BitTorrent Traffic Localization using a Multi-layered DHT

Kyung-Hwa Kim
Columbia University
New York, New York
khkim@cs.columbia.edu

Moritz Steiner
Alcatel-Lucent, Bell Labs
Holmdel, New Jersey
moritz@bell-labs.com

Matteo Varvello
Alcatel-Lucent, Bell Labs
Holmdel, New Jersey
matteo.varvello@alcatel-lucent.com

***Abstract*—Peer-to-Peer (P2P) file sharing applications such as BitTorrent generate a huge amount of expensive inter-ISP traffic. Traffic localization can be achieved by modifying the peer-selection strategy either at the client or at the *trackers* (BitTorrent servers). Both approaches require cooperation between ISPs and P2P networks which is unlikely to occur. Moreover, tracker-based traffic localization mechanisms are becoming ineffective as the role of the trackers is being replaced by Distributed Hash Tables (DHTs). In this work, we design and demonstrate a tracker-less traffic localization mechanism that guarantees a desired user download speed. Our design enforces communication between local peers by mean of a multi-layered DHT, where the original DHT network is extended with several topology-based DHTs. As a proof of concept, we implement the multi-layered DHT as an extension for the Vuze BitTorrent client. This demonstration shows the benefits of the multi-layered DHT while downloading several files from the Vuze network**

I. Introduction

Network traffic generated by Peer-to-Peer (P2P) file sharing applications (such as BitTorrent) accounts for a large fraction of the entire Internet traffic. The majority of P2P traffic is exchanged among end-hosts located at different ISPs, causing considerable monetary loss for the ISPs. This is due to the common design rationale of most P2P applications: peers download from peers located world-wide in order to increase path diversity and avoid bottlenecks [8].

Currently, *client-based* and *tracker-based* strategies are used to localize BitTorrent traffic. Client-based designs [2], [3], [6] modify the peer-selection strategy at the client in order to favor local peers from the peers discovered via the central trackers, the DHT and the peer-exchange-protocol[1]. These designs achieve low traffic localization as they only select the few, if any, local peers from the peer-set received. Tracker-based designs [11], [1] leverage the trackers (i.e., BitTorrent central servers which coordinate file exchanges) to return to a peer only peers located within its ISP. These designs achieve much higher traffic localization than client-based designs. However, these approaches will soon become ineffective as trackers are rapidly being replaced by DHTs [7].

In this work, we design and demonstrate a novel traffic localization strategy that leverages the DHT. This design is complementary to tracker-based traffic localization schemes.

[1]The interested reader can find more details about BitTorrent peer discovery mechanisms in [4].

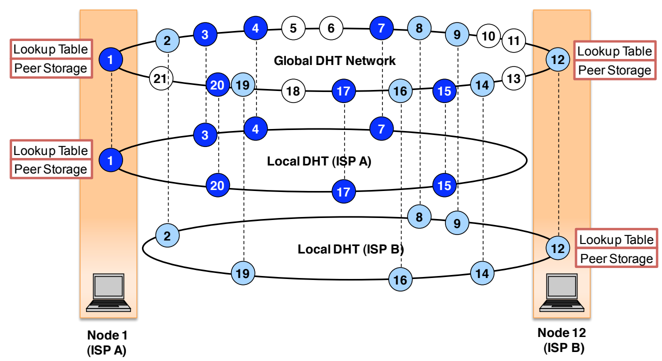

Fig. 1. Multi-layered DHT Architecture.

Our design is a multi-layered P2P system that contains a global DHT and multiple local DHTs. A peer joins the global DHT regardless of its ISP, whereas only peers located at the same ISP join the same local DHT. Our design has three main advantages: (1) local DHTs enable traffic localization; (2) the global DHT preserves the opportunity to contact worldwide peers in case the lack of fast local peers increases download durations; and (3) our design is backward compatible with BitTorrent as we do not modify the current DHT architecture.

II. Design Overview

Our design rationale is to enforce communication between local peers by mean of a local DHT. However, rather than modifying the existing DHT network (as in the hierarchical DHT [5] and topology-aware DHTs [10]) we suggest to maintain several extra DHT networks that we call *local DHTs*.

A local DHT is a logical overlay network which has the same functionalities of a classic DHT. However, its participants are limited to the peers which are located at the same ISP. Consequently, there is a single global DHT network formed by all active peers and there are multiple disjoint local DHTs which are constructed by each group of peers located at the same ISPs. Peers join both the global and the local DHT using the same identifier in the hash-space (see Figure 1).

A. Bootstrapping

We use the global DHT in order to bootstrap the local DHTs. Every peer publishes a special key-value pair in the global DHT: the key contains the name of the ISP at which the peer is located and the value contains the peer network address.

When a peer does not know any active local peer, it retrieves from the global DHT a list of local peers using its ISP name as a key.

B. Storing Peers and Lookup

The local and global DHTs are disjoint and independent networks. Thus, the operations of a local DHT such as routing, querying and announcing are separated from the operations of the global DHT. A peer aiming to announce that it holds a file or a portion of it needs to send *announce* messages to both DHTs. Peers which receive the announce messages store the peer information in different peer storages (global or local) depending on which DHT interface they received the message from. For example, Figure 2 shows that peer 7 in the local DHT-A and peer 7 in the global DHT, which are physically the same peer, have different peer lists for File 1 in their global and local storages, respectively.

In order to perform a lookup operation, a peer first sends *lookup* messages to the local DHT. The peer list returned from the local DHT contains local peers only as the local DHT only index local peers. This ensures that, if available, local peers are favored over global peers. If the peer discovers enough local peers to guarantee a download speed equal or faster than a *desired download speed*, which can be configured by the user in advance, the peer does not need to discover any additional global peer. Conversely, if the download speed is lower than the desired one, the peer initiates a lookup operation in the global DHT to retrieve information about worldwide peers. Note that given a peer has the same identifier in the local and global DHT, a global lookup can leverage the previous local lookup instead of completely traversing the global DHT.

C. Discussion

The hierarchical DHT [5] and topology-aware DHTs [10] focus on lookup localization in order to reduce the lookup latency. Conversely, we focus on localizing file transfer traffic, which is obviously heavier than lookup traffic. In addition, the local DHTs in our design do not have to be merged to form a global DHT as for the hierarchical DHT. Because our approach is adding an independent layer of local DHTs to an existing global DHT, our design is backward compatible with BitTorrent. We believe this feature is an important contribution of our design as it allows to direclty target millions of BitTorrent users. Nevertheless, the flexibility of our design allows the local network to be either a structured or unstructured P2P network.

III. Demo Setup

Vuze [9] is the most popular open-source BitTorrent client. In order to prove the backward compatibility of the multi-layered DHT, we develop it as a Vuze plugin, i.e., we do not modify the Vuze source code. In this demonstration, we visually show the benefits (traffic localization and download speed) of the multi-layered DHT using our Vuze plugin.

The goal of this demonstration is to compare a classic DHT (single DHT in the following) with a multi-layered DHT in terms of traffic localization and download speed. We proceed as follows:

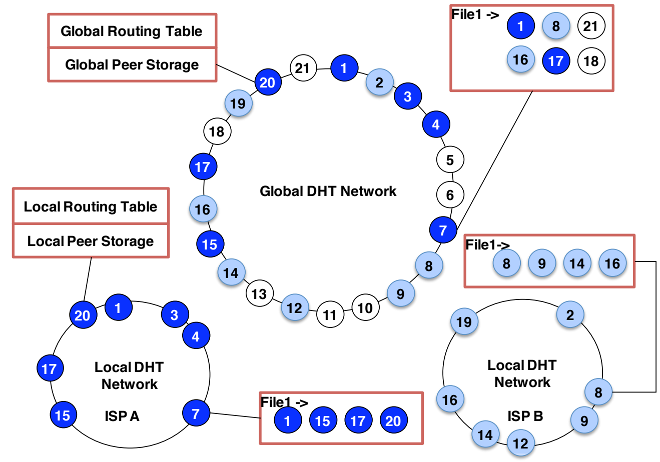

Fig. 2. Peer storage.

1) Single DHT - choose several popular files in the Vuze network; machine A located at ISP-A initiates the files download.
2) Single DHT - plot the download speed and ratio of local/unlocal traffic.
3) Turn on the 'multi-layered DHT' function at several peers located in ISP-A. These peers construct their own local DHT using the global DHT for bootstrapping. From these machines, begin to download the files previously requested my machine A.
4) Turn on the 'multi-layered DHT' function at the peer we run on machine A. Resume the file downloads. Plot the new download speed and traffic localization.

References

[1] V. Aggarwal, A. Feldmann, and C. Scheideler. Can ISPs and P2P users cooperate for improved performance? *CCR*, 37(3):29–40, 2007.

[2] Choffnes, Bustamante. Taming the torrent: a practical approach to reducing cross-isp traffic in peer-to-peer systems. In *Proceedings of ACM SIGCOMM*, SEATTLE, WA, AUGUST 2008.

[3] Dabek, Cox, Kaashoek, Morris. Vivaldi: a decentralized network coordinate system. In *Proceedings of ACM SIGCOMM*, AUGUST 2004.

[4] Di Wu, Dhungel, P., Xiaojun Hei. Understanding peer exchange in bittorrent systems. In *Proceedings of P2P*, Aug 2010.

[5] Prasanna Ganesan, Krishna Gummadi, and Hector Garcia-Molina. Canon in g major: Designing dhts with hierarchical structure. In *Proceedings of ICDCS 2004*, Washington, DC, USA, 2004. IEEE Computer Society.

[6] Ledlie, Gardner, Seltzer. Network coordinates in the wild. In *Proceedings of NSDI*, AUGUST 2007.

[7] Matteo Varvello and Moritz Steiner. Traffic localization for dht-based bittorrent networks. In *Proceedings of IFIP/TC6 Networking*, Valencia, Spain, May 2011.

[8] Michael Piatek, Harsha V. Madhyastha, John P. John, Arvind Krishnamurthy, Thomas Anderson. Pitfalls for isp-friendly p2p design. In *Proceedings of ACM HotNets*, New York, NY, October 2009.

[9] Vuze, Inc. Vuze. http://www.vuze.com.

[10] Marcel Waldvogel and Roberto Rinaldi. Efficient topology-aware overlay network. In *Proceedings of Hotnets-I*, 2002.

[11] Haiyong Xie, Y. Richard Yang, Arvind Krishnamurthy, Yanbin Grace Liu, and Abraham Silberschatz. P4P: provider portal for applications. In *Proceedings of SIGCOMM*, Seattle, WA, USA, August 2008.