

Improving BitTorrent Traffic Performance by Exploiting Geographic Locality

Chen Tian¹ Xue Liu² Hongbo Jiang¹ Wenyu Liu¹ Yi Wang¹

¹Department of EIE, Huazhong University of Science and Technology, Wuhan, Hubei, China

²School of Computer Science, McGill University, Montreal, Quebec, Canada

¹{tianchen,hxj,liuwuy,ywang}@mail.hust.edu.cn, ²xueliu@cs.mcgill.ca

Abstract—Current implementations of BitTorrent-like P2P applications ignore the underlying Internet topology hence incur a large amount of traffic both inside an Internet Service Provider (ISP)’ national backbone networks and over cross-ISP inter-networking links. These traffics not only occupy costly bandwidth, but also increase user perceived response latency. ISP-Biased Neighbor Selection proposes to exploit peers’ topological locality by biased neighbor selection, in which a peer chooses the majority of its neighbors from peers within the same ISP. In this paper, we propose to further exploit peers’ geographic locality. First we improved ISP-Biased Neighbor Selection (ISP-Biased+) to take into consideration network locality (or, city locations) within the same ISP. When required neighbor number is relatively much less than seeds available, ISP-Biased Neighbor Selection+ performs much better than original approach, proved by simulations. Next, we propose that a peer could also choose its neighbors from peers of different ISPs within the same city with priority: assist by a well-know Chinese operator’s unique ISP-internetworking content distribution network (CDN), these local cross-ISP traffics can be routed through local CDN cite. Using simulations, we show that cross-ISP traffic burden can be completely shifted to CDN local links and backbone traffic. At the same time, user perceived delay can be significantly reduced.

I. INTRODUCTION

Peer-to-peer network is the most popular platforms for Internet applications like file sharing and streaming media. Nowadays P2P applications, leading by BitTorrent, account for over 70% of the traffic in Internet [12]. The problem is that current implementations of BitTorrent-like P2P applications ignore the underlying Internet topology and set up data transfer connections among randomly chosen sets of peers. As a result, a significant amount of traffic are generated both inside an ISP’ national backbone networks and over cross-ISP inter-networking links [2], [14]. These traffics increase the operating cost of ISPs significantly, plus that traffic between an ISP and the outside world is very costly [3]. In addition, it is undesirable that they increase user perceived response latency. As a result, ISPs often adopt locality enhancing approaches, such as bandwidth limiting (throttling), gateway peers, and caching. Nonetheless, all these passive methods are not satisfactory. For example, the main effect of throttling is only to slow down content transfer and deteriorate user experiences.

Recently, ISP-Biased Neighbor Selection [4] proposes to exploit peers’ topological locality by actively helping peers’ neighbor selection: a peer chooses the majority of its neighbors from peers within the same ISP. By taking Internet

topology into consideration, this approach can significantly reduce cross-ISP traffic, as well as improve the download performance. However, this approach did not take ISP backbone traffic into consideration, which is also costly. In this paper, we first improve ISP-Biased Neighbor Selection to take access network locality (city locations) within an ISP into selection consideration (ISP-Biased+). Our experiments demonstrate that, when required neighbor number is relatively much less than seeds available, ISP-Biased Neighbor Selection+ performs much better than original approach.

The next generation networks of A well-know vendor in China¹ currently being built is supposed to change the Internet topology again, at least in China. By inter-networking to almost every ISPs in the same city, this CDN becomes the biggest, if not only, ISP-internetworking network in China. Assisted by this unique network, we can extend peers’ locality from topological field to geographic field: a peer could choose the majority of its neighbors not only within the same ISP, but also from peers of different ISPs but within the same city. Our experimental results show that cross-ISP traffic can be completely shifted to CDN local links and backbone. At the same time, user perceived delay can be significantly reduced. To our best knowledge, this is the first scheme exploring local ISP inter-networking benefits.

The remainder of this paper is organized as follows. We give the background and exploit the relative works in Section II. In Section III we present the basic ideas of ISP-Biased+ neighbor selection and CDN-assist P2P neighbor selection and routing. Then we demonstrate and compare the effectiveness of these approaches by simulations in Section IV. In Section V, we give some discussions. Finally, we represent our conclusions and future works in Section VI.

II. BACKGROUNDS AND RELATED WORKS

A. BitTorrent and its locality

BitTorrent-like P2P application [9] is designed to distribute large files to a large user population efficiently. To distribute a file, the provider must run a *tracker* for this file to keep track of all peers downloading this file. Potential downloader P must first contact the tracker to join the BitTorrent dedicated overlay *torrent* of this file. This overlay network consists of the tracker,

¹Due to privacy concerns, our collaborators request us not to disclose their company name

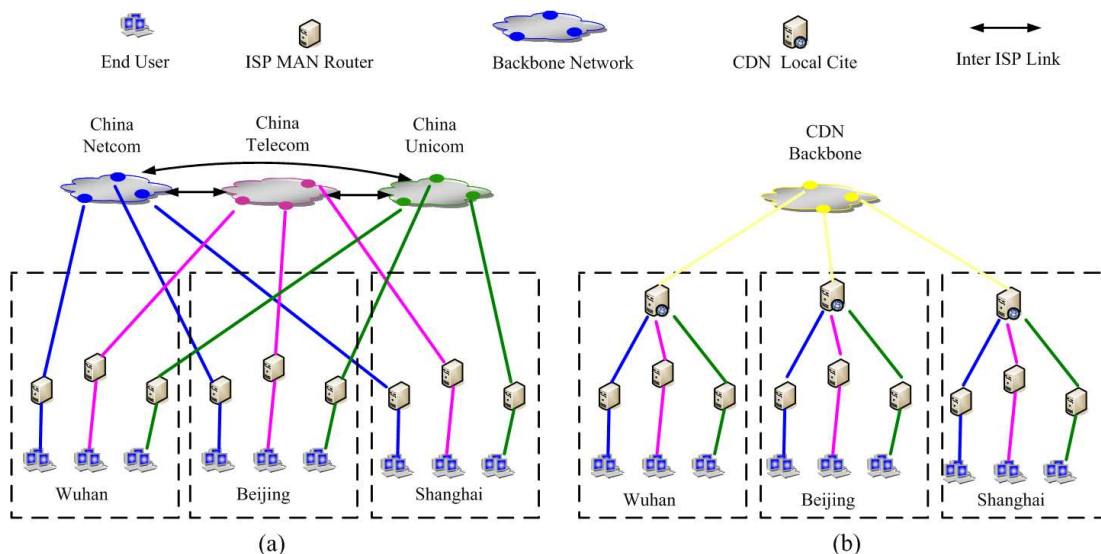


Fig. 1. (a) ISPs network topology (b) New architecture of next generation CDN in China

the provider, and all peers already downloaded or interested in downloading the file. The tracker then randomly selects N peers out of all the peers in the torrent as P 's neighbors, and sends the list back to P . P then initiates connections with those neighbors. Later on, another peer joining the torrent may get P as one of its neighbors returned by the tracker and initiate connections with P . In short, P 's neighbors in the torrent include both peers that P initiates connections to and those that initiate connections to P . The problem is: peers of BitTorrent-like P2P applications ignore the underlying Internet topology and choose neighbors randomly among all peers in the network. This is the reason of BitTorrent's high backbone and cross-ISP traffic. What's more, selfish P2P routing often conflicts with ISP controlled network routing policies [13] and may cause much loss to ISPs.

To control this traffic, ISPs often deploy *throttling*, or bandwidth limiting. Throttling reduces the rate of cross-ISP traffic at the cost of increased download time hence it deteriorates user experiences. It does not address the fundamental problem of the BitTorrent-like P2P application: neighbor selection. Besides bandwidth limiting, caches [1], and gateway peers [5] are also adopted in some ISPs. It is found that for these solutions, to avoid increasing download latency, the caches or gateways need to have much higher bandwidth than individual peers, and the bandwidth requirement should grow together with the number of peers inside the ISP [4].

ISP-Biased Neighbor Selection (ISP-Biased) [4] developed a new approach to enhancing BitTorrent traffic locality. The main technique is biased neighbor selection: a peer chooses its neighbors mostly from those within the same ISP, and only a few neighbors that are outside the ISP. Hence, peers within the same ISP form a cluster and are highly connected. By taken Internet topology into consideration, ISP-biased neighbor selection can reduce cross-ISP traffic significantly and improve the download performance. However, this approach did not

take ISP backbone traffic into consideration, which is costly. In Japan, over 62% backbone bandwidth is consumed by P2P traffic [14]. Suppose a peer of China Netcom in Beijing tries to choose neighbors within the ISP, as shown in Figure 1(a). It is obvious that other peers in Beijing should have priority over those in Shanghai or Wuhan. It is also obvious that access network locality should be taken into consideration, especially in giant Chinese cities with millions of users in one ISP. In this work we improve this approach by further considering city locality inside an ISP.

Yet question arose, as we also have many ISPs in the same city, can we exploit the information of geographic locality too? If these ISPs' local access networks are connected together, a large amount of traffic could be shifted from ISPs' national backbone networks and cross-ISP internetworking links to local area networks. And, users of different ISPs but in the same city are more possible to share same interests, hence organize them into clusters may be more efficient.

B. Content Distribution Network

Content Distribution Network (CDN) technology has been proposed to deliver content to proxies placed at strategic locations on the Internet. To make the content as close as to users, CDN proxies are traditionally been deployed along with edge routers of ISPs, such as Akamai [6] and ChinaCache [7].

There are 6 major ISPs in China, including China Netcom, China Telecom, China Unicom, China Mobile, CTT, and CERNET. Besides, there are many other small ISPs: Great Wall Net, Founder Net, etc. Users in any ISPs want to watch the live network broadcast of 2008 Beijing Olympic Games, but it is obvious that deploying separate proxy clusters for each ISP are costly and uneconomical.

The new architecture of a leading CDN currently being built in China is shown in Figure 1(b). CDN's local cite is the center of a local star network in a city: each ISP is connected to this

CDN cite by a high speed local link. Above city level, all nodes are connected by CDN’s private fiber backbone. Deliberately or not, this CDN in fact becomes the biggest, if not only, ISP-internetworking network in China. This new generation CDN may answer our question: now we have an applicable way to exploit peers’ geographic locality for BitTorrent-like P2P applications.

III. EXPLORING GEOGRAPHIC LOCALITY

The main technique we use is biased neighbor selection and CDN-assist neighbor routing. First, we classify all candidate neighbors of a peer to 4 categories: (1) Same ISP, Same City; (2) Different ISP, Same City; (3) Same ISP, Different City; (4) Different ISP, Different City.

We use Figure 2 to illustrate the chance of a peer being chosen as neighbors of another peer in each category. Original BitTorrent-like P2P applications choose neighbors randomly among all peers in the network. As shown in Figure 2(a), peers in each category have equal chance of being selected. In ISP-Biased Neighbor Selection [4], peers in same ISP have priority: $(1, 3) > (2, 4)$. This is illustrated in Figure 2(b).

Our ISP-Biased+ Neighbor Selection take access network locality (city locations) within same ISP into selection consideration: peers in the same access network have priority over those in others cities. The result is shown in Figure 2(c): most chosen neighbors are within the same ISP, and the priority is $(1) > (3) > (2, 4)$.

Further, we propose that a peer could also choose its neighbors from peers of different ISPs within the same city with priority: assist by a next generation CDN’s unique ISP-internetworking content distribution network, these local cross-ISP traffics can be routed through local CDN cite. Therefore, peers within the same city form a cluster and their cross-ISP traffic can be routed by CDN local network.

We sum up the CDN-assist algorithm: peers are chosen in the priority $(1) > (2) > (3) > (4)$ as shown in Figure 2(d). Neighbor traffic of category (1) are carried inside an ISP’s city network; category (2) are carried by CDN local network between ISPs; category (3) can be carried by CDN backbone, but its performance should be comparable or inferior to ISP’s own backbone, plus cost consideration, this traffic should be carried by ISP itself; as inter-ISP internetworking in China is notorious, category (4) should be carried by CDN backbone to improve user perceived response latency.

IV. PERFORMANCE EVALUATION

A. Simulation Setup

A discrete event simulator is developed to evaluate the performance of various algorithms. We have implemented and simulated all 4 algorithms: random, ISP-Biased, ISP-Biased+ and CDN-assist. Due to size limitation, only results of latter three are given.

There are 6 ISPs and 3 cities in the simulation scenarios. The network topology is organized similar with Figure 1(b). To be simplified, we assume every ISP and city is equal in network size. However, the numbers of ISPs and cities

participating in each simulation can be adjusted. Based on a well known profile [15], devices are connected with 3 types of link models:

- Local link, connect devices in same city, delay 20ms
- Backbone link, connect core devices of one company into backbone, delay 100ms
- Cross-ISP link, connect different ISPs, delay 300ms

As our main object is to compare neighbor selection and routing algorithms in traffic performance improvement, no bandwidth limiting or other P2P shaping approach are implemented in the simulations. A number of resources are simulated, and existing seeds of one resource is set to proportional to whole P2P system sizes or, ISP number (as we assume equal network sizes for ISPs). To be exactly, we set total existing seeds number of a resource is 10 times ISP number and randomly distribute them to ISPs and cities.

In each simulation, we randomly choose a peer to randomly request a new resource with given required number of neighbors, denote by N . By different algorithms, this peer gets its neighbors and starts to download. The main evaluation criteria are: network Traffic Burden and response Averaged Delay. Because the locations of each neighbor decide network traffic in this download process, we use average number of neighbor in each category to represent Traffic Burden. The arrival times of every first packet from each chosen neighbor is recorded and are averaged as Averaged Delay. Each data point shown in a graph is the mean of 10,000 simulation runs.

B. Evaluations

First, we measure Traffic Burden and Averaged Delay as functions of required number of neighbors. We have 3 ISP and 3 cities in this scenario as shown in Figure 1(b). There are total 30 existing seeds for each resource. We change required neighbor number from 6 to 20.

Averagely, each ISP has 10 seeds. ISP-Biased and ISP-Biased+ should have no difference in Cross-ISP traffic. When required number less than 10, ISP-Biased+ should perform better than ISP-Biased due to its local-city first policy. When required neighbor number exceeds 10, all available seeds inside requester’s own ISP are used, and their performance should be comparable.

Our assumptions are verified in Figure 3(a) and (b). When N is less than 10, ISP backbone traffic of ISP-Biased+ is significantly lower than that of ISP-Biased. In both algorithms, when N becomes bigger than total seeds and ISP could provide, both ISP local and backbone traffic reach an upper bound and cross-ISP traffic soared rapidly. Because available neighbors inside the peer’s belonging ISP is limited, it should turn to outside seeds.

In contrast, there is no cross-ISP traffic in Figure 3(c). This traffic is completely transferred to CDN local links and backbone traffic. Averagely, each city has 10 seeds. That’s why local traffic reaches an upper bound of 10. When N becomes bigger and local available neighbors exhausted, backbone traffics of ISP and CDN soared. However, these traffics cost




Fig. 2. Chance of a peer being chosen as neighbors




Fig. 3. (a)ISP-Biased (b)ISP-Biased+ (c)CDN-assist neighbor selection and routing

lower than cross-ISP traffic and the response delay is much lower to users too.




Fig. 4. Response Delay Comparison

Figure 4 of Response Delay Comparison again confirm our analysis. When N is less than 10, ISP-Biased+ performs better than ISP-Biased as it has more low latency local city traffic. When N is bigger than 10, performances of ISP-Biased and ISP-Biased+ are almost the same. On the contrary, CDN-assist always has superior low latencies.

Next, we measure Traffic Burden and Averaged Delay as the functions of number of ISPs. We change ISP numbers from 2 to 6. Accordingly, the available seeds number increase from 20 to 60. The required neighbor number N is fixed to 15. Because ISP-Biased and ISP-Biased+ performs almost the same when N is bigger than 10, here we only compare ISP-Biased+ with CDN-assist.

As N is fixed to 15, almost all available seeds inside requester’s own ISP should be used in each request even seeds locations are randomly distributed. Due to ISP-Biased+ neighbor selection’s design limitation, it should not be benefit from more ISP inter-connected. This is confirmed by Figure 5(a). Traffic patterns are irrelevant to inter-connected ISP numbers in the city.

On the contrary, the more ISP in the city, the better performance of BitTorrent traffic in CDN-assist algorithm. As shown in Figure 5(b), local traffic increases rapidly and backbone traffic decreases at the same rate when ISP number increases. It is also not strange that response delay of our approach outperforms ISP-Biased neighbor selection significant in Figure 3(c).

V. IMPLEMENTATION AND DISCUSSION

A first major issue is that how we can collect these geographic locality information? and by whom? and how to preserve the privacy issues? As the role of neighbor selection is assigned to Tracker servers, they could maintain a common locality information database collected from ISPs, CDN providers, etc. Users should have the right of choice: provide their topology information(e.g. IP addresses) to gain preferable performance, or preserve privacy at the cost of performance. Moreover, the solution favors interactions between the closest peers. In practical, a fixed share of neighbors should be deliberately chosen from different categories. This feature aims to guarantee the dispersion/distribution performance of the BitTorrent network.




Fig. 5. (a)ISP-Biased+ neighbor selection (b)CDN-assist neighbor selection and routing (c)Response Delay Comparison

Upgrade tracker and client software to support ISP-Biased+ and CDN-assist neighbor selection is simple. But the implementation of CDN-assist also requires efforts of both ISPs and CDN providers. As CDN-assist approach favor city-level cross-ISP traffic by carrying them inside CDN, edge routers should modify their route tables to support P2P packets redirect. Sited with edge routers, we believe already existing P2P traffic shaping devices is capable of this work. Compared with its potential benefits, these costs are trivial.

People may question why commercial CDN providers should take part in this service. Recall that one primary purpose of CDN is to reduce repeated network transmissions from content source to network edge. Commonly, only companies or organizations that could afford the services have the privilege of using CDNs to distribute their content. As stated above, P2P applications have occupied 70% of the traffic in Internet. There are many examples in the area of peer-to-peer (P2P) applications, such as YouTube [8] video sharing service, which enable individual consumers to act as content providers. Opening up a CDN so that P2P peers can take advantage of the CDN's infrastructure will bring efficient and high quality content distribution to users and reduce network cost to ISPs. All these benefits in turn open a big chance for CDN itself.

VI. CONCLUSIONS AND FUTURE WORKS

Current implementations of BitTorrent-like P2P applications ignore the underlying Internet topology hence incur a large amount of costly traffic both inside ISPs' national backbone networks and over cross-ISP internetworking links. These traffics also increase user perceived response latency. ISP-Biased Neighbor Selection proposes to exploit peers' topological locality by biased neighbor selection, in which a peer chooses the majority of its neighbors from peers within the same ISP. In this paper, we further exploit peers' geographic locality. First we improved ISP-Biased Neighbor Selection to take access network locality city locations within the same ISP into selection consideration. When required neighbor number is relatively much less than seeds available, ISP-Biased Neighbor Selection+ performs much better than original approach. Next, we propose that a peer could also choose its neighbors from peers of different ISPs within the same city with priority: assisted by a next generation CDN's unique ISP-internetworking content distribution network, these

local cross-ISP traffics can be routed through local CDN cite. Our experimental results show that cross-ISP traffic burden can be completely shifted to CDN local links and backbone traffic. In addition, user perceived delay can be significantly reduced. We plan to develop prototypes of ISP-Biased+ and CDN-assist algorithms and evaluate their performance in the real CDN platform we have mentioned above.

VII. ACKNOWLEDGMENT

The project is supported by the National Natural Science Foundation of China (No.60572063) and Chinese National High-Tech Research and Development Plan (No. 2007AA01Z223).

REFERENCES

- [1] CacheLogic. Cachelogic - advanced solutions for peer-to-peer networks.
- [2] Ashwin Bharambe, Cormac Herley, and Venkata N.Padmanabhan. Understanding and deconstructing bittorrent performance[A]. In Proceedings of 2005 SIGMETRICS[C], 2005.
- [3] William B. Norton. The evolution of the u.s. internet peering system, 2003.
- [4] Ruchir Bindal, Pei Cao, William Chan, Jan Medval, George Suwala, Tony Bates, Amy Zhang, Improving Traffic Locality in BitTorrent via Biased Neighbor Selection, IEEE International Conference on Distributed Computing Systems (ICDCS), July 2006
- [5] Thomas Karagiannis, Pablo Rodriguez, and Konstantina Papagiannaki. Should internet service providers fear peer-assisted content distribution. In Proceedings of the Internet Measurement Conference 2005, 2005.
- [6] Akamai Technologies, Akamai Web Page, URL:<http://www.akamai.com>, 2006.
- [7] Beijing Blue I.T. Technologies Co., Ltd. (ChinaCache), ChinaCache Web Page, URL:<http://www.chinacache.com>
- [8] YouTube Inc, YouTube Web Page, URL:<http://www.youtube.com>, 2006
- [9] Bram Cohen. Incentives build robustness in bittorrent, 2003.
- [10] Gkantsidis, C.; Rodriguez, P.R.; Network coding for large scale content distribution INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE Volume 4, 13-17 March 2005 Page(s):2235 - 2245 vol. 4
- [11] Mea Wang; Baochun Li; How Practical is Network Coding? Quality of Service, 2006. IWQoS 2006. 14th IEEE International Workshop on
- [12] Guo, L.; Chen, S.; Xiao, Z.; Tan, E.; Ding, X.; Zhang, X.; A performance study of BitTorrent-like peer-to-peer systems, Selected Areas in Communications, IEEE Journal on Volume 25, Issue 1, Jan. 2007 Page(s):155 C 169
- [13] Wang, J.H.; Dah Ming Chiu; Lui, J.C.S.; Modeling the Peering and Routing Tussle between ISPs and P2P Applications. Quality of Service, 2006. IWQoS 2006. 14th IEEE International Workshop on
- [14] Cho, Kenjiro (IJ); Fukuda, Kensuke; Esaki, Hiroshi; Kato, Akira; The impact and implications of the growth in residential user-to-user traffic. Source: Computer Communication Review, v 36, n 4, October, 2006, p 207-218
- [15] http://www.qwest.com/legal/docs/Qwest_Internet_Network_SLA_102103.pdf