

Measurement and Analysis on P2P Swarms Based on Coupon Collector's Problem*

Xiaobin Wang, Bi Shi, Xi Chen, Kai Lin, Zhe Yang

School of Computer Science and Technology, Soochow University, Suzhou, Jiangsu Province, China
{1244517416, 407641892, 275639596 & 35417406}@qq.com, yangzhe@suda.edu.cn

Abstract - Active measurement on peer-to-peer system is the best way to understand the current situation and its evolution trends. And it also provides the real measurement basis to support the research on system simulation and optimization. Based on coupon collector's problem in random sampling, rapid swarm covering model is proposed to guide the active measurement on BitTorrent swarm, which can predict the number of measurement requests is fast convergence in the mathematical expectation. And with the optimization of lowering expectations and requesting concurrently, the request repetition rate is 67.12% off, while still covering the 95% peers in a swarm. Moreover, the measurement time can be further shortened by interacting with peers by DHT and PEX. Based on the theoretical model, a low cost measurement platform is set up to trace the 49,854 Torrentz swarms and 98 HDChina swarms for a long time. Furthermore, the geographical characteristics are also analyzed.

Index Terms - P2P swarm; active measurement; coupon collector's problem; geographical characteristics; user preference

1. Introduction

The popularity of P2P applications not only provides people with high-efficient, scalable ways to share information, but also causes problems in network planning, QoS management, as well as information security. P2P systems choose its neighbors without considering the underlying topology, and therefore generating massive of cross AS traffics. Some researchers proposed to use topology-aware or locality-aware technologies to optimize its neighbor selection algorithm [1-6]. However, their assumptions on the peer distribution model are too ideal, e.g. Blond et al assumed that peers in the BT (BitTorrent) network obey uniform distribution [5]. However, the actual P2P peer distribution is extremely skewed [7].

In this paper, we proposed a model to fast snapshot P2P swarms based on an optimized solution to the Coupon Collector's Problem, and implemented a low-cost proactive probing platform that is capable of the long-term measuring on the global BT systems. According to the data collected, we analyzed the scale of BT swarms, snapshot times, as well as geographical distribution, usage preferences, and transport performance of peers, which provides fundamental facts for simulations of BT systems.

2. Related Works

Active measurements make use of particular P2P protocols to simulate clients that join in the system to extract peer

information and network behaviors. Saroiu et al first used modified open-source client software to conduct proactive measurement. They measured the topological characteristics of the Gnutella and Napster networks [8]. Liang et al designed a spider against KaZaA, and analyzed its network behavior as well as source pollution problem [9-10].

Recently, with the popularity of BT applications which have contributed nearly 35% of the entire Internet traffic, and exceeded the total traffic of all other P2P networks. As such, measuring BT systems, and analyze its network behavior has become a research focus. Izal et al built Tracker snapshots using spiders and analyzed peers download time in the swarm, and evaluated its performance of the communication algorithm [11]. Pouwelse et al built snapshots of multiple trackers using spiders, and analyzed the popularity, reliability, and life cycles of swarms, as well as their pollution rate, and the overall download performance [12].

Later then, many researchers started to use public experiment platform or built their own environment to conduct large-scale measurement to BT systems in real network environment. Wang et al deployed 200 nodes on the PlanetLab [13]. Hobfeld et al deployed 219 and 153 nodes on PlanetLab and G-Lab respectively [14]. Otto et al used various machines in the Ono and NEWS project [1], and Zhang et al used 51 computers in their experiment environment. They ran modified BT clients on these nodes, and used massive of torrents to continuously communicate with Trackers to get IP and port information of other peers. Liu et al used 8893 torrent files [7], Wang et al used 70, 000 video torrents [13], Zhang et al used 1, 192, 303 torrents of English resources [15]. Results showed that the distribution of nodes in ASes is extremely skewed even for popular swarms. They are distributed among more than 150 ASes, and some only have 1 or 2 nodes, which is not beneficial for applying locality-aware selection algorithm.

3. Active measurement model

A. Coupon Collectors' Problem

It is possible to extract peers, and build snapshots from BT Swarms from Tracker, DHT, and PEX. Limited by the protocol, one can only get part of peer information per request, which may contain duplicated items. Therefore, this is actually a simple random sampling question with repetition, which is a classical Coupon Collectors' Problem [6]. A Coupon

* This work is supported by NSFC (61070170), the High School NSF of Jiangsu (11KJB520017), the Technical Project of Suzhou (SYG201238), and the College Student Innovation Project (111028525)

Collectors' Problem is such that, there're n different types of coupons in a box, each coupon has an equal probability of selection, and the amount of coupons of each type is unlimited. Every time, we take out k coupons and repeat until all n types are collected, consider:

Problem 1: if we take out 1 coupon every time ($k=1$), then how many times should we repeat to get all n types.

Problem 2: if we take out multiple coupons every time ($k>1$), then how many times should we repeat to get all n types.

Assume we've already got i types of coupons, let p_i be the probability of getting the next new type of coupon, x_i be the times that we have to repeat to get the next new type. We have $P(x_i = z) = (1-p_i)^{z-1} p_i$, where x_i obeys geometric distribution, with a mathematical expectation of $E(x_i)=1/p_i$. Let X_n^k be the times of repetitions until we get all n types of coupons if we take k coupons out every time. We have $X_n^k = \sum x_i$. According to the linear property of its expectation, we have,

$$E(X_n^k) = \sum_{i=0}^{n-1} E(x_i) = \sum_{i=0}^{n-1} 1/p_i \quad (1)$$

Let $k=1$, $p_i = (n-i)/n$, therefore,

$$E(X_n^1) = \sum_{i=0}^{n-1} E(x_i) = \sum_{i=0}^{n-1} \frac{n}{n-i} = n \sum_{i=1}^n \frac{1}{i} = n \cdot H_n \quad (2)$$

where H_n is the harmonic series to the degree of n . When $n \rightarrow \infty$, we have $\ln n + \gamma + O(1/n)$, where γ is the Euler-Mascheroni constant, and $\gamma \approx 0.5772156649$. Therefore,

$$E(X_n^1) = n \ln n + \gamma n + O(1) \quad (3)$$

If we take $k=1$, i.e. we get only one node per time, we will get $E(X_n^1)$ nodes in total when the snapshot is done. In fact, Trackers will reply with the size of current swarm, and k ($k=50$) randomly selected nodes. Besides, nodes replied from DHT and PEX are normally greater than 1, i.e. $k>1$. Thus, if $k>1$, $p_i = 1 - C_i^k / C_n^k$ ($i < k$, $C_i^k = 0$), According to formula (1):

$$E(X_n^k) = \sum_{i=0}^{n-1} \frac{1}{1 - C_i^k / C_n^k} = \frac{1}{k} E(X_n^1) \quad (4)$$

B. Fast snapshot model for swarms

According to formula (4), for a single fully covered snapshot to a swarm of size n , we are supposed to get information of $E(X_n^1)$ peers, which contains massive duplicates, and the duplication rate r is,

$$r = E(X_n^1) / n = \ln n + \gamma + O(1/n) \quad (5)$$

As such, the duplication rate r is proportional to $\ln n$. It is extremely not good for the measurement of large swarm. For a swarm of 5000 peers ($n=5000$), we need to acquire 45,472 peers on average to build a fully covered snapshot, and the

amortized duplication rate $r \approx 9.1$. If Trackers reply with 50 peers ($k=50$) per request, we need to make 910 requests. In order to offload Trackers, BT protocol requires the minimum time between each two requests to be $\Delta t=300s$. As such, to finish one fully covered snapshot, we need at least 3 days. However BT swarms are dynamically changing, and the snapshot acquired in 3 days are of absolutely little value.

Considering the dynamic property and the measuring cost, it's important to finish one snapshot in a relatively short time T , at the smallest cost possible. As such, we propose a fast snapshot model, and optimize the solution to the Coupon Collectors' Problem using two techniques to lower the duplication rate, shorten the time used to do snapshot, and therefore ensuring the completeness and timeliness of the collected data.

First, we should reduce the expected value without ever losing timeliness. The solution to the Coupon Collector's Problem has a really high duplication rate due to its effort to try to acquire all n types of coupons, and the cost is relatively high to get the last few types of coupons. But if we transform the problem to "acquire u of the total n types of coupons, where $u/n \rightarrow 1$ ", we can greatly reduce the times of duplicated experiments, while still ensure data completeness.

Let random variable $X_{n,u}^k$ be the average times of experiments, if we collect k peers, from a swarm of size n , per request, until u of n types of peers are acquired. According to formula (2) and (3), if n and u are sufficiently large, the average number of peers collected when $k=1$ is,

$$E(X_{n,u}^1) = n(H_n - H_{n-u}) = n \ln \frac{n}{n-u} \quad (6)$$

If we are able to get k ($k > 1$) peers per request, according to Theorem 1, the average times of request should be:

$$E(X_{n,u}^k) = \frac{1}{k} E(X_{n,u}^1) = \frac{n}{k} \ln \frac{n}{n-u} \quad (7)$$

For a swarm with 5,000 peers, if we take $u/n=0.95$, then we only need to collect 14,949 peers to cover 95% of the total peers. The duplication rate has dropped to 3.0. However, if we collect them from Trackers, we'll still need nearly 300 requests, which would cost at least 1 day.

Therefore, the second method we propose is to use multiple hosts to request in parallel. If the minimum interval between two consecutive requests is Δt , and the entire snapshot process takes a time of t , during which, if we want to accomplish $E(X_{n,u}^k)$, then the number of hosts should be,

$$y = \frac{E(X_{n,u}^k)}{t/\Delta t} \quad (8)$$

Therefore, for a swarm with 5,000 peers ($n=5,000$), if we need to achieve 95% ($u/n=0.95$) of coverage in 30 minutes ($t=30min$), and if the interval between consecutive requests is

5 mins ($\Delta t=300s$), each request is replied with 50 peers ($k=50$), then the average number of hosts need is $y=40$.

4. Experiments and Analysis

A. The Active Measurement Platform

According to our fast snapshot model, we designed and implemented a proactive measurement platform for BT swarm as shown in Fig 1. The platform is consisted of 45 machines, and has 4 different types of nodes, including 2 spider hosts, 2 database hosts, 40 measuring hosts, and 1 coordination host. Spiders, measuring and coordination hosts are running on machines with Dual-Core Pentiums 4 2.4GHz CPU, 4G memory, 250GB hard disk, 1Gbit NIC, and Windows XP SP2 installed. The database is running on Dell PowerEdge 1950 server with a single Quad-Core Xeon CPU, 8G memory, 146GB×2 SAS hard disk, and Windows 7 installed.

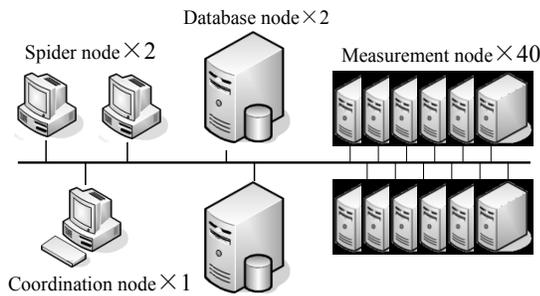


Fig 1 Active Measurement Platform.

Spiders are responsible for resolving HTML code and acquiring BT torrents from torrentz.eu automatically. Similarly, we manually selected some torrents from HDChina. HDChina is a famous PT (Private Tracker) website. The measurement nodes are running a customized Vuze client. Vuze is an open-source BT client. By modifying the source code, we made it possible to participate in BT swarms, and collect peers information. The database is responsible for storing torrents and peers information from the spider and measuring nodes. In order to increase data availability and system reliability, we adopt two servers replicating each other. The coordination node is responsible for task scheduling and data analysis. According to the system load, it assigns torrents to each measuring node, and analyzes the data collected in real time. We conducted such measurement experiment from July 15, 2011 to March 31, 2012 using the platform above, analyzed 49,854 Torrentz swarms, as well as 98 HDChina PT swarms, and collected 14,178,812 and 13,310 peers respectively.

B. Swarm Size

We first measured the distribution of the scale of BT swarms. Fig 2 shows the cumulative distribution of the scale and peers number of those 49,854 swarms from Torrentz. One obvious characteristic is that most majorities of swarms have a small scale, and the average size of 96.20% of the swarms is less than 103 peers. The scale of swarm roughly obeys the

Pareto Law, i.e. up to 81.5% of swarms collected contribute only as little as 34.83% peers. For the swarms we measured, nearly 81.5% of which have less than 300 peers, and their average size is 284. This is due to our torrents are acquired from Torrentz that covers 34 commonly used torrents publishing sites, which has more users; therefore the swarm scales are larger. Especially for those swarms of size among 10^3 - 10^4 , which only have a percentage of 3.68%, but contribute nearly 31.36% peers.

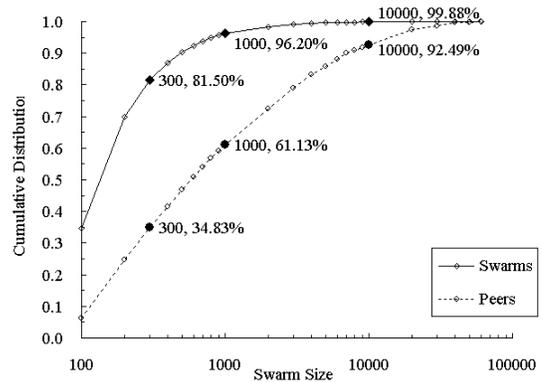


Fig 2 Cumulative distribution of the scale and peers count of Torrentz swarms.

Fig 3 shows the cumulative distribution of the scale and peers count of those 98 PT swarms from HDChina. Similar to Torrentz swarms, the scale of most of which are relatively small. Nearly 92.22% of swarms have no more than 300 peers, and the average size is 135. Since HDChina is a regional PT website, all downloading behaviors are conducted among limited user groups. Therefore, the scale is relatively small.

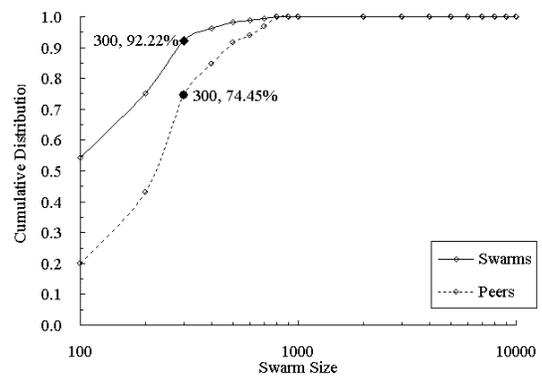


Fig 3 Cumulative distribution of the scale and peers count of HDChina swarms.

C. Geographical Distribution of Peers

Using the snapshots acquired, we leverage GeoIP database [15] to analyze the geographical distribution of peers as shown in Fig 4. The database is able to map IP addresses to their national belongs with a 99.5% accuracy. We sort the countries and regions with respect to peers count, and draw this diagram of geographical distribution of peers. As we could see, Fig 4 exhibits exponential distribution, which implies BT peers are

widely distributed all around the world. Those 14,178,812 peers we acquired from Torrentz belong to 222 countries or regions, and 62.40% of which belong to top 10 countries. Peers of the other 206 countries or regions are less than 1%. Peers from HDChina exhibits much more obvious geographical preference, where the extracted 13,310 peers belong to 30 countries and regions.

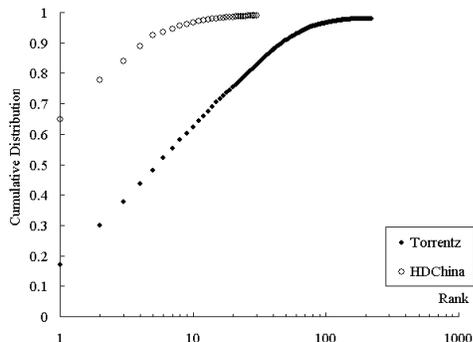


Fig 4 Geographical distribution of peers

Table 1 shows the top 5 countries as well as their belonging peers count. According to Table 1, users of regional PT websites like HDChina are mainly distributed among 30 countries or regions around mainland China, which is pretty concentrated.

TABLE I Geographic distribution of peers (Top 5)

Rank	Torrentz		HDChina	
	Code	Percentage	Code	Percentage
1	US	16.98	CN	64.90
2	IN	12.97	TW	12.92
3	UK	7.87	JP	6.18
4	CA	5.84	HK	4.97
5	PH	4.43	MY	3.43

Fig 4 and Table 1 also show that, although our measuring platform is deployed on Soochow University campus, the geographical distribution characteristics only have a relation to the source of torrents, but not the location of measuring. The reasons are: First, peers replied from Tracker, DHT, and PEX are randomly chosen, which doesn't tend to have much relation to the location of measuring. Second, BT users are distributed among all over the world, that peers downloading the same content belong to different regions. Third, the torrents extracted from Torrentz, and used in our experiment are from 34 popular torrents publishing websites with massive of users which don't exhibit obvious geographical preference. In contrast, HDChina is a regional PT websites with strong language preference, has a limited user group, therefore, their distribution is concentrated.

5. Conclusions and Future Work

We proposed a fast snapshot model based on an optimized solution to the Coupon Collectors' Problem, and implemented a low-cost proactive measuring platform to conduct long-term experiment on the global BT system. We analyzed the scale of swarms as well as the average time to build a complete snapshot to validate the correctness our fast snapshot model and active measurement platform. However, there still exist some problems, e.g. extensive snapshot time for large swarms. Although such swarms have a pretty small proportion, however, the study of large swarms has its value in investigating the performance and behaviors of peers. Furthermore, limited by the strict entry pass and sharing ratio mechanism, it's really difficult to conduct active measurement on PT systems. These are what we need to address in future.

References

- [1] T. Karagiannis, P. Rodriguez, K. Papagiannaki. Should internet service providers fear peer-assisted content distribution? [C]. In Proceedings of IMC'05, Berkeley, CA, USA.
- [2] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. G. Liu, and A. Silberschatz. P4P: Provider Portal for Applications [J]. SIGCOMM Computer Communication Review, vol. 38, no. 4, pp. 351-362, 2008.
- [3] V. Aggarwal, A. Feldmann, and C. Scheidele. Can ISPs and p2p users cooperate for improved performance? [J]. SIGCOMM Computer Communication Review, 2007, 37(3): 29-40.
- [4] D. R. Choffnes and F. E. Bustamante. Taming the Torrent: A Practical Approach to Reducing Cross-ISP Traffic in Peer-to-Peer Systems[C]. SIGCOMM'08, August 17-22, 2008, Seattle, Washington, USA, pp: 363-374.
- [5] S. L. Blond, A. Legout, W. Dabbous. Pushing BitTorrent locality to the limit [J]. Computer Networks, 2011, 55, pp. 541-557.
- [6] R. Bindal et al. Improving traffic locality in bittorrent via biased neighbor selection[C]. In Proceedings of the 26th IEEE International Conference on Distributed Computing Systems, (Lisboa, Portugal), 2006.
- [7] J. C. Liu, H. Y. Wang, K. Xu. Understanding peer distribution in the global internet [J]. IEEE Network, 2010, 24(4), pp.40-44.
- [8] 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, San Jose, CA, USA, January 2002,p.1-15.
- [9] [J. Liang, R. Kumar and K. W. Ross. The KaZaA Overlay: A Measurement Study[J]. Computer Networks,2005.p.1-25.
- [10]Liang , R Kumar, et al.. Pollution in file sharing systems[J].INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE. 2005. p.1174-1185.
- [11]M. Izal, G. Urvoy-Keller, E. W. Biersack, et al. Dissecting bittorrent: Five months in a torrent's lifetime [C]. In Proceedings of Passive and Active Measurements (PAM), (Antibes Juan-les-Pins, France), 2004.
- [12]J. Pouwelse, P. Garbacki, D. Epema, and H. Sips. The bittorrent p2p file-sharing system: Measurements and analysis [C]. In Proceedings of IPTPS'05, (Ithaca, NY, USA), 2005.
- [13]H. Wang, J. Liu, and K. Xu. On the locality of bittorrent-based video file swarming [C]. In Proc. of the 8th International Workshop on Peer-to-Peer Systems (IPTPS'09), Boston, MA, Apr.2009 April 2009.
- [14]T. Hobeld, F. Lehrieder, D. Hock, etc. Characterization of BitTorrent swarms and their distribution in the Internet [J]. Computer Networks, 2011, 55, pp. 1197-1215.
- [15]C. Zhang, P. Dhungel, D. Wu, etc. Unraveling the BitTorrent Ecosystem [J]. IEEE Transactions on Parallel and Distributed Systems, 2011, 22(7), pp. 1164-1177.