Analyzing Video Services in Web 2.0: A Global Perspective

Mohit Saxena, Umang Sharan, and Sonia Fahmy
Department of Computer Science
Purdue University
{msaxena,usharan,fahmy}@cs.purdue.edu

ABSTRACT

Serving multimedia content over the Internet with negligible delay remains a challenge. With the advent of Web 2.0, numerous video sharing sites using different storage and content delivery models have become popular. Yet, little is known about these models from a global perspective. Such an understanding is important for designing systems which can efficiently serve video content to users all over the world. In this paper, we analyze and compare the underlying distribution frameworks of three video sharing services – YouTube, Dailymotion and Metacafe – based on traces collected from measurements over a period of 23 days. We investigate the variation in service delay with the user's geographical location and with video characteristics such as age and popularity. We leverage multiple vantage points distributed around the globe to validate our observations. Our results represent some of the first measurements directed towards analyzing these recently popular services.

1. INTRODUCTION

The Internet is emerging as a prime broadcast medium offering Television, Radio, Cinema, and the exchange of videos for personal as well as commercial use. With the increasing demand, numerous websites offering a variety of options for sharing user-generated multimedia files have become available. YouTube, Dailymotion, and Metacafe are among the most popular video sharing services on the Internet today; they were the top three video streaming sites in the SeoMOZ Web 2.0 survey 2007 [2].

The new services leverage a video delivery technique known as *pseudo-streaming* [3]. This differs from traditional web streaming in that the video content can be played back as it is being progressively downloaded. The content, unlike traditional streaming, is delivered using HTTP/TCP through generic Web servers. While there are extensive studies of traditional and live web streaming [4, 5] and their workloads [6, 7], and content distribution using content delivery networks [8, 9] or peer-assisted approaches [10, 11], there is little data on the recent pseudo-streaming services. Only a few recent studies have analyzed how user-generated content is viewed and distributed on the Internet [3, 12, 13, 14]. However, these studies either passively monitor the pseudo-streaming multimedia traffic to characterize media access patterns, or they analyze the distribution of user-generated content available on these websites by crawling them.

In this paper, we present the first active measurement study that compares the performance of three prominent video sharing services – YouTube, Dailymotion, and Metacafe[1] – from the view-point of geographically diverse users. We address the following questions: (1) How is the performance of the new services affected by the structure of the storage and content-delivery models they employ, (2) How do these services currently take into account the individual content (video) characteristics such as age and popularity, and (3) What conclusions can be drawn about the efficiency of these approaches in serving users all over the world?

We select service delay as the simplest metric for capturing the performance of the pseudo-streaming service itself, independent of the client-side system or application. We define *service delay* as the time taken to fetch 1 MBytes of a Flash Video (FLV) stream in 50 KByte increments. We quantify the variations in service delay on the basis of user characteristics such as geographical location, and video characteristics such as the age and the popularity of the media content.[2] Age of a video file is defined as the time elapsed since the video was uploaded to the website. Popularity of a video file is defined as the number of times the video has been viewed since it was uploaded. We believe that the results presented in this paper are important to understand the frameworks on which these services are built, and how they can be improved towards a better user experience.

2. BACKGROUND

Services such as YouTube, Dailymotion, and Metacafe are popular because they bring together two of the most important features of Web 2.0: social networking and content sharing. These services are based on user-driven content whereby any user can register, and upload, view, and share video clips. Each uploaded file is filtered and processed and thereafter, it is available on the website for public viewing. Table 1 gives a brief summary of these three services.

Presently, YouTube is the market leader in online sharing of video and amongst the top five websites on the Internet

[1] www.youtube.com,www.dailymotion.com,www.metacafe.com
[2] Collected datasets are available at [1].

NOSSDAV '08 Braunschweig, Germany

	YouTube	Dailymotion	Metacafe
Unique Visitors (x 10^6/month)	70	10	10
Videos Watched (x 10^6/day)	100	25	15
Alexa Rank (Feb '08)	3	31	179
File Formats (for uploads)	WMV, MOV, MPG RAM, AVI, ASF	AVI, WMV, MP4, ASF RAM, MPG, MOV	AVI, WMV, MPG, QT MOV, 3GP
How to Upload	Website, Cellphone	Website, Cellphone, Webcam	Website, Desktop application
Video Data (2006)	45 TB	*unknown*	*unknown*

Table 1: Overview of YouTube, Dailymotion, and Metacafe.

in terms of traffic according to Alexa[3] rankings. Dailymotion, on the other hand, was founded at the same time as YouTube and has gradually become quite popular due to high video and audio quality, and high streaming speeds. Metacafe is the oldest among the three video streaming services. Its initial growth was slow, but by 2006 the website traffic increased rapidly, and now, Metacafe has nearly 10 million unique visitors per month.

3. METHODOLOGY

YouTube, Dailymotion, and Metacafe are highly dynamic services with thousands of new videos being uploaded daily. Video meta information such as number of views, age of file, and ratings are also frequently updated. Such vast databases require scalable systems for storage, and robust content delivery networks for serving media. In this section, we describe our framework for inferring the architecture and performance of the three services (Section 3.1), then discuss the utility of multiple vantage points (Section 3.2), and summarize our traces (Section 3.3).

3.1 Data Collection Framework

Our measurements proceed in two phases. In the first phase, we measure the latency associated with serving video streams, and in the second phase, we gather meta information about video streams and location information about the content servers where files are stored.

We define *service delay* as the time taken to fetch 1 MBytes of a Flash Video (FLV) stream in 50 KByte increments (order of FLV frame size)[4]. We compute the mean service delay for a video by averaging the service delays of three consecutive 1 MByte chunks, in order to eliminate the effect of variable-sized files. We fetch 3 MBytes per video, because we observe that the interesting variations in service delay measurements occur in the first few video blocks regardless of its total size, as we will show in Section 4.2. Our metric captures the network latency and the server processing time associated with a video. Since all three services use the FLV format to serve videos, we do not include the time taken to decode and render the stream when measuring service delay. Incremental service delay proves to be a good metric to objectively mimic and evaluate the performance of a video pseudo-streaming service, though it may not be a direct quantifier of the end-user video quality perception. We selected this metric because it makes our measurements independent of the client-end application (including any flow control or jitter buffer scheme employed), as our focus is on comparing the relative performance *of the three services* from different clients, independent of the client system and application properties. This is achieved as there is no simultaneous playback when we capture the stream over the TCP socket connection.

[3] http://www.alexa.com

[4] Our results with 10 KByte increments showed consistent trends; we omit them from the paper for brevity.

In the first phase, we crawled the three websites for video links to build a link database, summarized in Table 2(a). For each link present in the database, we downloaded the video from the corresponding website and measured the mean service delay. Both the link farming phase as well as the service delay measurement phase were fully automated. To account for the variable loads on the websites, content servers and network conditions, we captured our traces at different times over a period of 23 days (during June-July and September-October 2007) and averaged the results. In order to minimize the possible limitations of the using a testbed, as mentioned in [16], we ran several experiments at different times of day and different days of week. We note that there is very little variability in our observations with respect to these factors.

In the second phase, we explored what happens behind the scenes, i.e., the behavior of content delivery networks, their locations, and their behavior relative to the content we download. We crawled the three websites to gather (1) the actual URL of the FLV stream associated with each video (which we get as part of the HTTP redirect response to our GET query), and (2) the IP address of the content server returned to us by our DNS query. Further, we retrieve the meta information about the video itself which includes (1) age of the video, (2) video popularity, and (3) run length. This phase utilized python scripts as well and the dataset is summarized in Table 2(b).

3.2 Vantage Points

To evaluate global performance, we collected service delay measurements and content distribution patterns from nodes at Purdue University, and from PlanetLab nodes in California (US), Brazil, UK, France, Italy, Switzerland, India, Japan, and China. Table 2 gives the statistics for four of these locations only (due to space constraints). We were careful to use PlanetLab nodes from different types of sites (academic institutions as well as commercial sites) [17]. We note that our traffic always traverses a commercial network regardless of the client network type; hence, our results are representative of real scenarios [18]. PlanetLab nodes were selected for good coverage of different vantage points around the globe. In addition, we chose PlanetLab nodes with varying bandwidth limits (from 500 Kbps to 100 Mbps) to capture the impact of client bandwidth in our service delay analysis. We utilized Pathneck [15] to analyze the bottleneck bandwidths from different client locations to the content servers. We show results with different client network types, bandwidth, and geographical location in our analysis in Section 4.2.

# Videos	US	UK	Japan	India
YouTube	4,360	4,579	4,591	2,588
Dailymotion	3,383	3,391	3,400	2,300
Metacafe	4,482	4,483	4,483	2,471

# Videos	US	UK	Japan	India
YouTube	116,312	29,835	58,176	939
Dailymotion	55,448	27,710	27,630	2,133
Metacafe	88,470	22,618	34,925	1,323

Table 2: Number of Videos (a) Service delay measurement phase (b) Meta information retrieval phase.

3.3 Data Summary

Over a period of 23 days (during June-July and September-October 2007), we fetched a total of 16,118 videos from YouTube, 12,474 videos from Dailymotion, and 15,919 videos from Metacafe in the service delay measurement phase. Videos were chosen from the most viewed, most recent, and the most relevant categories so as to uniformly cover videos with different age and popularity. Table 3 shows the distribution of the videos in our data-set relative to their age, popularity, and run-length (duration). Our data-set includes videos as recent as 4 minutes old to those which are more than 3 years old. Similarly, we have video samples with zero views to those with more than 60 million views. Video run-length varied from almost 0 seconds to nearly 10 hours. We again note that we do not intend to analyze the content properties or its distribution, nor compare it with traditional media streaming workloads, as this has been studied in [3, 12].

In the second phase, we crawled 245,247 videos from YouTube, 137,936 videos from Dailymotion, and 177,156 videos from Metacafe to retrieve the meta information of these videos, and to gather information about the content delivery system through which these videos were distributed.

Table 2 lists the number of videos fetched during the two phases from a sample of locations. We recorded 2,405 distinct YouTube content servers and Google Web Caches, and 1,252 distinct Dailymotion proxy and content delivery network (CDN) servers serving content on behalf of Dailymotion. For Metacafe, although its CDN servers were distributed all over the world, we observed only 92 distinct IP addresses from which the videos were being served. We used the Maxmind Geo-API[5] and ip2location[6] databases to map IP addresses to their respective geographical locations.

Though the entire data collection process was automated, we faced two difficulties. First, due to the dynamic nature of the freshly posted media on the three service websites (where uploaded files are being continuously filtered for objectionable content), we had to prune dead links and update our link farms continuously. Second, it was imperative to manually revise location information in the traces and logs. Though we utilized a good database for mapping IP addresses to their geographical locations, we had to cross-check it with another database to rectify incorrect mappings in some cases. Research projects, e.g., NetGeo [19], can help increase the accuracy of such mapping tools.

4. RESULTS AND ANALYSIS

In this section, we analyze the data obtained as described in Section 3. Our aim is to answer the questions posed in Section 1 on service quality, user location, and content age and popularity. We first infer the underlying content delivery frameworks of the three services. Then, we give a comparative analysis of service delays, and their dependence on content age and popularity.

[5]http://www.maxmind.com
[6]http://www.ip2location.com

4.1 Delivery Frameworks

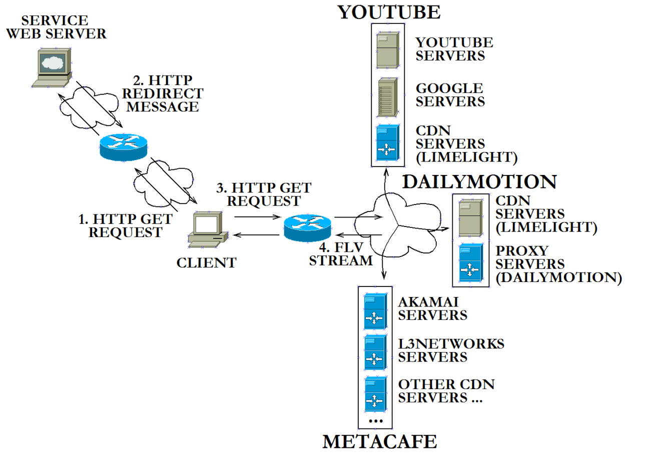

Figure 1: Content delivery frameworks.

Fig. 1 summarizes the content delivery frameworks that we inferred for the three services. From our traces, we observe that most of the YouTube content is served by YouTube's servers located in San Mateo (nearly 77%), followed by Google servers and caches (in Mountain View: 22%), and the rest (only 1%) from Limelight Network's CDN servers. Dailymotion's content is served by its proxies (in France: 85%), and the rest by CDN servers (Limelight). Metacafe uses only CDN servers (Akamai, L3 Networks) for distributing its content. When a video request is sent to Metacafe, it returns a redirection to a URL of the kind akvideos.metacafe.com or l3videos.metacafe.com followed by content related information. In contrast, YouTube either returns a URL for a YouTube server or an IP address mapping to cache.googlevideo.com. Dailymotion uses URLs of the kind proxy-xx.dailymotion.com or limelight-xx.cdn.dailymotion.com for redirection.

We observe that while YouTube and Dailymotion use a selective model to push their content, Metacafe pushes all its content to CDN servers, and relies on DNS-level load balancing to redirect the user requests. The distribution of content servers for the three services as observed in traces collected from multiple vantage points is shown in Fig. 2. Fig. 2(a) gives the distribution of the locations from which YouTube content is being served in our traces: San Mateo (YouTube servers), Mountain View (Google Caches and Servers), and from a few locations in Tempe, Arizona and the Netherlands (Limelight Networks CDN). It is interesting to note that even though YouTube is the largest video repository in the world, almost all of its content is being served from just a few locations in California, USA, to users around the globe. Fig. 2(b) depicts Dailymotion's mixed approach

	Age (minutes)			Popularity (# views)			Duration (minutes)		
	Min	Max	Median	Min	Max	Median	Min	Max	Median
YouTube	420	1,051,200	302,400	0	60,585,402	45,939	0	583	3.9
Dailymotion	4	1,252,800	172,800	0	2,074,886	626	0	369	3.2
Metacafe	20	1,576,800	216,000	0	6,936,944	3143	0	87	1.9

Table 3: Video characteristics and distribution.

of serving content using its proxy servers in France, and the rest using CDN servers. Fig. 2(c) illustrates that Metacafe uses a CDN-only approach – all its content is pushed to CDN servers distributed all over the world.

4.2 Service Delay Analysis

Fig. 3 depicts the cumulative distribution of service delays for YouTube, Dailymotion, and Metacafe from different vantage points for our data-set described in Table 2. An interesting observation from Fig. 3(a) is that the median service delay for YouTube as observed from our US traces is nearly 6.5 seconds; that for Dailymotion is 1.25 seconds, and in the case of Metacafe it is ∼1 second. This implies that, on the average, YouTube delivers 1 MByte of video content nearly 6 times slower, compared to Dailymotion and Metacafe. Similar trends are observed in our measurements from the UK.

However, in case of measurements from Japan (Fig. 3(b)), Dailymotion and YouTube incur nearly the same service delay. The median service delay for Metacafe is around 2.5 seconds – one-third of the time taken by Dailymotion and YouTube. These trends are consistent with our observation in Section 4.1 that Metacafe has CDN servers located in Tokyo, Japan. Dailymotion videos are served from its proxy servers in France and from CDN servers located in US, which explains its curves. This result clearly illustrates the impact of the proximity of the end-user to the content server.

Service delays collected from a client located at a *commercial* site in India are reported in Fig. 3(c). We observe that the service delay for all the three services has increased in this case which is also attributed to the comparatively low client bandwidth available at this site. For Metacafe, the service delay increased considerably because most of its videos were being served from the US in this case, even after the presence of its CDN servers in Tokyo. The increase in median service delay of YouTube and Dailymotion also illustrates the impact of the last-mile access bandwidth and client network type on the performance of these services. These results suggest that *with increasing last-mile bandwidths, the latency of these services will be dominated by their content distribution frameworks.*

The trends in Fig. 3(a), 3(b) indicate the highest service delays are for the most popular service among the three, which is YouTube. This was initially surprising, but investigation of the FLV bit rates (which indicate of the video playback quality) and the typical video playback time explains why a typical user does not perceive this.

We estimate the bit rates of the flash video content as the ratio of the video's file size (obtained from its HTTP header) and its run length (obtained from the meta-information of the video). Table 4 shows these estimates for the videos accessed. We can infer that the typical run-length of a 1 MByte FLV video stream varies between 20–30 seconds. This means that even though, on the average, YouTube is slower than

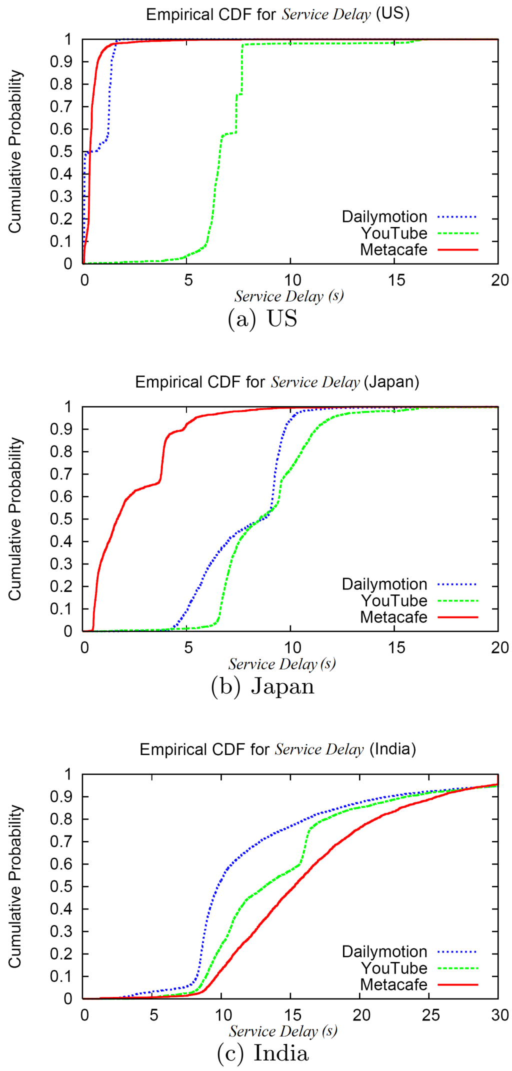

Figure 3: Distribution of service delays for 1 MByte chunks from different locations.

Dailymotion and Metacafe, a user cannot perceive this difference in typical cases. This is mainly due to the large difference between the service delay and the playback duration of the FLV stream for all three services today. At the same time, we note that as these services move on to serving high definition videos, this difference will be perceived by the end-user.

From the perspective of a user who wishes to jump to a point forward in the video stream, these service delay comparisons become more relevant even today. Many services including YouTube provide VCR functionality, whereby users

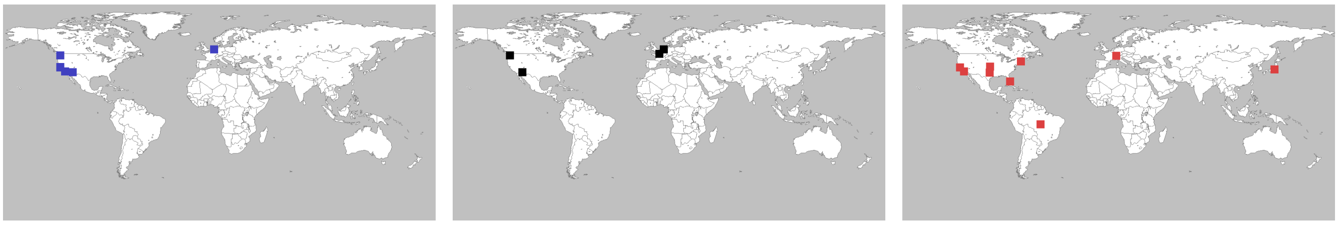

(a) YouTube (b) Dailymotion (c) Metacafe

Figure 2: Content server locations for different services.

can jump forward in a video stream, and the video is then buffered from the point indexed by the user in the stream. The real question now is how fast the service can respond when the video stream starts buffering for the first time or when it is re-indexed by the user. To investigate this, we analyze the service delay on a much finer granularity.

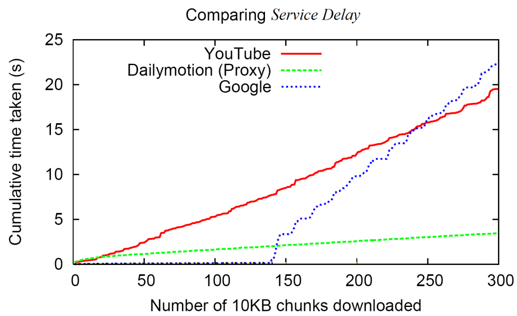

Figure 4: Cumulative service delay comparison.

Fig. 4 gives the cumulative service delay to fetch flash videos in incremental steps of 10 KBytes, as it was recorded in our US traces. For the sake of clarity, we do not plot the results for Metacafe and Dailymotion videos (available on CDN servers) here, because they incur much lower cumulative delays as compared to YouTube, Google, and Dailymotion videos (available on Dailymotion proxy servers). We plot different curves for YouTube videos fetched from YouTube servers and those retrieved from Google's domain. Fig. 4 shows that for the first few blocks, Google videos are retrieved very rapidly, compared to YouTube and Dailymotion. However, there is a sudden spike in the curve when nearly 1.4 MBytes of the FLV stream have been retrieved. This trend is observed in almost all the YouTube videos fetched from Google servers. We believe that this is an optimized content delivery method because once an initial part of the video has been buffered at the client, the rate can be stabilized to a lower value rather than transmitting the entire video stream as fast as possible (which would require significantly more server-side resources). This approach is usually associated with server side Web Caches and accelerators. This result indicates that for sequential viewing, *it is not always imperative to stream at constantly high rates to provide a good end-user experience.*

	FLV Video Bit Rates (Kbps)			
	Min	Max	Median	Mean
YouTube	29	1,340	314	289
Dailymotion	30	1,024	456	402
Metacafe	51	1,415	378	338

Table 4: FLV video encoding bit rate distribution.

4.3 Content Pushing: Popularity Model

Fig. 5 depicts the cumulative distribution of the videos according to their popularity. The plot is based on the meta information collected for each video while crawling the three services from Purdue University. Our YouTube sample set consisted of 54,426 videos, out of which 41,945 were found on YouTube servers, 12,191 on Google caches and 290 on Limelight Network servers. Fig. 5(a) shows that almost 80% of the videos available on Google caches have views greater than 100,000 (high viewership). In contrast, the median number of views for videos available on YouTube servers is 23,806. The Dailymotion sample set contained information about 27,725 videos out of which 23,694 were available on Dailymotion's proxy servers in France and 4,031 were found on CDN servers. Fig. 5(b) shows that more than 97% of Dailymotion videos found on CDN servers had views greater than 13,015 (high viewership). On the other hand, more than 90% of videos available on Dailymotion proxies had views less than 4,500. Note that we define high viewership for Dailymotion to be almost 7.6 times lower than that of YouTube. This is to account for the difference in the traffic received by these sites as mentioned in Table 1 (YouTube receives almost seven times more views per month than Dailymotion). For Metacafe, we do not give such a plot as all the videos are available on a CDN server, e.g., Akamai and Level3 Networks, irrespective of popularity. Similar trends were observed for our traces from other locations.

Our results suggest that both YouTube and Dailymotion follow a *selective push* model to allocate more resources to popular content so as to improve the typical end-user experience. In contrast, Metacafe follows a *blind push* model as a file is pushed and replicated over the CDN as soon as it is uploaded. The blind push model is what enables Metacafe to have the lowest service delay globally. The selective push model reduces the degree of replication of the content, but may lead to higher service delays. This result highlights the *classical space-time (degree of replication-service delay) tradeoff, and shows how the popularity of the content decides the extent to which it is pushed.*

4.4 Content Pushing: Age Model

Age is defined as the time elapsed since a video was published on the service website. For analyzing how the services adapt according to the age of the video, we analyzed the "most recent" category of videos. Out of the 7,775 YouTube videos that we crawled in this category from Purdue, we ob-

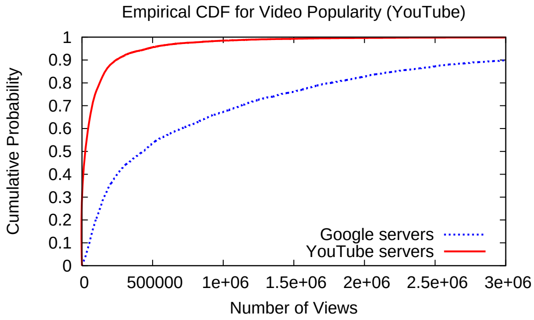

(a) YouTube

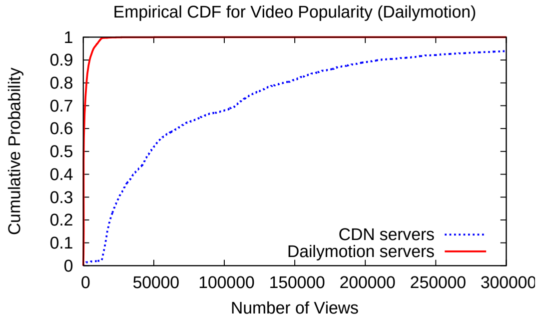

(b) Dailymotion

Figure 5: Distribution of videos according to popularity.

served that 7,681 were available on YouTube servers and 94 on Google servers. Further, 88% videos on YouTube and 73% on Google in this category were uploaded within the last seven days. Another important observation about the recently uploaded videos which were available on Google servers is that most of them received a high number of views soon after they were published. Similar trends are observed for Dailymotion. This result indicates that *although age is an important video characteristic, the services currently give more weight to video popularity in their content pushing models.*

5. CONCLUSIONS

We have presented a comparison of three popular video sharing services in Web 2.0 (YouTube, Dailymotion, and Metacafe), utilizing the PlanetLab experimentation platform. We have inferred the content delivery frameworks that these services use, as well as how the meta information associated with a video impacts its storage and distribution. Our results show how content providers and distributors employ differentiated distribution services by leveraging the meta-information of the content. We plan to investigate patterns introduced by local popularity of content, and incorporate new performance metrics such as jitter. We believe that our traces and results will be useful to researchers and developers wishing to gain insight into the structure of these content delivery services. This will enable enhancing the end-user experience of today's user, and serving high definition video content in the future.

6. REFERENCES

[1] http://www.cs.purdue.edu/homes/msaxena/datasets

[2] http://www.seomoz.org

[3] Gill, P., Arlitt, M., Li, Z., Mahanti, A.: YouTube traffic characterization: A view from the edge. In: Proc. of IMC, San Diego, CA, USA (October 2007)

[4] Wang, Y., Claypool, M., Zuo, Z.: An empirical study of RealVideo performance across the Internet. In: Proc. of SIGCOMM IMW, San Francisco, USA (November 2001)

[5] Sripanidkulchai, K., Maggs, B., Zhang, H.: An analysis of live streaming workloads on the Internet. In: Proc. of SIGCOMM IMC, Italy (October 2004)

[6] Arlitt, M., Jin, T.: Workload characterization of the 1998 world cup web site. IEEE Network **14**(3) (2000) 30–37

[7] Arlitt, M.F., Williamson, C.L.: Internet web servers: workload characterization and performance implications. IEEE/ACM Trans. Netw. **5**(5) (1997) 631–645

[8] Saroiu, S., Gummadi, K., Dunn, R., Gribble, S., Levy, H.: An analysis of Internet content delivery systems. In: Proc. of OSDI, Boston, USA (December 2002)

[9] Krishnamurthy, B., Wills, C., Zhang, Y.: On the use and performance of content distribution networks. In: Proc. of SIGCOMM IMW, New York, NY, USA, ACM Press (2001) 169–182

[10] Cheng, B., Liu, X., Zhang, Z., Jin, H.: A measurement study of a Peer-to-Peer video-on-demand system. In: Proc. of IPTPS, WA, USA (2007)

[11] Huang, Y., Xiao, Z., Chen, Y.F., Jana, R., Rabinovich, M., Wei, B.: When is P2P technology beneficial to IPTV services? In: Proc. of NOSSDAV, IL, USA (June 2007)

[12] Cha, M., Kwak, H., Rodriguez, P., Ahn, Y.Y.: I Tube, You Tube, Everybody Tubes: Analyzing the world's largest user generated content video system. In: Proc. of IMC, San Diego, CA, USA (October 2007)

[13] Cheng, X., Dale, C., Liu, J.: Understanding the characteristics of internet short video sharing: YouTube as a case study. Technical report, Cornell University (Jul 2007)

[14] Jink, M., Suh, K., Gu, Y., Kurose, J.F.: Watch global cache local: YouTube network traces at a campus network - measurements and implications. In: Proc. of ACM MMCN, San Jose, USA (Jan 2008)

[15] Hu, N., Li, L.E., Mao, Z.M., Steenkiste, P., Wang, J.: A measurement study of internet bottlenecks. In: Proc. of INFOCOM, Miami, USA (March 2005)

[16] Spring, N., Peterson, L., Bavier, A., Pai, V.: Using PlanetLab for network research: myths, realities, and best practices. In: SIGOPS Oper. Syst. Rev. (2006)

[17] Banerjee, S., Griffin, T., Pias, M.: The interdomain connectivity of PlanetLab nodes. In: Proc. of PAM Workshop, France (April 2004)

[18] Pucha, H., Hu, Y.C., Mao, Z.M.: On the Impact of Research Network Based Testbeds on Wide-area Experiments. In: Proc. of IMC, Rio de Janeiro, Brazil (October 2006)

[19] Moore, D., Periakaruppan, R., Donohoe, J., k claffy: Where in the world is netgeo.caida.org? In: Proc. of INET, Japan (Jul 2000)