Analyzing Peer-to-Peer Traffic's Impact on Large Scale Networks*

Mao Yang, Yafei Dai, and Jing Tian

Peking University, Beijing, P.R. China
{ym, dyf, tj}@net.pku.edu.cn

Abstract. Currently, peer-to-peer file sharing systems are playing a dominant role in the content distribution over Internet. Therefore understanding the impact of peer-to-peer traffic on large scale networks is significant and instrumental for the design of new systems. In this paper, we focus on Maze, one of most popular P2P systems on CERNET. We perform a systematic characterization of Maze's [1] traffic impact on CERNET. We investigate the traffic volume and bandwidth on different spatial levels aggregation. According to our log-based analysis, we claim that current P2P systems have much room to improve in reducing backbone network consumption. Locality-aware content delivering mechanism can reduce the traffic on backbone network effectively. Moreover, a system with less free-rider[3] will further reduce the traffic consumption. Thus the designers of P2P system should pay more attention on incentive mechanism to reduce free-rider.

1 Introduction

Currently, peer-to-peer file sharing systems are playing a dominant role in the content distribution over Internet. The dramatically increasing traffic make the ISPs worry about the abuse of backbone bandwidth by P2P systems. In this paper, we analyze P2P traffic based on public traffic log dataset of Maze. Maze[2] has been one of the largest non-commerce P2P file sharing system over CERNET (China Education and Research Network), which is developed, deployed and operated by our academic research team. Based on its open log dataset, we can leverage Maze as a large-scale measurement platform. CERNET is an ISP which connects thousands of universities and research institutes throughout China. In CERNET, a university can be regarded as an intranet with high bandwidth. We named these intranets *zones* and the *zones* are linked by the backbone network of CERNET. There are more than 200,000 active MAZE users on CERNET every month, thus we think Maze can be an ideal platform for our measurement and analysis.

The goal of our work is to help people to understand the characteristics of P2P traffic across large scales networks and its impact on Internet backbone network. Besides, we also want to find some mechanisms that can save the backbone bandwidth utilization. The followings are of the interesting questions that we want to research and understand:

* Supported by National Grand Fundamental Research 973 program of China under Grant No.2004CB318204; This work is partially supported by an Intel Sponsored Research Project.

V.N. Alexandrov et al. (Eds.): ICCS 2006, Part IV, LNCS 3994, pp. 412 – 419, 2006.

a) How do the P2P users distribute across internet on different spatial (User, IP, and zone) levels? b) What is the characterization of P2P traffic volume and traffic bandwidth? c) Do current P2P applications waste too much backbone bandwidth?

Different from previous work, we adopt a new methodology in our research. We aggregate users by zones, because zone is a more suitable level than IP-prefix or AS in analyzing the traffic impact on ISPs' backbone network. To our best knowledge, our study is the first research work base on zone level. Further more, we have most detailed logs on file transfer transactions, which enables us to perform a systematic measurement and an accurate simulation.

Based on our analysis results, we learn the following lessons:

1. The traffic volume distributions are different on different level aggregation.

2. The intra-zone bandwidth is stable and high while the inter-zone bandwidth is unstable and low. The average bandwidth decreases when the traffic is heavy.

3. The current P2P systems' content delivery model wastes much backbone bandwidth. Then we discuss potential improvement in saving backbone bandwidth of three mechanisms and claim that the P2P system is a good application model for ISPs if the systems can reduce the number of free-rider and adopt some local-aware content delivery mechanism.

The paper is structured as follows. First, we describe our research methodology in Section 2. We then discuss the host and traffic distribution in Section 3. In Section 4, we analyze the traffic impact on backbone network. Section 5 is the related works. Finally, we conclude in Section 6.

2 Methodology

Though continuous logs of Maze traffic are maintained, we perform our analysis on a log segment gathered during the period of three weeks from 09/09/05 to 09/30/05. During this period, more than 190,000 active users participated in more than 26 million file transfers. The total data traffic volume exceeded 460 Terabytes.

Table 1. Data set of Maze Traffic

Log during	9/9/2005 - 9/30/2005
# of records	26,615,75
# of unique users	190,645
# of unique IP	369,724
Total traffic volume (GBytes)	460,000
Average flows / Second	253MBytes

The data gathered for this study consists of a collection of user points during this period and the detailed traffic log. When two peers report the completion of a file transfer to the server, our log keeps only the data from the uploading peer. Each traffic log entry contains the following: *uploading peer-id*, *downloading peer-id*, *log upload time (server)*, *transfer start time (source)*, *transfer end time (source)*, *bytes* *transferred*, *file size*, *download peer's IP*, *upload peer's IP*, and *file md5 hash*. The bytes transferred can be different from the file size if the transfer was interrupted, or if the transfer is sourcing from multiple peers.

2.1 Map the IP Address to Locations

The 164,056 (86%) users and 152,136 IP addresses in Maze come from CERNET and this paper only analysis the users on CERNET. We aggregate IP addresses into Zones by using the WHOIS service of CERNET. As we have mentioned before, a zone refers to an intranet of a university / college or a research institution etc, and so the intra-zone transfers always have high bandwidth and do not consume any backbone bandwidth. The addresses space of CERNET currently spans 2752 zones. Most zones only own less than 32*256 IP address. There are only three zones own a whole B-Class IP addresses space.

3 Host and Traffic Distribution

3.1 Host Distribution

Figures 1 plots the cumulative distributions of users and IPs associated with Users ranked in decreasing order of number. There are 170000 unique users span on 846 different zones. We observe skews in the distributions of users after the zone aggregation. 52% users are in the top 20 zones. The same thing discovered in the users' IP addresses, 50% unique IP addresses are in the top 20 zones. The Figure 2 illustrates the number of unique IP address verses the number of unique users in each zones. The average the host density (# of user / # of IP) is 1.07. We observe that are especial large which means the users on these zones are using NATs.

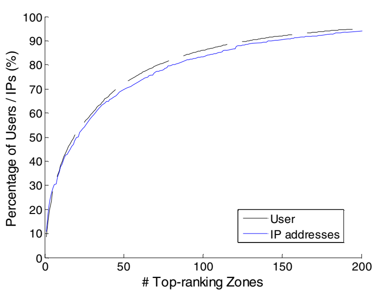

Fig. 1. The cumulative distributions of users and IPs associated with Users ranked in decreasing order of number

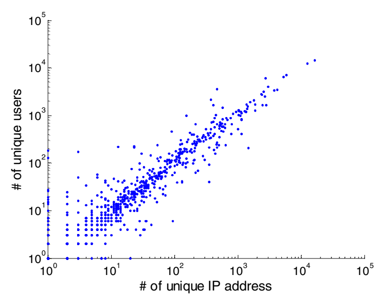

Fig. 2. The number of unique IP address vs. the number of unique users

3.2 Traffic Volume Distribution

To understand the impact of P2P systems on backbone of CERNET, we should analyze the distribution of the traffic volume. Because of the distribution of IP addresses is similar to user distribution, we analyze the traffic volume distribution only from user layer and zone layer. There is 36.20% intra-zone traffic volume and the other traffic is among zones which consumes the backbone bandwidth.

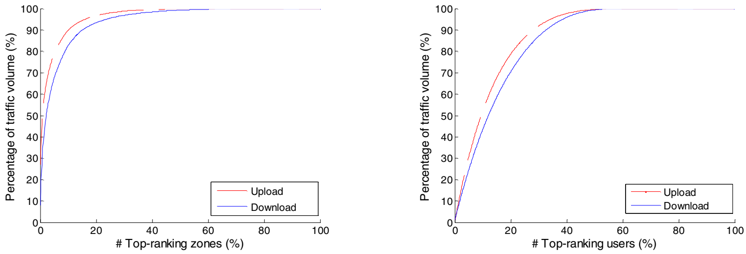

Fig. 3. The cumulative distributions of inter-zone traffic volume associated with Users / Zone level ranked in decreasing order of volume. Left: aggregate on user, Right: aggregate on zone.

Figures 3 illustrates inter-zone traffic volume aggregated on user level. Top 50% users are responsible for the whole download and upload traffic volume, and the top 20% users are responsible for over 50% traffic volume. We observe the distribution of upload volume is more skewed than the download volume, which means there are more super nodes acting as server in the system. Neither distribution of upload nor download volume follows the Zipf's law, which is not a straight line in log-log scales.

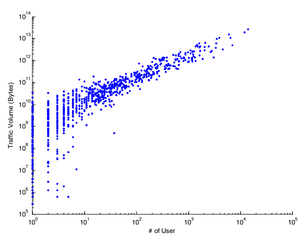

Fig. 4. Number of user vs. the total intra-zone traffic volume for each zone

What surprises us is that the distribution is quit different from user level. There are fewer zones (30%) providing the upload and more zones (60%) are responsible for download. As we investigate, this problem is cause by the limitation of IP addresses in CERNET. The users which are behind NATs or firewalls can hardly be accessed by users from other zones. Many zones in CERNET have firewalls because of the limited IP addresses space, so the users from these zones cannot not serve the user in other zones.

Figure 4 present the number of user versus the total intra-zone traffic volume for every zone. The zone which has more users induces more intra zone traffic. This is reasonable: the current P2P file sharing systems including Maze adopt a bandwidth-first mechanism to encourage user to download from proximity in a higher priority.

3.3 Bandwidth Characteristics

This subsection discusses the bandwidth characteristics of P2P systems (both on user layer and zone layer). We define the average bandwidth between a pair user as: If there were n files transferred between two peers and the file size is S_i , the transfer time is D_i. We define the average bandwidth (or transfer speed) between two peers: $AverageBandWidth = sum(S_i) / sum(D_i)$. This metric help us to understanding the traffic quality between each peer.

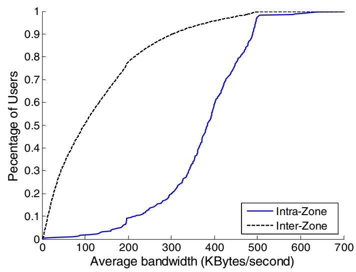

Fig. 5. The cumulative distributions of average bandwidth

Figure 5 plots the cumulative distributions of average bandwidth of intra-zone links and inter-zones links. In current P2P file sharing systems, the uploading peers always limit uploading bandwidth to the downloading peers. In Maze, the free-rider's downloading bandwidth will be limited to less than 200KBps by uploading peers, and the default max bandwidth for every link is 500KBps. Thus, the intra-zone average bandwidth ranges mainly from 200Kbps to 500kBps. We observe there is only 10% bandwidth less than 200KBps for intra-zone links, and the percentage increases to more than 70% for inter-zone links. The transfer intra-zone can provide more high speed service, and the P2P systems might enough peers exchange their data intra-zone.

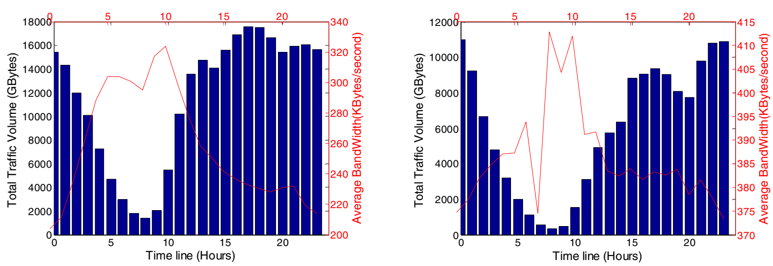

Fig. 6. The total traffic volume and average bandwidth aggregate by hours. Left: Intra-zone right: inter-zone.

We observe the traffic volume has a strong daily pattern (Figure 6.). The max flux was observed on 3 pm and 11 pm. The average bandwidth has strong correlations with the traffic volume, especial for the inter-zone traffic. When the network workload is heavy, the bandwidth for P2P decreases 10% for intra-link and 30% for inter-link, which means the bandwidth intra-zone is more stable.

4 The Traffic Impact on Backbone Network

In this section, we focus on the traffic impact on backbone network. A large scale P2P file sharing system consists of millions of users in an overlay network. Users exchange their content to each other in the overlay network. If two users delivery their content between different zones, it will consume the backbone bandwidth of ISPs.

The current P2P systems such as BitTorrent and Maze support some inner mechanisms to reduce the abuse of backbone network. The basic mechanism is bandwidth-first mechanism. It means when a peer are upon a selection of potential uploading peers, it try to select the peers who have higher bandwidth links to it. It is an approximate mechanism to implement locality-aware mechanism which lets user adapts file downloading to match the physical networks. As we know, there are more than 60% traffic is from external peers in Maze. Are there any potential in saving backbone bandwidth?

We propose some mechanisms which can reduce traffic on backbone as follow:

Origin mechanism (used by Maze): When peer Alice requests a file, the central index server will tell this peer some peers who have this file. Alice will request this file from those peers and download with bandwidth-first mechanism.

A) Locality-aware mechanism: If there are some online peer has this file in the same zone with Alice, Alice prefers downloading this file from those local peers.

B) Locality-aware on No free-rider condition: free-rider refers to the peer does not upload any file to other peer even if he has some content. If peer Bob downloaded a file, we assume he wills storage this file more than one month and can service this file to other peers whenever he is online. These assumptions will help us to understand the impact of free-rider to P2P systems.

C) Perfect proxy mechanism: We assume there are perfect proxies in every zone. When peer Bob download a file outside the zone, the proxy of this zone will cache this file forever. Alice need not download file outside of zone again if the proxy has it. This mechanism is like the traditional CDN (Content Distribution Network) solutions.

To understanding the potential of those proposed mechanisms, we conduct a trace-driven simulation. The steps of our simulation are: a) Sort the whole transfer record logs based on transfer start time. b) Parse the transfer record, when peer Alice downloads File f from peer Bob from time $t1$ to time $t2$, we assume Peer Alice has owned file f in time $t2$ and Bob has owned file f in time $t1$ c) Read the transfer records one by one, when peer Alice requests a file f at time t, perform the three optimization mechanisms, calculate the hit rate (by volume) in local zone. The detail as follow:

A) If there are some other peers is in the same zone with Alice and they are uploading f at time t, Alice apt to download only from those peers. We called locality-ware mechanism.

B) If there are some other peers is in the same zone with Alice and they have owned f before time t, Alice apt to download only from those peers.

C) If the proxy of Alice's zone has owned f before time t, Alice apt to download only from the proxy.

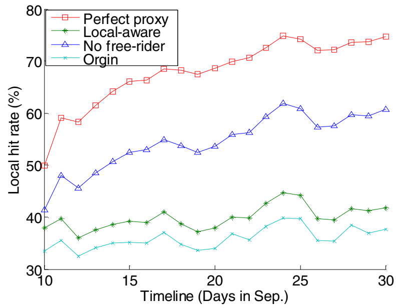

Fig. 7. The percentage hit rate in local zone in timeline (days)

Figure 7. plot the average percentage hit rate in local zone, the hit rate means the percentage of traffic volume just on local zone. The original mechanism has a stable hit rate around 36.20%, which means that more than 60% traffic occurs on of the backbone network. The locality-aware mechanism will increase the hit rate by 5% percentage. It also has a stable hit rate.

If there are no free-riders in the system, the hit rate will have a sharp increase. The average hit rate is 54.80%, and the hit rate will increase following the time from 40% to above 60%. This demonstrates if all the users maintain the download files more than 3 weeks, the hit rate will exceed 60%. Unfortunately, all of P2P file sharing systems have the free-rider problem. Many users remove their downloaded file from system and refuse to service other users. Our experiment proves that the potential improvement of P2P system.

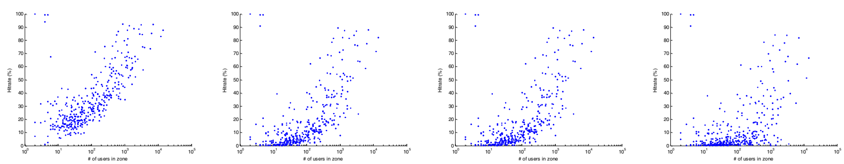

Fig. 8. The hit rate for every zone (*x*-axis is the number of user in a zone). Left to right (Original, Locality-aware, No free-rider condition, Perfect proxy).

The perfect proxy mechanism is an ideal model, and it shows the upper bound of improvement. The average hit rate is 68.19%, and the hit rate reaching 75% at the end of simulation.

Figure 8. illustrates the hit rate in different zones. We find that the hit rate is influenced by the zone size. The zones with large population have higher hit rate. Some large zones have hit rate more than 50% even in original mechanism. This demonstrates a large population zones are friendly to ISPs, they do not need too much backbone network bandwidth. The zone with low population can hardly get a high hit rate even in perfect CDN model. We also observe the high population zone has the similar hit rate in no free-rider condition with the perfect proxy. This demonstrates if there are enough users in a zone, there is high probability that somebody has your desired.

We conclude several optimizations from this simulation can be done by the designers of P2P system: a) Design a better locality-aware download mechanism. Tell peers the replica of file in his / her local zone, the bandwidth-first mechanism is not accurate. b) Design some incentive systems to reduce the free-riders in the system, which can encourage users not to remove the downloaded file away from system and encourage users to stay online longer. c) Encourage new users to join in the network, or place some super nodes in medium or small size zones. The P2P system in a large scale zone will consume less percentage backbone bandwidth than ISPs supposed.

5 Related Works

Subhabrata et al. [4] analyze the P2P system traffic in (IP, prefix, AS) level, and focus on the workload model on flew-level data. Several measurement studies have characterized the basic traffuc of P2P. Saroiu et al. [5] analyzed the behaviors of peers inside the Gnutella and Napster. Krishna et at. [6] demonstrate that KaZaA traffic did not exhibit Zipf-like behavior. Thomas et al. [7] is most similar work with us. They aggregate BitTorrent users on AS level, and demonstrated the "locality-aware" solution will reduce the bandwidth usage between ISPs.

6 Conclusion

Through the analysis of Maze's log data, we achieved a comprehensive understanding of the characterizations of P2P system's traffic volume and bandwidth. We also analyze the P2P traffic impact on the backbone network. We conclude that the current P2P systems consume too much backbone bandwidth, but the situation can be improved. And ultimately P2P system is a good solution for content delivery.

References

1. http://maze.pku.edu.cn.
2. Mao Yang, Ben Y. Zhao, et al. Deployment of a large scale peer-to-peer social network, Proceedings of the 1st Workshop on Real, Large Distributed Systems.
3. M. Yang, Z. Zhang, X. Li, Y. Dai, "An Empirical Study of Free-Riding Behavior in the Maze P2P File-Sharing System," In *Proceedings of IPTPS*, Ithaca, NY. February 2005.
4. S. Sen and J. Wang. Analyzing peer-to-peer traffic across large networks. In Proceedings of the Second SIGCOMM Internet Measurement Workshop (IMW 2002), Marseille, France, November 2002.
5. S. Saroiu, P. K. Gummadi, and S. D. Gribble. A measurement study of peer-to-peer file sharing systems. In Proceedings of Multimedia Computing and Networking (MMCN) 2002, January 2002.
6. Krishna P. Gummadi, Richard J. Dunn and et al. "Measurement, Modeling, and Analysis of a Peer-to-Peer File-Sharing Workload". Proceedings of the 19th ACM Symposium on Operating Systems Principles (SOSP-19), Bolton Landing, NY.
7. Thomas K., Pablo R., et al. Should Internet Service Providers Fear Peer-Assisted Content Distribution? Internet Measurement conference 2005.