

A Longitudinal Characterization of Local and Global BitTorrent Workload Dynamics

Niklas Carlsson¹, György Dán², Anirban Mahanti³, and Martin Arlitt⁴

¹ Linköping University, Sweden

² KTH Royal Institute of Technology, Sweden

³ NICTA, Australia

⁴ HP Labs, USA and University of Calgary, Canada

Abstract. Workload characterization is important for understanding how systems and services are used in practice and to help identify design improvements. To better understand the longitudinal workload dynamics of chunk-based content delivery systems, this paper analyzes the BitTorrent usage as observed from two different vantage points. Using two simultaneously collected 48-week long traces, we analyze the differences in download characteristics and popularity dynamics observed locally at a university campus versus at a global scale. We find that campus users typically download larger files and are early adapters of new content, in the sense that they typically download files well before the time at which the global popularity of the files peak. The noticeable exception is music files, which the campus users are late to download. We also find that there typically is high churn in the set of files that are popular each week, both locally and globally, and that the most popular files peak significantly later than their release date. These findings provide insights that may improve the efficiency of content sharing locally, and thus increase the scalability of the global system.

1 Introduction

Today, the Internet is heavily used for content delivery. To reduce the content delivery cost, many scalable delivery techniques have been proposed. These include both server-based and peer-to-peer based designs. One promising approach is to split large files into smaller chunks, and allow clients to download these chunks in parallel from servers and/or other clients (i.e., their peers). While this chunk-based approach has proven highly effective, an open problem is how to best manage large-scale content replication systems. For example, a content distributor may want to allocate and manage resources to best match the current workload characteristics. One of the main reasons that this problem remains is due to the adage that “You cannot manage what you do not measure.” In particular, the large scale of many of these content delivery systems have made the task of measuring and understanding the relationships between the workloads observed in different parts of the network very challenging.

As a first step towards understanding these distributed workload dynamics, this paper analyzes BitTorrent usage as observed both locally and globally. BitTorrent is the

* This is the author’s version of the work. To appear in *Proc. Passive and Active Measurement Conference (PAM)*, Vienna, Austria, Mar. 2012. The original publication is available at www.springerlink.com.

original and most successful chunk-based protocol. While the use of BitTorrent for sharing copyrighted material has resulted in the shutdown of some large trackers and file hosting sites, the protocol is still responsible for a sizable fraction of Internet traffic.

In this study, we measure BitTorrent content popularity from two different vantage points over a 48-week period. First, we used passive measurements to capture all non-encrypted peer-to-tracker communication to/from a large university campus. Second, we simultaneously used active measurements to perform weekly “scrapes” of 721 unique BitTorrent trackers to extract a list of the files each maintained information on (i.e., tracked). Overall, the two datasets provide the longitudinal popularity patterns for more than 50 thousand and 10 million torrents, respectively. Our analysis is the first longitudinal multi-torrent analysis, and the first to capture the differences in the download characteristics and popularity dynamics of a large set of files observed both locally and globally. What makes this comparison particularly interesting is that there is a significant overlap in the files observed in the two datasets; 93.1% of the files downloaded on campus were observed in the globally collected tracker-based dataset. This overlap is a testament to the scope of our extensive global measurements.

Some of our key findings include: the campus users typically download larger files, particularly movies and TV shows; they typically download files well before the time at which the global popularity of the files peaks, with the exception of music files, which the campus users are late to adopt and show little interest for. These findings provide insights into what content in chunk-based systems should be “cached” locally. This would not only improve the experience for users on campus, but potentially increase the scalability of the global system as well.

While we find that there typically is high churn in the set of files that are popular each week, both locally and globally, we observe that most of the popular files in fact peak well after their initial use. This is in contrast to the behavior predicted by flash crowd models, which typically suggest that the popularity peaks close to the release of a torrent [10, 19]. Our results indicate that BitTorrent experiences the effects of the “rich-get-richer” phenomenon [1], which suggests that the future popularity of a file is proportional to the current popularity (to date) and files with many downloads in the past therefore will receive relatively more downloads in the future. This is particularly encouraging for chunk-based systems, which rely on peers to upload chunks that they already have downloaded, and often perform worse under flash crowds.

2 Related Work

Many measurement studies have looked at the characteristics of P2P content popularity. A significant number of works focused on the characteristics of content popularity over a fixed time interval, e.g., the number of downloads or the concurrent number of peers in swarms. Some of these works performed the measurements locally at ISPs or university campuses using deep packet inspection [9, 17]. Other works used global measurements to characterize the popularity of contents over a time interval, e.g., through monitoring search requests [12, 11], or by performing tracker scrapes [7, 20, 19]. Most of these studies show that the popularity of contents shows Zipf-like characteristics.

A different set of studies considered the performance dynamics of individual swarms over time. These works were either based on measurements of global content popularity [2, 13, 14], or on measurements of local content popularity [10, 16]. Such studies (e.g., [10, 16]) indicate that typical swarms stop growing very soon after their introduction, while a recent study by Zhang et al. [19] suggests that most swarms continue to grow after their major flash crowd ends.

Content popularity has been considered in many other contexts too. The majority of these works have considered the file popularity distribution over a time interval and for a specific system. For example, Zipf-like file popularity was found to apply to Web objects [4], to user generated content [6, 5, 8] and to on-demand streaming media [18]. Other studies have looked at the temporal evolution of popularity based on a single system, e.g., for the case of social media [15, 3].

In contrast to the above, our work is the first to capture and to compare the temporal popularity dynamics as observed locally (through passive measurements) and globally (through active measurements) for a content distribution system. Furthermore, to the best of our knowledge both datasets analyzed here are among the most comprehensive measurement datasets of torrent popularity.

3 Methodology

3.1 BitTorrent Overview

To facilitate scalable and efficient file sharing, BitTorrent splits content (files) into many small chunks. Each chunk may be downloaded from other peers downloading the same content. Volunteer trackers maintain state information about all peers that currently have chunks of a specific file. A peer that wants to download a file can learn about the peers sharing that content by contacting such a tracker. Upon request, the tracker provides the peer with a subset of the known peers. In exchange, the peer has to provide the tracker with information about its download progress. Additionally, the peers must inform the tracker about changes in their status (i.e., when they join or leave).

Peers typically learn about trackers by downloading a `.torrent` file from a regular Web server, or from a torrent search engine. Users may download these `.torrent` files from many sites, and many trackers may be involved in helping the local peers. The message format and information communicated to the trackers is protocol specific. In this paper, we leverage the protocol specifications to capture the information communicated to the trackers in the peer-to-tracker communication.

3.2 University Campus Measurements

Our first dataset is a trace of all non-encrypted peer-to-tracker communication at the University of Calgary, which has approximately 35,000 faculty, staff and students. The data were collected between Sept. 15, 2008 and Aug. 17, 2009. The measurements were collected using a `Bro`¹ script that extracts the application-layer information about all HTTP transactions across the University's Internet link in real time.

¹ <http://www.bro-ids.org/>

For the purpose of this paper, we focus on one specific type of HTTP transaction. In particular, we look at all peer-to-tracker HTTP requests. These transactions can easily be identified, as the URIs of these requests contain the strings “peerID”, “info_hash”, etc. The URI also contains other useful information, such as the amount of content downloaded and uploaded thus far to/from the peer. As a download typically consists of a sequence of these requests, we can track the file sharing progress as reported (to the different trackers). We do not analyze the data that the peers downloaded, or otherwise try to assess the indirect impact the content itself may have on the observed file popularities and system dynamics.

Our anonymized traces do not allow us to identify individual users across downloads, but we could identify that most of the downloads were associated with recent versions of the Vuze (52.78%) and μ Torrent (23.7%) clients. Other common clients were Transmission (10.6%), the mainline client (5.7%) and BitComet (4.2%). A more fine grained classification augmented with information about which clients had implemented various extension protocols, suggested that up to 97.5% of the downloads may have been done by a peer that implemented the peer-exchange protocol (PEX) and the DHT functionality. While these functionalities may reduce the peers’ incentive for peer-to-tracker communication, these standard clients are expected to follow the BitTorrent standards and periodically communicate with the trackers.

3.3 Global Tracker-based Measurements

Our second dataset is tracker-based, and provides a global view of content popularity. This dataset is obtained using two kinds of measurements. First, we performed “screen scrapes” (i.e., systematically extracted information from Web pages that would normally be viewed on a screen) of the torrent search engine www.mininova.org, the most popular torrent search engine according to www.alexa.com at the time of the measurements (Alexa rank of 75, August 1, 2008). From the initial screen scrape we obtained 1,690 tracker addresses. Second, we scraped all the 1,690 trackers for peer and download information of all the torrents they maintain. This allowed us to efficiently obtain the number of leechers, seeds, and completed downloads as seen by each tracker. We performed the tracker-scrapes weekly from Sept. 15, 2008, to Aug. 17, 2009, the same 48-week period as the University measurements. All scrapes were performed at 8pm GMT. We removed redundant tracker information for trackers that share information about the same swarms of peers, and identified 721 independent trackers.

3.4 Overlap between Datasets

During our 48-week measurement period, the University users used BitTorrent to download 56,963 unique files, while the 721 unique trackers maintained state information about 11.2M unique torrents. Using the mininova screen scrapes, we obtained file size and category information for 911,687 distinct files. Table 1 summarizes our datasets.

While the local peers contacted 2,371 different trackers (about 3 times as many as the monitored trackers), a total of 45,404 files (93.1% of all files observed locally) are observed in both the local and the global datasets. This large overlap is due to the extensive global dataset and the high skew in tracker usage. This is illustrated in

Table 1. Summary of datasets.

Property	University (tracker communication)	Global (tracker scrapes)	Mininova (screen scrapes)
Trackers	2,371	721	1,690
Torrents	56,963	11.2M	911,687
Downloads	1.73M	37.0B	–
HTTP requests	249M	–	–
Start date	Sep. 15, 2008	Sep. 15, 2008	Sept. 2008
End date	Aug. 17, 2009	Aug. 17, 2009	Aug. 2009
Frequency	All requests	Weekly scrapes	Twice

Figure 1, which shows the number of files and downloads associated with each tracker. The lines in this figure clearly have Zipf-like characteristics, suggesting that a few of the trackers are responsible for much of the load generated by the University. However, there is also a long tail of trackers that are responsible for a significant fraction of the total load (although each of these trackers is not responsible for much of the load). Finally, two screen scrapes (one at the beginning and one at the end of the measurement period) provide the file size statistics and category information for 16,119 (33.1%) of these files.

4 Content Download Characteristics

We first take a closer look at the differences in the downloaded content. Figure 2 shows the cumulative distribution function (CDF) of the file sizes observed in the two datasets. While the majority of downloads in both datasets are for files between 100 MB and 10 GB, there is a very large difference in the files size distributions. Overall, the files downloaded in the University network are much bigger. (Note that x-axis is in log scale.) This difference suggests that globally, many users have lower access bandwidths than the relatively high access bandwidths of the local university users. As a result, the global users may be more reluctant to download larger files.

Figure 3 shows the number of downloads per category in each dataset. The categories used here are those defined by Mininova. Figure 3 shows that the university users download a much larger fraction of movies and TV shows (i.e., video) than the global users, which download a lot more music (these tend to be much smaller files). This further explains the large difference in file sizes observed locally and globally. To set these results in context, Figure 4 shows the file-size distribution for the five most downloaded file categories in our datasets. While the distributions differ significantly between categories, all of the categories have highly variable files sizes.

5 Download Popularity Dynamics

We next compare the download dynamics of the two datasets. Figure 5 shows the cumulative distribution function of when the downloads take place relative to the time that each (individual) file’s popularity peaked globally. We calculate the peak popularity of each file as the mid-point of the week that it had the most downloads. In the case of ties,




Fig. 1. Tracker load and content per tracker, as




Fig. 2. File size distributions for university and global dataset.




Fig. 3. Breakdown of downloads per category (based on Mininova classification).




Fig. 4. File size distributions category breakdown using Mininova dataset.

we pick the earlier week. Approximately 40% of the downloads in the global dataset happen during the peak week, 20% before the peak week, and 40% after the peak week. Furthermore, there is a substantial tail of views both before and after the peak week.

Figure 5 also shows the results for the university dataset. There is a substantial difference in when the downloads take place on campus. For example, more than 70% of the downloads in the university dataset take place before the global peak week. In fact, roughly 40% of the downloads takes place at least 10 weeks before the global peak week. This is particularly interesting as it may suggest that users at well-connected universities can be used for predicting content that will become increasingly popular.

To assess the generality of this observation, Figure 6 breaks down the percentage of downloads that took place on or before the week that the global popularity peaked for each distinct file in a particular category. The fraction of early downloads is much higher in the university dataset for almost all categories, particularly for anime, games, and TV shows. For music files, on the other hand, the University users seem to be late to follow current trends. We hypothesize that the combination of high speed Internet access and tech-savvy users alters the user behavior, thereby changing the workload. We note user groups that are early adapters of a content type provide good content-sharing opportunities for that content, and that peer storage could be prioritized based on the content types for which the peer have been found to be an early downloader.




Fig. 5. Time of downloads relative to the global peak week of that file.




Fig. 6. Downloads at or before global peak week for different file categories.




Fig. 7. Time of the local peak week relative to the global peak week.




Fig. 8. Average time until peak for files of different popularity.

We next consider the difference in the actual peak week, as observed at the university and globally. Figure 7 shows the CDF of the relative difference in weeks. In addition to the CDF for all files, we also show three lines for the set of videos that are among the 10, 100, and 1,000 most popular files, as observed globally. These results confirm our earlier observation that the university users are early downloaders; however, they also show that the differences in peak times are biggest for the files that are the most popular (e.g., in the top 10 curve). This suggests that well-connected university users may be particularly early to download content that will become highly popular. The significant number of early local peaks for the “all curve”, suggests that there is also a substantial tail of niche content that the University users are early to download. Thus, there may be lessons to be learned about future (global) popularity, by observing content consumption trends on (well-connected) university networks.

While previous studies have observed early flash crowd behavior immediately after the release of content [10], the fact that we observe big differences in the peaks of popular content suggests that the global peak of popular content often happens well after the release of the content. Figure 8 shows the average time (in weeks) until the global peak for content of different popularity, as a function of the number of downloads during the files’ peak week and over the full measurement period, respectively. Here, we applied logarithmic binning, with four bins per scale-factor 10. This figure supports




Fig. 9. Time until global peak.




Fig. 10. Rich get richer analysis.

our hypothesis that popular content in fact often peaks well after its release (or in this case, conservatively, after we first discover it having been released). While local results are omitted for brevity, this observation is consistent both locally and globally.

Figure 9 shows the CDF of the time until peak for the most popular files in the global dataset. We note that the most popular files (top 100) peak well after we first discover them; e.g., more than 60% of these files peak at least 20 weeks after we first observe downloads. As suggested by the previous figure, for less popular files there are more files that peak early. However, also for these files there is a substantial amount of files that peak several weeks after first observed.

The fact that most of the popular files peak rather late suggests that the popularity dynamics of these files are not determined by any initial flash crowd behavior when the file is released [10], but that popular files instead often build popularity over time. Preferential attachment and rich-get-richer models have been applied to various domains, and suggest that files that have gained many downloads typically gain even more downloads due to the fact that they are popular. To validate if the global popularity in fact follows this characteristic, Figure 10 shows the number of downloads in week 27 of a video’s lifetime, as a function of the number of downloads in week 26, and as a function of the total number of downloads up to and including week 26. This corresponds to the views at the half-year point of the video’s lifetime. Again, we used logarithmic binning (with four bins per scale factor 10). We note that the previous week’s downloads (in week 26) on average is a very good predictor of the number of downloads during the following week. In fact, the slope is very close to one. Also, the current number of downloads (up until the current week) shows strong correlation; however, the direct relationship is much weaker.

In general, ignoring the age of the videos, the Pearson’s correlation coefficient between consecutive weeks (after log transformation) is on average 0.35 and 0.68 (with standard deviations of 0.12 and 0.15) for the university and global datasets, respectively.

To summarize, our results suggest that rich-get-richer models may help describe the BitTorrent popularity, and that a (popular) file’s popularity often peaks long after it is released. This observation is in contrast to previously proposed flash crowd models that include a decreasing number of downloads per time unit following the initial flash crowd, but is similar to recent observations by Zhang et al. [19]. This suggests that search mechanisms are not as effective as they could be. The adoption of other tech-




Fig. 11. Hotset churn analysis.

nologies (e.g., RSS feeds, online social networks) may also reduce the delay between when content is released and when its popularity peaks.

6 Hotset Churn

We next assess the weekly churn in content popularity. For this analysis we focus on the “hotset”, the set of most downloaded files in a given week. Figure 11(a) shows the overlap in the hotset observed in consecutive weeks at the university or global dataset, and the fraction of files in the university hotset that simultaneously are in the global hotset. While there are high variations, we note that there typically is high churn in the hotset, with on average only 20-40% weekly overlap. The overlap between the hotsets of the two datasets for any given week is even lower, with less than a 10% overlap.

Among the set of files that at some point is in the hotset, we found that 33.4% (19.2%) of the files at the university (globally) enter the hotset multiple times. While the majority of files (84.8% and 93.7%) only remain in the hotset for a single week, 1.6% (0.7%) of the files remained in the hotset for at least five weeks. We note that the smaller percentage of long-term popular files in the global dataset is likely due to a higher overall churn (as indicated by more one-timers and one-week cases). As a fraction of the number of files in the hotset, however, the datasets appear much more similar. These observations are illustrated by Figure 11(b), which shows a given weeks overlap with the hotset of week 20. In a few weeks the hotset is almost entirely replaced, both datasets contain files that remain in the hotset for a longer period, and some number of files (in both datasets) appears to move in and out of the hotset multiple times.

7 Conclusions

Using two simultaneously collected datasets, capturing the download characteristics and the content popularity dynamics observed both at a university campus and by a large set of global BitTorrent trackers, this paper analyzed the differences in workload dynamics observed locally versus at a global scale. We find that users on a well-connected university network download larger files and are early adopters of new content; they download files well before the time at which the global popularity of the files peak.

These characteristics suggest design modifications to content replication systems such as BitTorrent, to localize the effects of “rich” users, and improve the scalability of the global system. Exploring these alternative designs is left for future work.

References

1. Barabasi, A., Albert, R.: Emergence of scaling in random networks. *Science* 286(5439), 509–512 (Oct 1999)
2. Bharambe, A.R., Herley, C., Padmanabhan, V.N.: Analyzing and Improving a BitTorrent Network’s Performance Mechanisms. In: *Proc. IEEE INFOCOM* (Apr 2006)
3. Borghol, Y., Mitra, S., Ardon, S., Carlsson, N., Eager, D., Mahanti, A.: Characterizing and modeling popularity of user-generated videos. In: *Proc. IFIP PERFORMANCE*. Amsterdam, Netherlands (Oct 2011)
4. Breslau, L., Cao, P., Fan, L., Phillips, G., Shenker, S.: Web Caching and Zipf-like Distributions: Evidence and Implications. In: *Proc. IEEE INFOCOM* (March 1999)
5. Cha, M., Kwak, H., Rodriguez, P., Ahn, Y., Moon, S.: I Tube, You Tube, Everybody Tubes: Analyzing the World’s Largest User Generated Content Video System. In: *Proc. ACM IMC* (2007)
6. Cheng, X., Dale, C., Lui, J.: Understanding the characteristics of internet short video sharing: Youtube as a case study. In: *Proc. IWQoS* (2008)
7. Dán, G., Carlsson, N.: Power-law revisited: A large scale measurement study of P2P content popularity. In: *Proc. International Workshop on Peer-to-peer Systems (IPTPS)* (Apr 2010)
8. Gill, P., Arlitt, M., Li, Z., Mahanti, A.: YouTube Traffic Characterization: A View from the Edge. In: *Proc. ACM IMC* (2007)
9. Gummadi, K., Dunn, R., Saroiu, S., Gribble, S., Levy, H., Zahorjan, J.: Measurement, modeling, and analysis of a peer-to-peer file-sharing workload. In: *Proc. SOSP* (2003)
10. Guo, L., Chen, S., Xiao, Z., Tan, E., Ding, X., Zhang, X.: Measurement, Analysis, and Modeling of BitTorrent-like Systems. In: *Proc. ACM IMC* (Oct 2005)
11. Hefeeda, M., Saleh, O.: Traffic modeling and proportional partial caching for peer-to-peer systems. *IEEE/ACM Trans. on Networking* 16(6), 1447–1460 (2008)
12. Klemm, A., Lindemann, C., Vernon, M.K., Waldhorst, O.P.: Characterizing the query behavior in peer-to-peer file sharing systems. In: *Proc. ACM IMC* (2004)
13. Legout, A., Urvoy-Keller, G., Michiardi, P.: Rarest First and Choke Algorithms Are Enough. In: *Proc. ACM IMC* (Oct 2006)
14. Menasche, D., Rocha, A., Li, B., Towsley, D., Venkataramani, A.: Content Availability in Swarming Systems: Models, Measurements and Bundling Implications. In: *ACM CoNEXT* (Dec 2009)
15. Mitra, S., Agrawal, M., Yadav, A., Carlsson, N., Eager, D., Mahanti, A.: Characterizing web-based video sharing workloads. *ACM Tran. on the Web* (2), 8:1–8:27 (May 2011)
16. Pouwelse, J.A., Garbacki, P., Epema, D.H.J., Sips, H.J.: The bittorrent p2p file-sharing system: Measurements and analysis. In: *Proc. IPTPS* (2005)
17. Wierzbicki, A., Leibowitz, N., Ripeanu, M., Wozniak, R.: Cache replacement policies for P2P file sharing protocols. *Euro. Trans. on Telecomms.* 15, 559–569 (2004)
18. Yu, H., Zheng, D., Zhao, B., Zheng, W.: Understanding User Behavior in Large-Scale Video-on-Demand Systems. *SIGOPS Oper. Syst. Rev.* 40(4), 333–344 (2006)
19. Zhang, B., Iosup, A., Pouwelse, J.A., Epema, D.: Identifying, analyzing, and modeling flashcrows in bittorrent. In: *Proc. IEEE Peer-to-Peer Computing*. Kyoto, Japan (Aug/Sept 2011)
20. Zhang, C., Dhungel, P., Wu, D., Ross, K.W.: Unraveling the bittorrent ecosystem. *IEEE Transactions on Parallel and Distributed Systems* 22, 1164–1177 (Jul 2011)