On the Locality of BitTorrent-based Video File Swarming

Haiyang Wang* Jiangchuan Liu* Ke Xu†
*{hwa17, jcliu}@cs.sfu.ca
School of Computing Science, Simon Fraser University, British Columbia, Canada
†xuke@csnet1.cs.tsinghua.edu.cn
Department of Computer Science and Technology, Tsinghua University, Beijing, China

***Abstract*— In the past few years, there have been tremendous interest in the peer-to-peer(P2P) content delivery. Although this communication paradigm does not need a dedicated server infrastructure, it dramatically increases the traffic over inter-ISP links. In particular, the most popular P2P application, BitTorrent(BT) generates a huge amount of traffic on the Internet.**

To address this challenge, P2P locality has been examined, which explores the access to local resources to optimize the inter-ISP traffic. However, most of these approaches have focused on a global strategy, and attempted to change the peer selection mechanism, which potentially affects the random topology of BT and thus reduces its robustness. The content and the peer diversities are seldom discussed, particularly the video file swarms of distinct characteristics.

In this paper, we for the first time examine the different BT contents and peer properties in regards to the locality issues through a large-scale measurement. We demonstrate the distinct characteristics of video file swarms, and find that the distribution of the AS clusters (a set of peers belonging to the same AS) follows the Mandelbrot-zipf law. Our results also suggest that the peer in a few ASes are more likely to form large AS clusters and most ASes on the Internet do not have enough potential for locality. Therefore, a global locality approach may not be our best choice. We then address the problem through a selective locality approach based on a novel peer prediction method.

I. Introduction

Peer-to-peer content delivery has become one of the most popular applications in recent years. BitTorrent, the most successful P2P file sharing system over the Internet, has been widely used for the distribution of large files. Although the P2P paradigm does not have to maintain a dedicated server infrastructure, it generates a huge amount of traffic over inter-ISP links. In particular, even though some BT peers are located in the same or nearby ISPs and downloading identical contents, they are unnecessarily connected through remote peers. Since the ISPs typically pay their peering or higher-level ISPs for global connectivity, the traffic between different ISPs is costly and presents significant network engineering challenges. To make the matter worse, the success of BitTorrent has also greatly motivated the design of new traffic-intensive applications, such as streaming service, over the Internet [1]. In fact, BitTorrent itself has already been extensively used for video file distribution, albeit in a download-and-play mode.

To alleviate the inter-ISP traffic problem, many solutions have been proposed beyond the straightforward blocking of the P2P traffic [2]. Among them, P2P locality [3] has been widely suggested, which explores access localities to reduce long-haul traffic. There is no doubt about the existence and benefit of BT locality and there have been a series of implementations. However, most of them are based on the modification of the global peer selection mechanism; the peer and content features are largely ignored during the locality process. Although the peers may gain some benefits from this modification, the robustness of the BT networks may be sacrificed due to the change of random topology. On the other hand, the distinct characteristics of different contents, particularly video file swarms, have yet to be measured and explored.

In this paper, we for the first time examine the BT locality problem with content and the peer diversities through a large-scale measurement. Our study suggests that the video contents are obviously quite popular in BT networks. Most of these video file swarms contain very large files that pose significant challenges to ISPs. On the other hand, based on the AS level measurements, we find an interesting relationship between different ASes by holding the peers that willing to download the identical contents. Our investigation indicates that the peers belonging to some ASes are more likely to appear in a large AS cluster. We further observe a Mandelbrot-zipf distribution [4] in the ratio of AS cluster size to *swarm size* (the total number of peers in a swarm), which indicates that the peers belonging to a few large AS clusters are indeed more eligible to be adjusted by a locality mechanism. Therefore, a selective locality mechanism is required to optimize the overhead and the robustness of the BT networks.

The rest of this paper is organized as follows: In Section II, we illustrate the related works. Section III presents our AS level measurement results. After the description of the main challenge, we discuss an AS relationship based approach in section IV to predict whether a peer is likely to appear in a large AS cluster. Finally, the paper is concluded in Section V.

II. Background and Related Work

P2P locality attracted attention from many researchers in recent years. The pioneer work of T. Karagiannis et al. [3] is the first study to address the locality issues in P2P systems. Aiming to solve the inter-ISP traffic problem, they studied both the real traces and simulation results. They also evaluated the benefit of several architectures and present the concept of locality in a particular solution. Blond et al. [5] showed, through a controlled environment, that high locality values (defined by [3]) enable up to two order of magnitude saving on inter-ISP traffic without any significant impact on peers' download completion time. The work from Xie et al. [6] suggested cooperation between peer-to-peer applications and ISPs by a new locality architecture, P4P. Large-scale test results showed that P4P can reduce both the external traffic and the average downloading time. Choffnes et al. [7] proposed Ono, a BitTorrent extension that leverages a CDN infrastructure, which can find the location of peers that are close to each other. Bindal et al. [8] also examined a novel approach to enhance BitTorrent traffic locality: *biased neighbor selection*. Using this method, a peer chooses the majority, but not all, of its neighbors from peers within the same ISP. Simulation results showed that it can greatly reduce the inter-ISP traffic of BT networks.

However, most of these pervious studies have focused on global strategies. The content and the peer diversities are seldom discussed. For example, the BT locality approaches are processed upon every single peer in the BT swarms, and potentially changed the random topology of the BT peers [9]. These modifications will not only raise a remarkable overhead but also affect the robustness of BT networks.

III. Measurement and Analysis of AS-level Characteristics

In this section, we for the first time examine the BT swarms of different contents in regards to locality. In our study, we investigate 30415 video metainfo files and 44317 non-video metainfo files. These metainfo files are mainly advertised by www.btmon.com from Feb 12 2007 to Aug 12 2008. We developed a script to automatically detect the 'href' field in a given HTML file and download the metainfo files ending with '.torrent'.

Within our data set, there are 316 bad metainfo files, 1027 unavailable swarms due to tracker failures, and 3340 swarms having less than 2 peers. None of these abnormal swarms are included in our study.

We carry out an Internet-based measurement using the PlanetLab [10]. We run a modified version of CTorrent [11] (a very typical BitTorrent client in FreeBSD) on more than 200 PlanetLab nodes. This client software was modified to log various peer level information including IP addresses. The modified CTorrent clients actively join each torrent and record the peers' IP within the peer set. Since the contents of many Internet swarms may involve copyrights, no real content were downloaded in our measurement. Moreover, a postprocess is applied to filter the peer information of probing nodes in the raw data.

Content size is a very important characteristic in all P2P systems. Figure 1 shows the distribution of content size among different data sets. We can see that the contents shared by BT video file swarms are mostly large. In video file swarms, the mean object size is approximately 1000MB and 90% of video

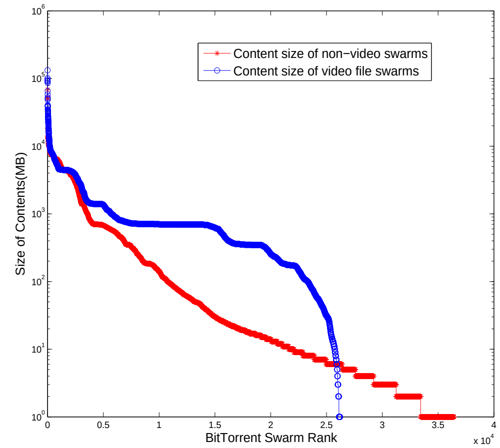

Fig. 1. Content size of BitTorrent swarms (sorted in descending order)

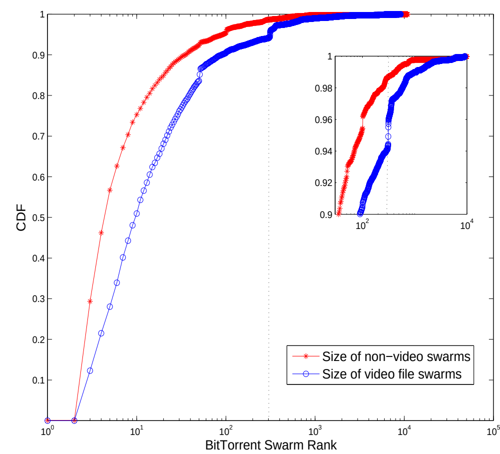

Fig. 2. BitTorrent swarms size (CDF)

contents are larger than 100MB. Moreover, there are 5% of the video contents with size being larger than 10GB, and the maximum video size reaches nearly 20GB. On the other hand, the size of non-video swarms is relatively smaller, with only 30% of the non-video contents being larger than 100MB. It also worth noting that over 50% of non-video contents are less than 20MB, whereas those small contents are very few in the existing video file swarms.

Figure 2 shows the cumulative distribution of the BT swarm size. This distribution is relevant to the popularity of different BT contents. We can see that although more than 95% swarms have less than 300 peers, the video file swarms are generally larger than non-video swarms.

According to these observations, we know that the video file swarms potentially generate more traffic due to its large content size and swarm size. If the peers of a video file swarm are uniformly distributed between different ASes, it is more likely to generate heavy traffic through the inter-ISP links .

To further investigate this problem, we randomly select 8893

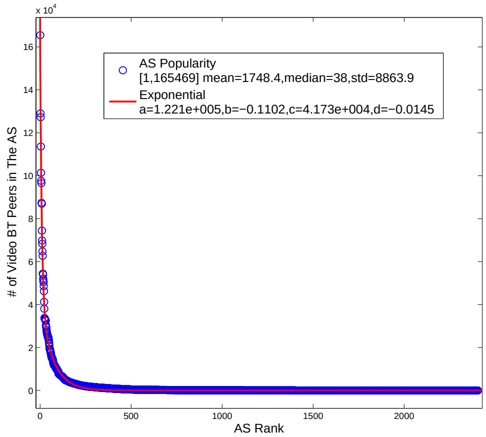

Fig. 3. AS popularity of existing video file swarms

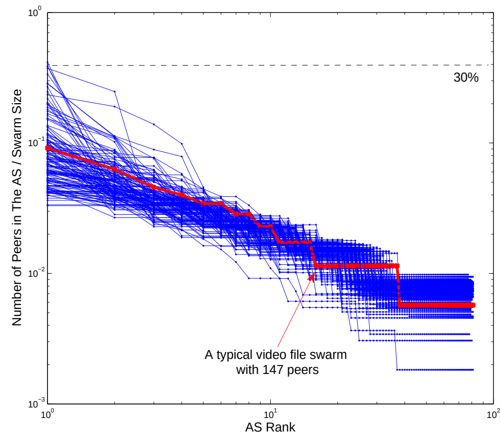

Fig. 5. Ratio between AS cluster size and swarm size (141 small swarms)

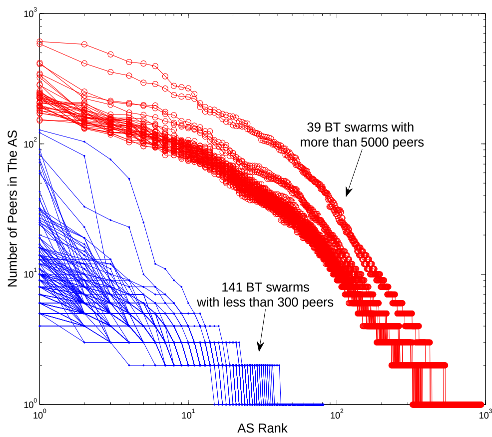

Fig. 4. Distribution of AS cluster size

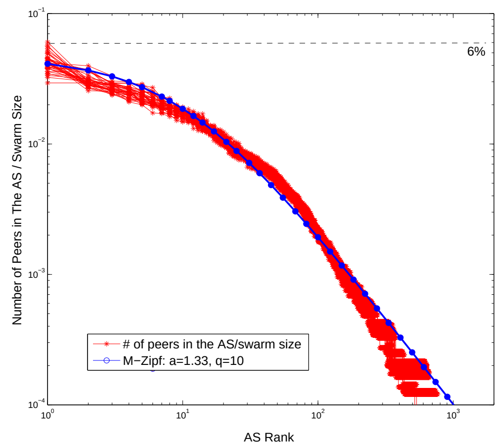

Fig. 6. Ratio between AS cluster size and swarm size (39 large swarms)

BT video file swarms, and collect the AS information of every peer in each swarm. This probing is based on the 'whois' command on the Linux system, and most replies are from 'whois.cymru.com'. From Figure 3, the AS popularity of video BT peers fits an exponential distribution; that is, among all the 2405 ASes in our measurement, most of them have less than 10000 peers in total. We also present the Top-10 ISPs/ASes with most video BT peers in Table I. These results give us further hints on the challenge and the potential requirements of P2P locality in these ASes.

We then investigate the AS distribution of different video file swarms in Figure 4. In this figure, 141 small video file swarms (with less than 300 peers) and 39 large video file swarms (with more than 5000 peers) are selected. Each point in the figure indicates the number of peers in an AS, and the values are sorted in descending order. We can see that the large BT swarms generally involve more ASes and their AS distributions are more uniform than that of small ones.

Figures 5 and 6 show the ratio between the AS cluster size and the swarm size. We can see that this ratio is quite high in small swarms: the largest AS cluster can even reach 30% of the swarm size. Therefore, given their small peer populations, these swarms already have strong locality features in nature. Consequently, the extra locality mechanism is not necessary for them.

In the case of large swarms, Figure 6 shows that although large AS clusters are more likely existing, the ratio to the swarm size is relatively low. In fact, the largest AS cluster only has less than 6% of peers in the AS. Moreover, we find that the distribution of this ratio can be fitted by a Mandelbrot-Zipf (MZipf) distribution with $\alpha = 1.33$ and $q = 10$. The MZipf distribution defines the probability of accessing an object at rank i out of N available objects as: $p(i) = K/(i+q)^{\alpha}$, where $K = \Sigma_{i=1}^{N} 1/(i+q)^{\alpha}$, α is a skewness factor, and $q \geq 0$ is a plateau factor. q is so called because it is the reason behind the plateau shape near to the left part of the distribution. This is intuitive because the size of AS is an upper bound on the AS cluster size. Moreover, the Zipf-like distribution indicates that, the size of most AS clusters are relatively small.

According to the definition of locality [3], although there

TABLE I
TOP 10 ISPS (BT VIDEO USER)

	AS#	Peers	AS Name-Internet Service Provider
1	3352	165469	TELEFONICA-DATA-ESPANA(TDE)
2	3662	129047	DNEO-OSP7-COMCAST CABLE
3	6461	127297	MFNX MFN-METROMEIDA FIBER
4	2119	113597	TELENOR-NEXTEL T.NET
5	19262	101390	VZGNI-TRANSIT-Verizon ISP
6	3301	97658	TELIANET-SWEDEN TELIANET
7	3462	96564	HINET-DATA CBG
8	4134	87392	CHINANET-BACKBONE
9	6327	86964	SHAW-SHAW COMMUNICATION
10	174	74453	COGENT COGENT/PSI

are many large AS clusters in large swarms, the locality of most peers is poor in nature. Therefore, the peers in a large AS cluster have both the potential and incentive to incorporate a locality mechanism; also, the optimization of these peers is of more interest to both ISPs and individual users. However, most existing locality approaches treat all peers in the swarm with equal importance and attempted to changed the global peer selection mechanism. We believe that the random peer selection is the core of the BitTorrent protocol. The common belief that BT is efficient, robust and scalable, is mostly based on the random topology of such a system [12][13][14]. Therefore, a global locality mechanism will not only involve more overhead but also degrade the robustness of existing BitTorrent networks [5]. Specifically, if we apply locality to all peers, the peer graph will be more clustered than that of random, and therefore, few peers will have the neighbors that belonging to the other ISPs; When the churn rate is increasing, the failure of these cross-ISP peers may lead to the swarm splitting problem that is harmful for content spreading.

On the other hand, the challenge to design a selective locality mechanism is also significant: It is well known that the locality mechanism must be processed before forming the huge swarms; During the early periods, however, it is hard to predict whether a peer will belong to a large AS cluster in the future.

Fortunately, according to our measurements, we find that the ASes are not independent with each other; they are highly related by holding different peer sets of the BT swarms. On the other hand, peers belonging to different ASes also have different features due to this relationship. In particular, some peers are more likely to form a large AS cluster than that of others. Such an relationship is potentially more useful among video file swarms because the video contents are more likely to have geographic localities due to the language variations. For example, few people in United States would like to download a video of Japanese version. This leads to the design of a prediction method for selective locality, as will be discussed particularly in the next section.

IV. DISCUSSION OF A POSSIBLE PEER PREDICTION METHOD

In this section, we will discuss the peer prediction method based on the AS level relationships. The main idea of this approach is that, based on the pre-knowledge of AS and swarm relationship, we can quantify the possible clustering characteristics of a given AS. In particular, if we know the peers belonging to some ASes are more likely to form a large AS cluster, we can apply a selective locality mechanism only at these peers. The peers belonging to other ASes, on the other hand, can be processed by the standard random peer selection to ensure the network robustness and connectivity. It is also worth noting that we assume certain stationarity of this property, which we expect to be further confirmed by the future measurement results.

We use $\aleph$ to denote all ASes in the network, and use $\Re$ to denote the set of existing video file swarms. We define two random variables A and S in our framework. A refers to different ASes, and the probability that A takes on the value a ($a \in \aleph$) is $P(A = a)$. S takes on values over the set of existing video file swarms $\Re$. We use T to refer to the frequency table of A and S. An elements in the table, $T(a, s)$, refer to the number of peers (in swarm s) that belong to AS a.

Two relationships can be built according to table T. The first is the conditional probability distribution $P(S|a)$, which represents that for a given AS a, the frequency of swarm S is belonging to a. This value can be computed by electing the column in the table T corresponding to a, and normalizing it by the sum of this column:

$$P(s|a) = T(s, a) / \sum_{a} T(s, a) \tag{1}$$

The second relationship is the conditional probability distribution $P(A|s)$, which represents that for a given swarm s, the likelihood of ASes A being used by a given swarm s. This value can be computed by electing the row in the table T corresponding to s, and normalizing it by the sum of this row (the computation detail is shown in Figure 7):

$$P(a|s) = T(s, a) / \sum_{s} T(s, a) \tag{2}$$

According to these two relationships, we can further compute the probability $P(A|a)$. $P(A|a)$ summarizes how AS a is associated with all other ASes A due to the swarm level relationship. By tally up how likely other ASes are also holding similar amount of peers from the same swarm, we sum over the contribution in proportion to how frequently swarm s is belonging to AS a:

$$\begin{aligned} P(A|a) &= P(A|s_1)P(s_1|a) + P(A|s_2)P(s_2|a) + ... \\ &= \sum_{s} P(A|s)P(s|a) \end{aligned} \tag{3}$$

After computing $P(A|a)$, we use entropy to quantify the amount of randomness in the probability distribution. Note that a low entropy implies that AS a is weakly associated with a large number of ASes. This occurs when the AS generally

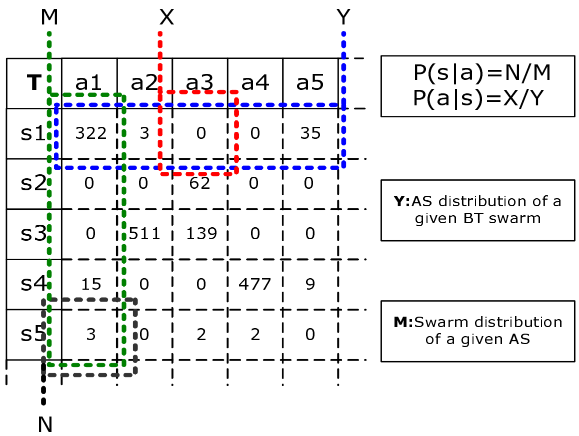

Fig. 7. Details of table T and different relationships

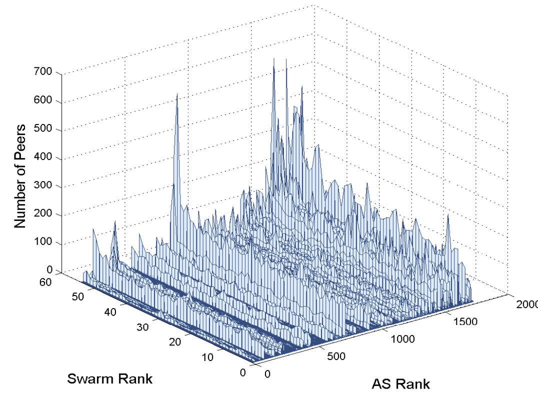

Fig. 8. Table T of 1747 ASes

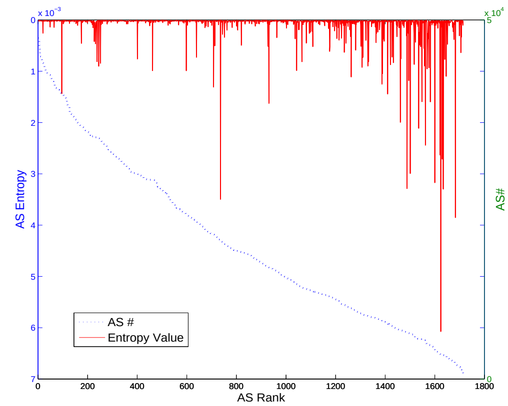

Fig. 9. The entropy value of different ASes

do not have large AS clusters. On the other hand, when the entropy value of an AS is very high, the peers belonging to this AS are very likely to form a large AS cluster. Therefore, we can compute the entropy of $P(A|a)$ as follows:

$$Entropy(a) = H(P(A|a)) = -\sum_{a' \in A} P(a'|a)logP(a'|a) \quad (4)$$

According to the entropy value of different ASes, a modified tracker protocol will carry out the following selective locality process when a BT peer is arrived (note that the entropy of each AS is preprocessed by computing table T according to Eq.1-Eq.4; these entropy values have already existed in the trackers before the execution of the following steps):

Step1: When a peer x arrives, obtain the AS# a of this peer by sending a 'whois' request;

Step2: For a given AS# a, compute the entropy of AS# a according to Eq.4;

Step3: If this value is larger than a pre-configured threshold e, send the peer set information (the sets of neighboring peers) to peer x by giving high priority to the neighbors that are in the same AS with x; otherwise, send the peer set information according to the standard random peer selection method.

We now illustrate a simple validation of the proposed method with real AS information (more evaluation results can be found in our technical report [15]). In order to compute the entropy values for different ASes, we randomly select 54 torrents which include different peers in the 1747 ASes. The frequency table (table T) of these swarms are shown in Figure 8. We can see that the swarm distribution of different ASes do have diverse features; in particular, some ASes are always holding more peers than that of others within all the swarms. This observation also confirms that there is certain stationarity in the distribution.

In order to quantify the characteristics of AS clustering, we compute the entropy values of these 1747 ASes and plot the results in Figures 9 and 10. According to these figures, most ASes have the entropy values smaller than 0.00004, and only a few ASes have very high entropy values. This result confirms our observation that only a few ASes are eligible to be adjusted by a locality mechanism (the statistical characteristics of the entropy values are presented in Table II). It is also worth noting that there is an sharp turning point in Figure 10 when the values of X-axis are around 1600. We carefully checked these 150 ASes (AS rank 1597 to 1747) that have relatively low entropy values; the conclusion is that very few peers in these ASes are likely to join the BitTorrent swarms, which may be due to two reasons: first, we only select 54 BT torrents in this experiment, and second, some ASes are indeed limiting the traffic of BitTorrent and other P2P applications.

According to the above results, we present the Top-10 ISPs/ASes with the highest entropy values in Table III. In general, the peers in these ASes are more likely to form a large AS cluster than that of others. Therefore, the selective locality should give higher priority to the peers belonging to these ASes. Although we may use more torrents to further improve these AS entropy values, these results are still quite acceptable, e.g., AS# 3352 and 2119 are both very popular ASes in Table I.

One potential problem of this approach is the requirement of global knowledge. In fact, although the trackers are holding the global peer information of most torrents, the entropy values may not be updated in real-time, because the overhead will be

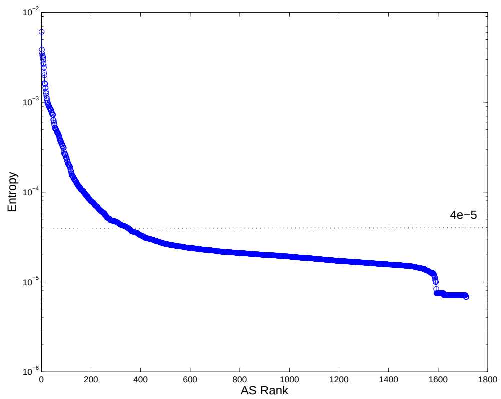

Fig. 10. Distribution of the entropy value

TABLE II
FACTS OF VIEWS, ENTROPY OF DIFFERENT ASES

E	min	max	mean	median	std
Entropy	6.82e-6	0.0061	8.27e-5	2.04e-5	3.16e-4

TABLE III
TOP 10 ISPS (ENTROPY VALUES)

	AS#	Entropy	AS Name-Internet Service Provider
1	3352	0.0061	TELEFONICA-DATA-ESPANA(TDE)
2	2119	0.0039	TELENOR-NEXTEL T.NET
3	19262	0.0035	VZGNI-TRANSIT-VERIZON ISP
4	3301	0.0033	TELIANET-SWEDEN TELIANET
5	6461	0.0033	MFNX MFN-METROMEIDA FIBER
6	4134	0.0032	CHINANET-BACKBONE
7	6327	0.0030	SHAW-SHAW COMMUNICATION
8	3320	0.0027	DTAG DEUTSCHE TELEKOM AG
9	3462	0.0026	HINET-DATA CBG
10	5089	0.0024	NTL NTL GROUP LIMITED

unacceptable during the possible flash crowds of peer arrivals. Fortunately, our observations have already shown that there is certain stationarity in the peer distribution of a given AS. Therefore, the entropy values can be computed infrequently by preprocessing the peer information. In general, the trackers only need to query the entropy value by the AS number and process the selective locality mechanism according to the querying results.

Note that our solution is beyond the simple use of AS popularity distribution. Although AS popularity distribution (Figure 3 and Table I) may provide some meaningful information for the validation, it is not feasible for the peer prediction. Specifically, the variation of peer number cannot reflect the relationship that we need to know between the ASes; for example, AS# 3662 and AS# 6461 have very similar popularity in Table I; yet the peers inside these ASes are not necessarily having similar clustering properties (AS# 3662 is quite popular in Table I but is not included in Table II).

V. CONCLUSIONS

In this paper, we studied the existing video file BT swarms in regards to the locality issues. We for the first time examined the problem through a large scale Internet-based measurement, focusing on content and peer diversities. According to our results, a global locality approach may not be our best choice. The peers in large AS clusters however are of the most importance during the locality optimization. Based on the relationships of different ASes, a possible peer prediction approach is discussed, serving as the foundation of a novel selective locality mechanism.

A distinguishing feature of our study in comparison to previous works is the focus on real-world measurement and high level features such as content and peer diversities. The different AS relationships are also quantified for the first time in the BitTorrent system. We will further enhance our solution by reducing its computation overhead and improving its accuracy for real deployment.

REFERENCES

[1] J. Liu, S. G. Rao, B. Li, and H. Zhang, “Opportunities and Challenges of Peer-to-Peer Internet Video Broadcast,” *Proceedings of the IEEE, Special Issue on Recent Advances in Distributed Multimedia Communications 96(1):11-24, 2008.*

[2] M. Dischinger, A. Mislove, A. Haeberlen, and K. P. Gummadi, “Detect Bittorrent Blocking,” in *Proc. ACM/USENIX IMC 2008.*

[3] T. Karagiannis, P. Rodriguez, and K. Papagiannaki, “Should Internet Service Providers Fear Peer-Assisted Content Distribution?” in *Proc. ACM/USENIX IMC 2005.*

[4] Z. Silagadze, “Citations and the Zipf-Mandelbrot’s law,” *Complex Systems 11:487-499, 1997.*

[5] S. L. Blond, A. Legout, and W. Dabbous, “Pushing BitTorrent Locality to the Limit,” *INRIA Tech, Rep. 2008.*

[6] H. Xie, R. Y. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz, “P4p: Provider Portal for Applications,” in *Proc. ACM SIGCOMM 2008.*

[7] D. R. Choffnes and F. E. Bustamante, “Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems,” in *Proc. ACM SIGCOMM 2008.*

[8] R. Bindal, P. Cao, W. Chan, J. Medved, G. Suwala, T. Bates, and A. Zhang, “Improving Traffic Locality in BitTorrent via Biased Neighbor Selection,” in *Proc. IEEE ICDCS 2006.*

[9] C. Dale, J. Liu, J. G. Peters, and B. Li, “Evolution and Enhancement of BitTorrent Network Topologies,” in *Proc. IEEE IWQoS 2008.*

[10] Planetlab. [Online]. Available: http://www.planet-lab.org/

[11] Ctorrent. [Online]. Available: http://ctorrent.sourceforge.net/

[12] A. AL-Hamra, A. Legout, and C. Barakat, “Understanding the Properties of BitTorrent Overlay,” *INRIA Tech, Rep. 2007.*

[13] G. Neglia, G. Reina, H. Zhang, D. Towsley, A. Venkataramani, and J. Danaher, “Availability in BitTorrent Systems,” in *Proc. IEEE INFOCOM 2007.*

[14] X. Yang and G. de Veciana, “Service Capacity of Peer to Peer Networks,” in *Proc. IEEE INFOCOM 2004.*

[15] H. Wang and J. Liu, “Statistics of BitTorrent-based Video File Swarming,” *Tech, Rep. School of Computing Science, Simon Fraser University, 2008.*