

A White-box Empirical Study of P2P-VoD Systems: Several Unconventional New Findings

Tieying Zhang*, Zhenhua Li[†], Huawei Shen*, Yan Huang[‡], Xueqi Cheng*

*Institute of Computing Technology, Chinese Academy of Sciences

Email: {zhangtiey, shenhuawei}@software.ict.ac.cn, cxq@ict.ac.cn

[†]Department of Computer Science and Technology, Peking University

Email: lzh@net.pku.edu.cn

[‡]Tencent Research, China

Email: galehuang@tencent.com

Abstract—P2P-VoD systems have gained tremendous popularity in recent years. While existing research is mostly based on theoretical or conventional assumptions, it is particularly valuable to understand and examine how these assumptions work in realistic environments, so as to set up a solid foundation for mechanism design and optimization possibilities. In this paper, we present a comprehensive measurement study of CoolFish, a real-world P2P-VoD system. Our measurement provides several new findings which are different from the traditional assumptions or observations: the access pattern does not match Poisson distribution; session time does not have positive correlation with movie popularity; jump frequency does not have a negative correlation with movie popularity as assumed in previous studies. We analyze the reasons for these results and provide suggestions for the further study of P2P-VoD services.

I. INTRODUCTION

With the rapid deployment of broadband access into households, Video-on-Demand (VoD) services have become one of the most popular Internet applications. VoD services provide users free interactivity, such as random access, pause and jump. However, due to its high bandwidth requirements, it is of extensive high-cost to provide VoD services. For instance, YouTube, the most popular on-demand video-sharing service, has to pay more than one million dollars' worth of bandwidth a month for transmission [8]. Moreover, it has to be equipped with a supercomputer with 400 nodes and a 10-gigabit Ethernet connection to provide video service.

As P2P technologies have obtained enormous success in file sharing [3], [7] and live streaming [11], [13], more and more VoD providers [4], [8], [11], [13] have begun to pay attention to developing P2P-VoD applications to reduce server costs and accelerate user downloading. Recently, researchers have focused on using multi-task downloading (MTD)¹ in P2P-Live and P2P-VoD systems to make full use of the end user's ability [15], [16], [18], [20]. In MTD systems, a peer can download multiple tasks in parallel with its current viewing video, which means peers may serve other peers with content they are not viewing. While MTD systems have been proved efficient, the key parameters in existing studies are mainly based on theoretical or conventional assumptions.

¹MTD is named VUD (view-upload decoupling) in works [15], [16] and "multi-movie system" in work [20].

In this paper, we focus on MTD systems and present a comprehensive measurement study of CoolFish [5], a real-world MTD P2P-VoD system, deployed over 28 provinces in China. Since the deployment of CoolFish, from Oct. 2008 to Feb. 2011, there has been over 5 million visits to the system. Using a large amount (3 months) of collected data, we conduct an extensive analysis of P2P evaluation on user behavior and system performance, and discuss the implications of the measurement. Due to the full control of CoolFish, we present a "white-box" view of user access patterns (e.g., jump) which would otherwise be impossible. Based on our analysis of the measurements, we propose several new findings and some interesting observations, especially the user interactivity (e.g., jump positions and jump frequency).

Our contributions can be summarized as follows:

- 1) We present several new findings in this paper. For example, we find the access number does not match Poisson distribution which is the traditional assumption; session time does not have positive correlation with movie popularity; and jump frequency does not have a negative correlation with movie popularity as assumed in previous studies.
- 2) The measurements are conducted via our real-world MTD system, CoolFish [5]. It gives a "white-box" view into user access patterns which would otherwise be impossible. CoolFish has been deployed over 28 provinces in China. It has received over 5.2 million user visits and the number of recent daily visits has exceeded 7000.

The rest of this paper is organized as follows. We give a system overview of CoolFish in Section II. Section III presents the system measurement and analysis in detail. Section IV reviews the related work. Finally, this paper is concluded in Section V.

II. SYSTEM OVERVIEW

A. Deployment

CoolFish is the first P2P-VoD system deployed in China Science & Technology Network (CSTNet), one of the four

major ISPs² in China. CSTNet is a nationwide network for the scientific and technical communities, government departments and hi-tech enterprises, with more than 1.5 million end users. It connects about 200 research institutes of the Chinese Academy of Sciences (CAS) and four campuses with more than 58,000 students. CoolFish has been very popular after released in CSTNet. From Oct. 2008 to Feb. 2011, there have been over 5.2 million user visits and the number of recent daily visits has exceeded 7000. In the hot time, there are 700 simultaneous online viewers. Table I presents the main log statistics of the system.

TABLE I: CoolFish system statistics from Oct. 2008 to Feb. 2011.

Parameter	Value
Total number of visits	≈ 5,200,000
Peak number of online users	> 700
Number of videos	> 1500
Recent daily data volume	≈ 6.08 TB
Average video bit rate	700 Kbps
Average video length	1.2 hours
Percentage of CSTNet users	78%
Percentage of NAT users	15%

The core of CoolFish is based on our previous VoD system [18], which is designed to use MTD to provide VoD service. CoolFish is able to support an average video bit rate of 700Kbps, which is about 50% higher than that of most commercial P2P-VoD systems with a video bit rate less than 450Kbps [4], [11]–[13]. The maximum bit rate supported by CoolFish is 2.5Mbps with high playback continuity (close to 1.0). This should be largely owed to the good network bandwidth of CSTNet. We can see from Table I that the majority of users are from CSTNet, with the maximum bandwidth of 24 Mbps [6]. This result also indicates that Hi-Definition video can be provided by P2P technology under good network conditions. Our current CoolFish release has implemented with over 80,000 lines of C++ codes in total. The users are distributed over 28 provinces in China. The user IP prefix distribution is shown in Fig. 1 from Google Analytics [1].



Fig. 1: User IP prefix distribution in CoolFish.

²The four major ISPs in China: China Telecom, China Netcom, CERNET, and CSTNet.

B. Architecture

CoolFish is a mesh-based network just like BitTorrent system. Fig. 2 represents the general architecture of the CoolFish system. The whole system includes a set of servers and peers (end users). The function of each component can be understood as follows:

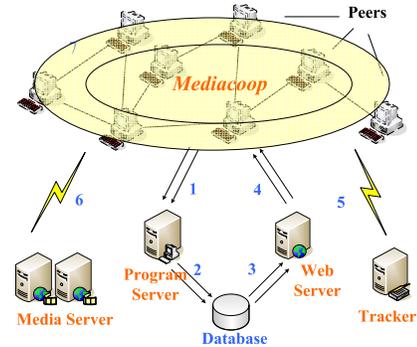


Fig. 2: System architecture of CoolFish.

- Peers report their programs (movies) information to Program Server, including video name, viewing progress, jump and pause position, startup and jump latency, and so on.
- Program Server writes the information into Database.
- Web Server reads the portal information from Database, shows the program list and online peers on the web site.
- Peers register themselves to the Tracker and gets the peer list.
- Media Server holds the full content of all the videos and provides them to the peers.
- Mediacoop provides distributed lookup mechanism.

III. OBSERVATION AND ANALYSIS

In this Section, we present the measurements and analysis based on large amount of data (three months) collected from CoolFish. We first present the overview of system scale in Section III-A. Then we classify the characteristics into three categories: user access, available time and on-demand interactivity, which are the most important parameters used in P2P-VoD systems [19].

A. Server Stress and Number of Online Users

Server stress (Server upload) is the outgoing bandwidth required at the Media Server. The lower the stress on the media server, the more scalable the system is. We evaluate the server stress with the change of the number of online users against hours. It can be seen from Fig. 3 that the number of online users fluctuates with hours of a day. The peak value appears at around 22:00 - 23:00, which is different from the previous statistical results (the peak time is around 19:00) [11]. This is because the majority of users in CoolFish are college students, and 22:00 is the most popular time for Chinese students to be on line. For the convenience of analysis, we define three representative time periods (RTP): hot time [19:00 - 01:00],

normal time [12:00, 19:00) and cold time (01:00 - 12:00). These three RTPs have different characters which will be used in the following sub-sections.

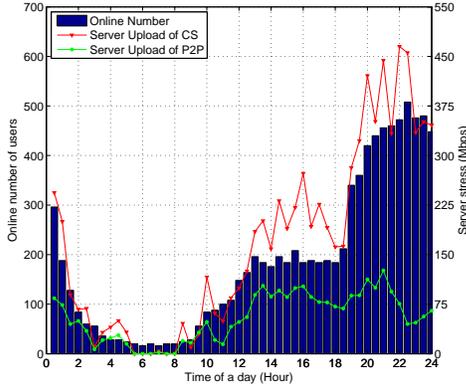


Fig. 3: Number of online users and server stress.

For server stress, we compare the architecture of P2P with the traditional client-server (C/S) network. For P2P, we take the advantage of the measurement result of the CoolFish system. For C/S, we explore server stress by using trace simulation method where we accumulate all the peers' downloading rates through CoolFish logs. The results in Fig. 3 show that the server stress of P2P is far less than that of C/S. Especially in hot time, P2P can reduce more than 70% server stress than C/S.

B. User Access

1) *Access Number*: The logs of CoolFish record the arrival time for all the users. Using these logs, we can easily compute the access number to the system. Fig. 4 shows the CDF of number of user accesses to the system within 24 hours in one day. A general trend is that the access number increases rapidly between 18:00 p.m. and 01:00 a.m., but increases slowly after 02:00 a.m. This trend is basically consistent with the representative time periods (RTPs) defined in Section III-A.

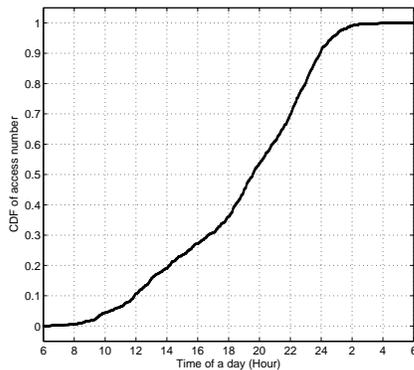


Fig. 4: CDF of access number in one day.

The access number model refers to the distribution of the number of accesses in unit time, which is one of the most important metrics for use in P2P-VoD area. In our measurement we collect the data of one month, one day, three RTPs (randomly selected), three hours (randomly selected) separately, and the unit time is 5 minutes. Fig. 5(a) represents the probability density functions (PDFs) of access number model for one month and one day. Traditionally, people think the access number model should follow Poisson distribution:

$$y = f(x|\lambda) = \frac{\lambda^x}{x!} e^{-\lambda}, x = 0, 1, 2, \dots$$

where λ is both the mean and the variance of the distribution. Therefore, we use Poisson distribution to fit the curve of access number. However, as Fig. 5(a) shows, the PDF of access number of one month does not match Poisson distribution. Further, the PDF of one day has a trend to fit the Poisson distribution but still has a big difference. Intrigued by this trend, we plot the PDF of access number for the three representative time periods (RTPs) and compare them with Poisson distribution curves with different λ (Fig. 5(b)). We can see from Fig. 5(b), the three RTPs have different PDFs. Although there are still some differences between Poisson distribution and the real data curve, we can find more similarities using Poisson distribution to fit the PDF curve. To further obtain the accurate model, we plot the PDFs of access number by hourly time division. Fig. 5(c) shows the results. We can see that the Poisson distributions are able to fit with observed data more accurately. For the three hours, we should set the Poisson distribution with different λ of 10, 12 and 14 respectively.

Here, we should analyze why there is a large difference for the PDFs of access number for monthly, daily, and hourly time divisions. Actually, the daily pattern is accumulated by hourly data. However, from Fig. 5(c) we can see that although the hourly data, in general, follows Poisson distribution, they are not perfectly matched. There are still some error-differences and outliers, which are also accumulated into daily and monthly patterns. The larger the error-difference is, the more unmatched the curve is, which is why when the sampling time area becomes larger, the access number pattern gets more and more unmatched to Poisson distribution. On the other hand, the implication from the result is useful to simulation experiments: the access number model should be set to Poisson distribution according to hourly division or smaller unit.

2) *Inter-arrival time*: Another characteristic we are interested in is the users' inter-arrival time to the system. Inter-arrival time is defined as the time interval between two successive users' arrival. For this metric, we use exponential distribution to fit inter-arrival time pattern. Exponential distribution is as follows:

$$y = f(x|\mu) = \frac{1}{\mu} e^{-\frac{x}{\mu}}, x > 0$$

where $\mu > 0$ is a scale parameter of the distribution. Again, we plot first the PDF of inter-arrival time for the data within a

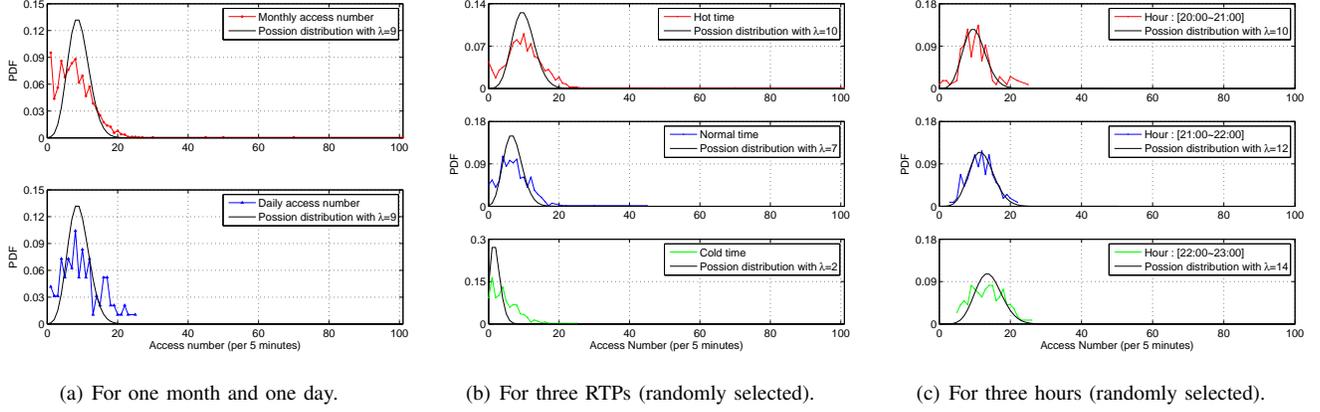


Fig. 5: PDFs of access number and fitting curves for different time divisions.

month, as shown in Fig. 6(a). Then, in order to understand this metric deeply, we plot the PDFs for the three RTPs. Fig. 6(b) shows the results. We can see that although there are some error-differences, the real data curve basically agrees with the exponential distribution curve (with $\mu = 70$). Further, Fig. 6(b) shows the detailed measurement for the three RTPs. We see that the exponential distribution with different μ can fit the real data pattern more accurately. For example, when the inter-arrival time is less than 100 seconds, the real data is mostly under the fitting curve in Fig. 6(a), while in Fig. 6(b) the fitting curves basically match the mean value. From this result, we can draw a conclusion that inter-arrival time patterns basically follow exponential distribution.

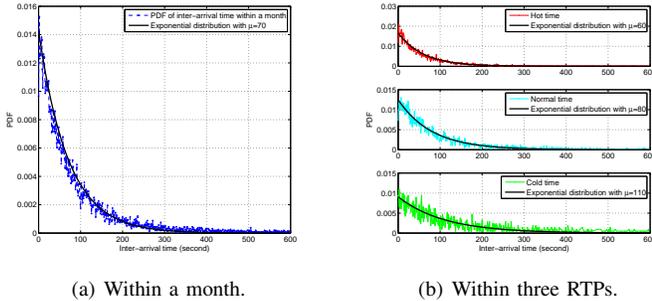


Fig. 6: PDFs of inter-arrival time and fitting curves for different time divisions.

C. Available Time

The available time of a user includes online time and session time. Online time is referred to the time duration from the user's joining the system to leaving the system, which is also called connection time. Session time is defined as the time duration of viewing a movie. When the user stops the movie or changes to another one, session time is recorded. The available time of a user is a quality measurement of use's availability. In this section, we first discuss the *Movie Popularity*, which is a related index for the following analysis. Then, we present

the measurement results of online time and session time and analyze the implications.

1) *Movie Popularity*: Movie Popularity reflects the users' desire to watch a movie. It is an important index to analyze user behavior. In Fig. 7, we show a measurement of the number of movies sorted by decreasing popularity and the corresponding CDF of the number of users. We can see that the top 20% of movies hold about 80% of users. This result does illustrate the popularity of P2P-VoD system that follows Pareto principle (80-20 rule) [9].

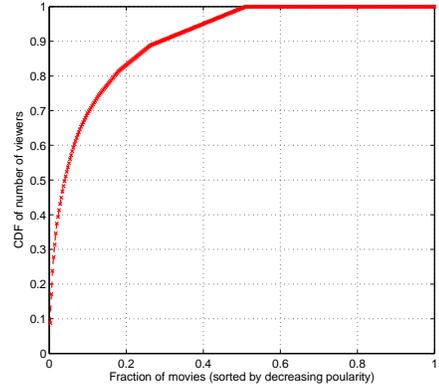


Fig. 7: Number of viewers with different popularity.

2) *Online time*: Online time measures how long a user stays in a P2P-VoD system, and it is the time duration when a user activates the P2P-VoD client software to the time when a user turns off the client software or when a user is in timeout. Online time reflects the time availability to assist each other. The knowledge of online time of peers is important because it can help us know the availability of peers and gives the instructions for further research. Fig. 8 shows the CDF of online time of peers staying in the P2P-VoD system during the three-month measurement period. From this figure, we can observe that there is a high fraction of peers (i.e., about 50%)

that stays in the P2P-VoD system more than 50 minutes, which provides the contribution of upload services to the community.

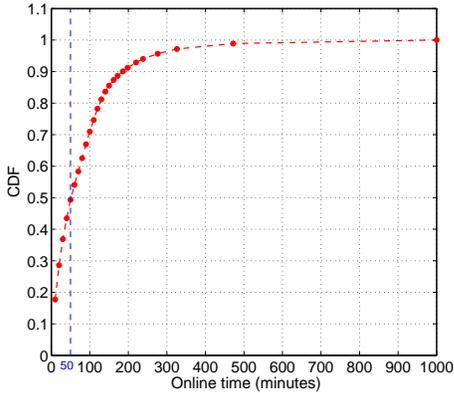


Fig. 8: CDF of Online time.

3) *Session time*: For the following analysis, we firstly define the Movie Popularity normalized as follows:

$$p_i = \frac{n_i}{n_{max}}$$

where $0 \leq p_i \leq 1$ is the popularity of movie i , n_i is the number of viewers watching movie i , n_{max} is the largest number among n_i .

Given the above statistics, we want to further find out whether peers can assist each other in the movie viewing. To answer this question, we examine the session time distribution of peers. Because the time lengths of different movies are not same (in CoolFish system, the movie length is from 45 minutes to 2.5 hours), the absolute value of session time cannot tell the whole story. Instead we use *session time index (STI)* defined as follows:

$$SessionTimeIndex(STI) = \frac{SessionTime}{MovieTime}$$

For example, if a movie length is 60 minutes (*Movie Time*), and a user takes 30 minutes to view it (*Session Time* is 30 minutes), the corresponding *STI* is $30/60 = 0.5$. *STI* describes the complete level of viewing a movie. To understand this metric, in Fig. 9, average *STI* is shown as a function of Movie Popularity. We see that for low Movie Popularity (0-0.2), *STI* has a negative correlation with the popularity. However, as popularity increases (0.2-0.6), this correlation becomes positive till the popularity is 0.6. When the popularity is larger than 0.6, *STI* tends to be stable. This is an interesting observation as most people think *STI* should have positive correlation with Movie Popularity, but why is there a negative correlation when popularity is very low? We try to explain it with the experience knowledge in VoD streaming. If a user chooses an unpopular movie, that means he pays close attention to this movie and wants to watch it with great interest. Thus, the viewing time is longer relative to others. This is why for the low popularity, *STI* is reversely high.

D. Interactive behavior

Interactivities are the vital differences distinguishing VoD streaming from live streaming, which is worthy for us to study in-depth. Interactive behaviors in P2P-VoD system mainly refer to jump and pause operations. In this section we focus on these two interactive features.

When viewing a movie, users can jump to different positions (target positions) of the movie. The knowledge of which positions users tend to jump to is important because it can help us to design the cache and schedule mechanisms. And the distribution of the target jump positions is also useful for simulation and emulation experiments. Fig. 10 illustrates the PDF of jump positions. When a user drags the track-bar to a different position of the movie (a jump event), system will generate a jump record. Fig. 10 shows the result. A general trend is that the PDF is inversely proportional to the jump position. The higher the jump position is, the lower PDF is. We also represent the liner fitting curve for the PDF of jump positions. The trend mentioned above is more obvious judging from the liner fitting curve. This illustrates that users perform more jumps to the beginning of the movie, and to the end of the movie with lower probability.

Work [4] has prove that 80% jumps have short distances ($< 300s$), which means most target positions are close to the original positions. Therefore, Fig. 10 also implies that users perform more jumps around the beginning of the movie. And as the movie continues, users perform jump operations with lower probability. Intuitively this result is in accord with the habit of user behavior. Usually when a user starts a new movie, he wants to get more information to decide whether to continue, which leads to more jumps at the beginning of the movie. When the user decides to continue to watch the movie, the number of jumps is relatively lower. Especially at the several minutes at the end of the video, the number of jumps is very low. This is because most users do not care about the content in the end of the movie, such as the cast and publisher credits. This result implies that the distribution of jumps does not follow uniform distribution which was assumed in the previous work [11]. The reason might be that the sampling in work [11] is only for one movie, while our measurement is based on three-month log data of about 400 movies.

Another interesting statistic we like to extract is the number of jumps for a given Movie Popularity. Most people [4] think that the number of jumps should have a negative correlation with Movie Popularity, because when the movie is more popular, it is more possible to attract viewers. Fig. 11 shows the relationship between Movie Popularity and number of jumps. We observe that in most of the time (popularity from 0.2-0.9) the number of jumps has stable values (from 1.9 to 2.6), and the most popular movie (popularity is 1.0) does have fewer jumps. However, we also find that when the popularity is very low (0 and 0.1), the number of jumps is significantly fewer. The reason is similar to the explanation discussed in **Session time**. If a user chooses an unpopular movie, it should be largely due to his interest and considered opinion.

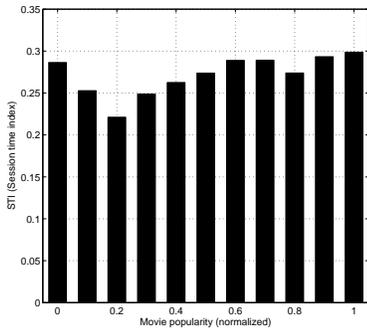


Fig. 9. STI over movie popularity.

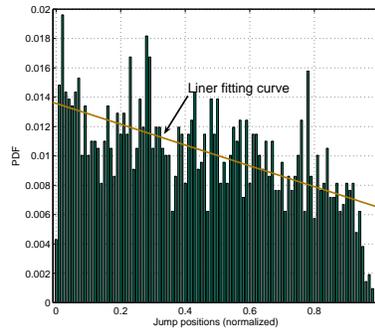


Fig. 10. PDF of target jump positions.

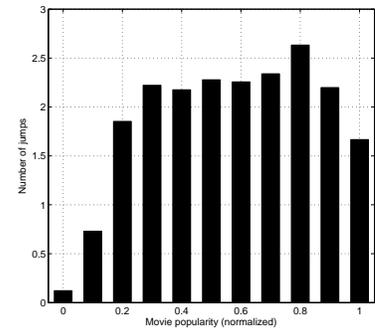


Fig. 11. Jump frequency over popularity.

Therefore, he will perform fewer track-bar drags.

As for pause, we find that most people (92%) have no pause operation in a session and the number of pause is only 0.6 in a session on average. Actually, we are not very interested in this metric, because the pause operations have no influence on the user viewing and system performance. When a user conducts pause operation, the data schedule continues to be proceeded.

IV. RELATED WORK

VoD services have attracted lots of study efforts. Yu et al. [17] present a comprehensive study of user behavior in a centralized VoD system, such as access pattern and video popularity. Huang et al. [10] study the implications of reducing the server load of the traditional VoD system by P2P technologies. They conduct in-depth analysis based on measurement and simulation using traces from MSN VoD system. However, they only study traditional VoD services while the real-world P2P-VoD features are needed. Recently, some works [2], [9], [14] focusing on P2P-VoD measurement and design are presented. However, they are either based on simulations or “black-box” measurement.

The closest works to ours are works [4], [11]. They use a “white-box” approach similar to our measurements. Hang et al. [11] presents a general architecture design of a deployed P2P-VoD system and carries out an in-depth measurement analysis of users’ behavior, scheduling strategies, and user satisfaction. While [11] provides useful insights into P2P-VoD system, it mainly focuses on the system design rather than the measurement. The study in [4] also present some measurement results, but it does not provide access and online time model.

V. CONCLUSION

While P2P technologies are widely deployed in VoD services, it is valuable to understand the P2P-VoD features in real-world system. Using three-month log data collected from CoolFish, we conduct an in-depth measurement of a real-world P2P-VoD system including access number, inter-arrival time, peers’ available time and on-demand interactivity. Based on the analysis of these measurements, we propose several new findings and some interesting observations: the access number model does not match Poisson distribution which is the traditional assumption; session time does not have positive

correlation with movie popularity; jump frequency does not have a negative correlation with movie popularity as presented in existing studies. We analyze the reasons for these results and provide suggestions for the further study of P2P-VoD services.

ACKNOWLEDGMENT

This research is supported by the National High-tech R&D Program of China (Grant No.2010AA012500) and the National Natural Science Foundation of China (Grant No. 60933005 and No. 60873245).

REFERENCES

- [1] G. Analytics. <http://www.google.com/analytics>. 2010.
- [2] S. Annapureddy, S. Guha, C. Gkantsidis, D. Gunawardena, and P. R. Rodriguez. Is high-quality vod feasible using p2p swarming. In *WWW '07*, New York, NY, USA, 2007. ACM.
- [3] BitTorrent. <http://www.bittorrent.com>. 2010.
- [4] B. Cheng, X. Liu, Z. Zhang, H. Jin, L. Stein, and X. Liao. Evaluation and optimization of a peer-to-peer video-on-demand system. *Journal of Systems Architecture*, 54(7):651–663, 2008.
- [5] CoolFish. <http://www.cool-fish.org>. 2010.
- [6] CSTNet. <http://www.cstnet.net.cn/bill.jsp>. 2010.
- [7] emule. <http://www.emule.com>. 2010.
- [8] L. Gomes. Will all of us get our 15 minutes on a youtube video? *Wall Street Journal*, 2006.
- [9] L. Guo, S. Chen, Z. Xiao, E. Tan, X. Ding, and X. Zhang. Measurements, analysis, and modeling of bittorrent-like systems. In *IMC '05*.
- [10] C. Huang, J. Li, and K. Ross. Can internet video-on-demand be profitable. *ACM SIGCOMM*, 2008.
- [11] Y. Huang, T. Fu, D. Chiu, J. Lui, and C. Huang. Challenges, design and analysis of a large-scale p2p-vod system. *ACM SIGCOMM Computer Communication Review*, 38(4):375–388, 2008.
- [12] Z. Liu, C. Wu, B. Li, and S. Zhao. Usesee: large-scale operational on-demand streaming with random network coding. In *INFOCOM '10*.
- [13] PPStream. <http://www.ppstream.com>. 2010.
- [14] N. Vratonjic, P. Gupta, N. Knezevic, D. Kostic, and A. Rowstron. Enabling dvd-like features in p2p video-on-demand-systems. In *Proceedings of Workshop on Peer-to-peer streaming and IP-TV 2007*, 2007.
- [15] D. Wu, C. Liang, Y. Liu, and K. Ross. View-upload decoupling: A redesign of multi-channel p2p video systems. In *INFOCOM 2009*.
- [16] D. Wu, Y. Liu, and K. Ross. Queuing network models for multi-channel p2p live streaming systems. In *Proceedings of IEEE INFOCOM*, 2009.
- [17] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng. Understanding user behavior in large-scale video-on-demand systems. In *ACM Eurosys'06*.
- [18] T. Zhang, Z. Li, X. Sun, and X. Cheng. Multi-task downloading for p2p-vod: An empirical perspective. In *Proceedings of ICPADS '10*.
- [19] T. Zhang, J. Lv, and X. Cheng. Mediacoop: Hierarchical lookup for p2p-vod services. In *Proceedings of ICPP '09*.
- [20] Y. Zhou, T. Z. J. Fu, and D. M. Chiu. Statistical modeling and analysis of p2p replication to support vod service. In *INFOCOM 2011*.