Computer Networks 55 (2011) 1242–1256

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

LANC: Locality-aware network coding for better P2P traffic localization

Guoqiang Zhang [a,*], Guoqing Zhang [a], Suqi Cheng [a,b,**]

[a] Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
[b] Graduate University of Chinese Academy of Sciences, Beijing, China

ARTICLE INFO

Article history:
Received 8 April 2010
Received in revised form 28 October 2010
Accepted 5 November 2010
Available online 7 December 2010
Responsible Editor: I.F. Akyildiz

Keywords:
Locality-aware network coding
P2P traffic localization
Network coding
Locality-awareness
Beyond P4P

ABSTRACT

As ISPs begin to cooperate to expose their network locality information as services, e.g., P4P, solutions based on locality information provision for P2P traffic localization will soon approach their capability limits. A natural question is: can we do any better provided that no further locality information improvement can be made? This paper shows how the utility of locality information could be limited by conventional P2P data scheduling algorithms. Even sophisticated data scheduling algorithms such as the local rarest first could not take full advantage of the locality information.

Network coding's simplified data scheduling makes it competent for improving P2P application's throughput. However, its potential for P2P traffic reduction has not been fully recognized. This paper proposes the locality-aware network coding (LANC) that uses locality information in both the topology construction and downloading decision, and demonstrates its exceptional ability for inter-domain P2P traffic reduction. The randomization introduced by network coding enhances the chance for a peer to find innovative blocks in its neighborhood. Aided by proper locality-awareness, the probability for a peer to get innovative blocks from its proximity will increase as well, resulting in more efficient use of network resources. Extensive simulation results show that LANC can significantly reduce inter-domain P2P traffic without sacrificing application-level performance. Aided by the same locality knowledge, the inter-domain traffic of LANC in most cases are less than 50% of the current best approach that does not use network coding. We also show that this property holds when the tit-for-tat incentive mechanism is introduced or the overlay topology changes dynamically.

1. Introduction

In recent years, peer-to-peer file distribution applications such as BitTorrent[1] have emerged as promising solutions and attracted wide-spread attentions among Internet users due to its scalability with user population and resilience to flash crowds as well as dynamic user arrivals and departures. However, they also pose serious challenges to the network operators because of the heavy traffic generated by them. Measurement data show that P2P applications have already become the dominant network bandwidth consumer, contributing more than 50% to the overall network traffic [1,2]. In some cases, the P2P traffic consumes so much bandwidth that it starves other services [3].

One major cause of the tremendous traffic is the inefficient use of network resources. P2P systems are based on an overlay topology on top of the underlying network, upon which content searching, and sometimes routing, is performed. This leads to problems for both the P2P applications and network operators. Being largely network-oblivious, overlay network construction either disregards or relies on limited knowledge of underlying network topology, which inevitably leads to inefficient overlay

* Principal corresponding author. Tel.: +86 10 62600721.
** Corresponding author.
E-mail addresses: guoqiang@ict.ac.cn (G. Zhang), chengsuqi@ict.ac.cn (S. Cheng).

[1] http://www.bittorrent.com/.

1389-1286/$ - see front matter © 2010 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2010.11.014

topology and routing, and consequently sub-optimal application-level performance. The problems for network operators are twofold: (1) the dynamics of P2P traffic makes ISPs hard to manage these traffic and unable to do effective traffic engineering [5,6]; and (2) P2P traffic often crosses network boundaries multiple times [7,8], which can either incur increased costs for its inter-domain traffic or cause traffic imbalance with its peers, potentially violating the contractual peering agreements [6]. In addition to the possible rising cost and contract violation, the increase of inter-domain traffic is even more serious in that inter-domain links, rather than the ISP backbone links, are the major source of network bottlenecks [9].

Solutions towards relieving the P2P traffic burden can be coarsely categorized into three phases. In the initial phase, the tough reality makes some ISPs to unilaterally throttle or rate limit P2P applications using techniques such as deep packet inspection (DPI). In reaction, the P2P application developers use techniques such as dynamic ports or message encryption to hide their traffic. Hence, a vicious circle emerges, which is dramatized as a *war* between the P2P content providers and ISPs in the IETF. In the second phase, P2P application developers and ISPs begin to independently seek approaches to localize the P2P traffic. The P2P applications can obtain coarse-grained locality knowledge via reverse-engineering to aid its overlay network building and downloading [16,12,8,13–15,17], whereas the ISPs can leverage the widely deployed caches to localize the P2P traffic [11,8,10]. Recently and in the third phase, P2P content providers and ISPs begin to cooperate to make efficient use of the network resources [4–6], which is believed to be the only way to relieve the tension between them, e.g., P4P [6].

As ISPs begin to provide locality information as network services, the performance gain arising from the accuracy of network locality information will soon approach its limit. A natural question arises: can we do any better, i.e., more efficient network resource utilization, beyond P4P? The chance lies in the way how data is scheduled and propagated in the network. Traditional P2P paradigm either uses random policy or relies on local information, e.g., local rarest first policy, to make data scheduling decisions, which may lead to biased distribution of data blocks in the network. As a result, no matter how accurate the locality information is, the data scheduling and propagation mode inherently limits further performance improvement.

Network coding [23] allows intermediate nodes not only to receive, replicate and forward, but also to do arbitrary coding operations on packets. By applying network coding to P2P content distribution, the data blocks transmitted on the network are no longer the original blocks, but coded blocks of multiple random original blocks. This way, network coding greatly simplifies the data scheduling and enhances the application-level throughput [31,33]. However, it is argued that the computation overhead incurred by the encoding and decoding processes may offset or even outweigh the benefits [24], which significantly impedes its wide adoption and deployment.

However, standing at ISPs' point of view, the randomization introduced by network coding improves the probability for a peer to find innovative blocks from its neighborhood, which, if aided by proper locality knowledge, can increase the chance of finding and retrieving innovative blocks from its proximate neighbors. Therefore, network coding could be a powerful tool to enhance the efficiency of network resource utilization, with the ability to breach the capability limits of state-of-art techniques. The main contributions of this paper are listed as follows:

(1) We show how the utility of locality information could be limited by conventional data scheduling algorithms, whereas this limitation can be greatly avoided by network coding based scheduling.
(2) We propose a novel locality-aware network coding (LANC) data scheduling which integrates the locality-awareness and network coding as an effective approach to reduce inter-domain P2P traffic redundancy. Simulation results show that aided by the same locality knowledge, LANC can reduce the inter-domain P2P traffic by more than 50% in most cases.
(3) We provide considerations for a practical LANC implementation, such as optimal parameter settings and balancing between the computation overhead and the system throughput.

The subsequent paper is organized as follows: Related work of network coding and P2P traffic localization are discussed in Section 2. A simple motivation example is presented in Section 3 to illustrate how data scheduling affects and how network coding improves P2P traffic localization. We present the model used for our study in Section 4. The simulator is described in Section 5 and the experimental evaluation results are presented in Section 6. Implementation considerations are discussed in Section 7. Finally, we conclude the paper in Section 8.

2. Related work

2.1. Network coding

Network coding is a paradigm shift from the conventional information transmission and processing mode. By allowing the intermediate nodes to perform arbitrary coding function on the input data, it emerges as a promising technology to realize the theoretical multicast capacity upper bound predicted by the max-flow min-cut theorem [23]. It is demonstrated that linear codes on Galois finite field can achieve this rate [25,26].

The concept of random network coding [28] paves the way for the practicality of network coding. Random network coding allows each node to autonomously select its code, as opposed to the centralized assignment in deterministic algorithms [25,27], hence facilitating the distributed implementation in a large and dynamic environment. Consequently, there is a gradual research shift from the theoretical point of view to the more practical settings [29–34], particularly applying network coding to P2P content distribution [31,33,34]. Avalanche has demonstrated