPI tools [1].

Delayer: The delayer adds certain amount of delay to the detected target traffic according to a delay insertion policy (that will be explained later). For this, we utilize dummynet [11]. Dummynet can create multiple communication links (i.e., *pipe*) with different setting for bandwidth and propagation delay. Dummynet also has a list of numbered rules (i.e., *ruleset*) to decide their fate. Therefore, after creating multiple pipes with different delay, ruleset is updated dynamically according to the delay insertion policy. Then, when the target traffic is given, the target traffic is forwarded to proper pipe for the delay insertion.

Delay insertion policy: Every second, the delayer calculates average upload throughput through the preferred links

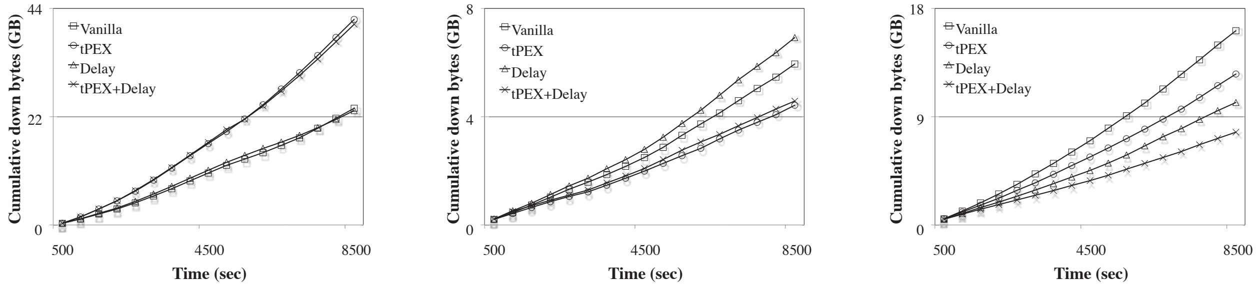

(a) Intra-domain traffic (b) Through peering links (c) Through transit links

Fig. 3. Cumulative amount of download traffic.

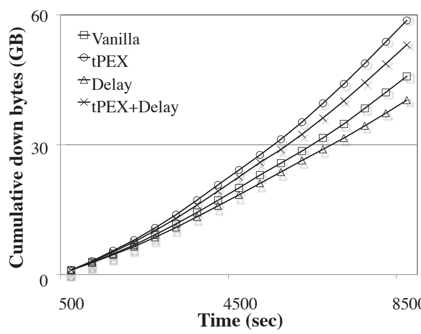

Fig. 4. Cumulative amount of download traffic (total).

from one remote peer to one local peer during previous 10 seconds. Then, the delayer finds the minimum upload throughput through the preferred links to each local peer. If upload throughput through non-preferred link from one remote peer to one local peer becomes larger than the minimum upload throughput through the preferred link to the same local peer, the delayer adds delay to degrade the upload throughput through the non-preferred link to lower than the minimum upload throughput. Due to the space limit, we skip to detail how to calculate the delay to be added. If there is no upload through the preferred link to certain local peer, no delay is added to the inter-domain traffic of non-preferred link going to the local peer. The delayer does not know when each local peer selects new peers for upload by applying TFT strategy. But, as a heuristic approach, every second, the delayer calculates the average upload throughput during previous 10 seconds, since TFT strategy considers the download throughput during previous 10 seconds.

C. Comparison with Existing Approach

Common procedure of most existing bilateral approaches [4], [5] is that users select their communication partners based on the guidance given by ISP. The given accurate network information allows users to enjoy better content download performance. In addition, the explicit collaboration between users and ISP enables ISP to control P2P traffic as it wants. However, these approaches have practical limitations that make deployment difficult. For example, they require the dedicated server (e.g., iTracker of [4]) to provide the guidance to the users and the explicit interface to communicate with the users. They also need to modify P2P system to reflect the guidance in selecting communication partners. ISP needs an elaborate way to generate the guidance not to disclose its sensitive information.

On the other hand, rather than depending on the explicit user collaboration that accompanies several practical limitations, our approach exploits existing features of BitTorrent. Therefore, we do not need to modify existing BitTorrent. We do not need to worry about the disclosure of sensitive network information. In addition, our approach improves user content download performance through the traffic localization while providing a way for control inter-domain traffic to ISP (see next Section). The improved content download performance may be able to be an incentive for users not to react.

III. Evaluation

We build ns-2 based simulation environment with 24-hour BitTorrent swarm measurement. The swarm we used includes 6,493 peers in total. The peers are distributed over 674 ASes. As a target AS, we chose one AS that has 535 local peers and 2 peering and 5 transit links. Two peering links are the preferred links. We set 100KB/s for uplink capacity and 300KB/s for downlink capacity of peers. In our simulation, user link capacity is a performance bottleneck. BitTorrent tracker returns 50 peers randomly selected. The topology-aware PEX injects at most 20 local peers every 600 seconds. Each peer has 4 unchoke slots. For performance comparison, we compare four different cases: BitTorrent without any scheme (Vanilla), topology-aware PEX only (tPEX), adaptive delay insertion only (Delay), and combination of two approaches (tPEX+Delay) as our proposed solution. For the clear presentation, we show the simulation results up to around 10,000 seconds.

We first examine traffic volume. Fig. 3 and Fig. 4 show the cumulative amount of downloaded bytes over time. The intra-domain traffic volume is mainly affected by the topology-aware PEX (Fig. 3(a)). Even though each local peer still knows remote peers after the topology-aware PEX works, the local peers tend to communicate with each other more often than before the injection of local peers due to the TFT strategy. The topology-aware PEX decreases the inter-domain traffic by increasing the intra-domain traffic (Fig. 3(b) and Fig. 3(c)).

On the other hand, the adaptive delay insertion does not have any noticeable impact on the traffic localization

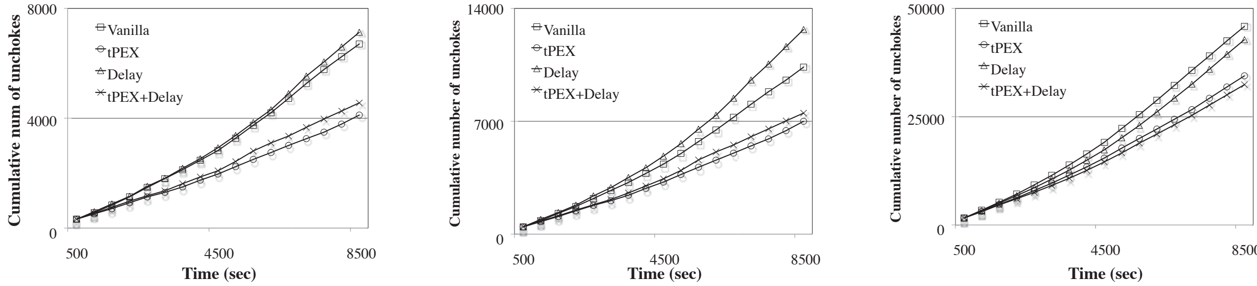

(a) Unchokes sent through peering links (b) Unchokes received through peering links (c) Unchokes received through transit links

Fig. 5. Cumulative number of unchokes through inter-domain links.

(Fig. 3(a)). Instead, the adaptive delay insertion decreases the inter-domain traffic through transit links (Fig. 3(c)). At the same time, the delay insertion increases the inter-domain traffic through peering links (Fig. 3(b)). This result shows that the adaptive delay insertion can affect the inter-domain traffic by adding some delay to inter-domain traffic of specific link. To support this observation and to examine how the delay insertion affects TFT strategy, we measured a number of unchokes (Fig. 5). The delay insertion increases a number of unchokes for peers beyond peering links (Fig. 5(a)). It means that more local peers upload chunks to remote peers through peering links rather than through transit links. In return, more local peers are unchoked (i.e., selected) by peers beyond peering links (Fig. 5(b)) and this leads to the increased amount of download through peering link (Fig. 3(b)). At the same time, less local peers are unchoked by peers beyond transit links (Fig. 5(c)) and this leads to the decreased amount of download through transit link (Fig. 3(c)). This shows that the delay insertion can redirect inter-domain traffic between inter-domain links by affecting peer selection based on TFT strategy.

Now, we examine the content download performance (Fig. 6). The number on top of each bar indicates a number of download completions. Basically, the delay insertion does not increase the intra-domain traffic while decreasing the inter-domain traffic through transit links. Even though the inter-domain traffic through peering links increases, amount of increase is less than the decrease of transit link case. Due to this reason, the adaptive delay insertion decreases the total traffic volume (Fig. 4) and thus degrades the download performance compared to Vanilla. Due to the improved traffic localization with the manipulated neighboring peer set, the topology-aware PEX increases the total traffic volume and thus improves the download performance. On the other hand, when the topology-aware PEX and the adaptive delay insertion are applied together (as our proposed solution), the combination of two approaches increases the total traffic volume. Our approach improves the download performance and also increases a number of download completions compared to Vanilla.

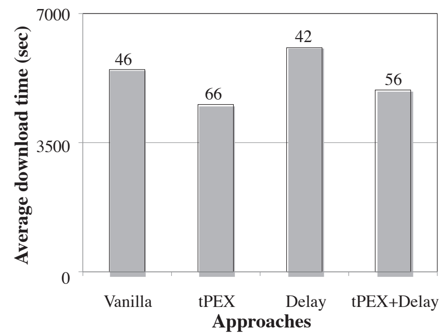

Fig. 6. Download performance.

IV. Conclusion

We propose one practical BitTorrent traffic control technique that exploits PEX and TFT strategy of BitTorrent. The topology-aware PEX achieves the traffic localization by manipulating neighboring peer set. The adaptive delay insertion redirects the inter-domain traffic by affecting TST strategy with the delay insertion. Simulation results verify that our approach can be a win-win approach. As future work, we plan to study an effect of various things including a number of peers injected by the topology-aware PEX and peer distribution over neighboring ASes.

References

[1] Sandvine, http://www.sandvine.com/.

[2] Kontiki, http://www.kontiki.com/.

[3] L. Jun, Z. Shunyi, L. Shidong, and X. Ye, Active p2p traffic identification technique, in: Proceedings of IEEE Computational Intelligence and Security, 2007.

[4] H. Xie, Y. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz, “P4P: Portal for (P2P) applications,” in *Proc. of ACM SIGCOMM*, 2008.

[5] F. Lehrieder, S. Oechsner, T. Hobfeld, D. Staehle, Z. Despotovic, W. Kellerer, and M. Michel, “Mitigating unfairness in locality-aware peer-to-peer networks,” *International Journal of Network Management*, vol.21, no.1, pp.3-20, 2011.

[6] B. Cohen, “Incentives build robustness in BitTorrent,” in *Proc. of IPTPS*, 2003.

[7] Masahiro Yoshida and Akihiro Nakao, “Measuring bittorrent swarms beyond reach,” in *Proc. of IEEE P2P*, 2011.

[8] M. Yoshida and A. Nakao, “A resource-efficient method for crawling swarm information in multiple bittorrent networks,” in *Proc. of International Symposium on Autonomous Decentralized Systems*, 2011.

[9] D. Wu, P. Dhungel, H. Xiaojun, C. Zhang, and K.W. Ross, “Understanding peer exchange in bittorrent systems,” in *Proc. of IEEE P2P*, 2010.

[10] SANDVINE Report, Global Internet Phenomena Spotlight (Asia-Pacific, Fixed Access), http://www.sandvine.com/general/document.download.asp?docID=54&sourceID=0.

[11] Dummynet, http://info.iet.unipi.it/ luigi/dummynet/.