

Viewing YouTube from a Metropolitan Area: What do Users Accessing from Residential ISPs Experience?

Parikshit Juluri[†], Louis Plissonneau[‡], Yong Zeng[†], Deep Medhi[†]
[†]University of Missouri–Kansas City, USA [‡]Orange Labs, France
{pjuluri, zengy, dmedhi}@umkc.edu, louis.plissonneau@orange.com

Abstract— In this work, we perform a controlled study on the perceived experience of viewing YouTube videos as observed from the end users’ point of view through their residential ISPs in a metropolitan area. This study is conducted using our tool *Pytomo*, which we developed to emulate the end users’ experience of viewing YouTube videos. *Pytomo* crawls and downloads YouTube videos to collect a number of measures, including information about the YouTube servers that are delivering them. This open-source tool was provided to a group of volunteers located in the Kansas City metropolitan area. These volunteers, who use different residential ISPs to access the Internet, were instructed to synchronously run the tool to collect the measurement data. Based on the data collected over specific time windows (separated by three months), we observed that there is a noticeable difference in the quality of experience depending on the residential ISPs. Furthermore, the content distribution policies for YouTube, for different residential ISPs vary and the round trip time is not the primary factor for choosing video servers.

I. INTRODUCTION

Video streaming services have become the most dominant traffic on the Internet [1]. *Progressive download* videos that can be embedded into any web-based applications (websites, blogs, and social networks) that enables users to watch videos while simultaneously downloading them. With *progressive download* videos, users expect videos to start playing almost instantaneously. Any delay or degradation is perceived instantaneously by the user, thus relating to her QoE.

One of the most popular video streaming services is YouTube. While there have been a number of studies on YouTube (see Section II), few work on precisely understanding the experience as seen by end users. Undoubtedly, there is enormous difficulty in trying to collect data on end users’ experience. First, there is a lack of tools to conduct such data collection. Secondly, there is the difficulty of finding sufficient volunteers who are willing to participate in such a study. In particular, we are interested in conducting a study based on residential users connected to comparable DSL and cable-based Internet providers as opposed to users connected from academic institutions or corporate environments. For users or computers that access YouTube from academic institutions or corporate locations, it is possible to run bots that could be easily monitored on several computers (in labs or testbeds) to collect data. Also, academic institutions or corporations have high speed access; hence, users having access from such

environments will notice quite a different experience than residential users. However, it is difficult to obtain a whole set of machines connected to residential ISPs in order to run the bots without actual user participation.

In order to study the perceived user experience we needed to collect the QoE metrics from the end users. Thus, to be able to collect the information associated with viewing YouTube by the end users, we recently developed the tool, *Pytomo* [2]. This tool *emulates* a user watching YouTube videos and collects the QoE metrics that can help analyze the perceived experience of a user. Apart from these metrics, we also collected information about video servers that were accessed.

The next important issue was to identify volunteers to collect data using *Pytomo*. We found that it is rather difficult to do a large-scale data collection by asking volunteers accessing YouTube from their residences, especially in a systematic manner. To our knowledge, this is the first work with actual users. For this work, we have collected measurement data from volunteer users located in the Kansas City metropolitan area. The users accessed YouTube using our tool from three different residential ISPs. This controlled study focused on the following:

1. to compare the quality of experience of the end users and YouTube distribution policies when accessing YouTube videos from the same geographic location
2. to collect the data for videos being accessed from actual residential users rather than from a controlled lab setting
3. to study the difference in the YouTube video distribution policies between two time slots

This paper is organized as follows. In Section II, we present the related work and discuss our approach. We then present our tool in Section III and discuss the different metrics that we collected. In Section IV, we discuss the data collection process and later present results and analysis in Section V. We then conclude with a summary in Section VI.

II. RELATED WORK

Several studies [3], [4] have crawled videos of websites such as YouTube to characterize the videos based on their meta-data. In [5], Gill *et al.* captured YouTube video requests collected from a campus network. These works studied the access patterns, distribution of the videos such as lengths and

types, evolution of the popularity of the videos, and statistical properties of the requests.

Reverse engineering techniques have been used to study the content distribution network (CDN) architecture and the video distribution system of YouTube. In [6] and [7], an active measurement infrastructure based on PlanetLab nodes was used to determine the DNS namespace used by YouTube servers. They collected information about the location of the video servers and cache servers and estimated the number of servers observed in different locations. In [8], the authors studied traces from the edge of five networks to understand factors that influenced the video server selection. Our approach differs from these works as we considered only residential users. In [9] a longitudinal study was undertaken to investigate the evolution of YouTube infrastructure over four years.

In [10], Finamore *et al.* packet traces were obtained to study the user experiences from observing the user interaction patterns. In [11], a crowd sourcing approach was used to obtain the QoE metrics of the YouTube videos, which requires the users to watch the videos and provide their feedback.

The set of works that is more closely related to the current work uses active analysis, where the performance metrics are collected while downloading videos. In [12], service delay was calculated while downloading chunks of 1 MB data from YouTube, Dailymotion, and Metacafe. They studied the service delay by obtaining the data in chunks of 50 Kbytes, and assumed that each chunk can play the same amount of video. However, we found that this assumption does not hold for many videos. In *Pytomo*, we parsed the FLV metadata to determine the precise play rate of the video, which gave a more accurate QoE measurement. In [13], user-side instrumentation was used to measure the quality metrics such as join time, buffering ratio, average bit rate, and rate of buffering events for video-on-demand and live content services. In [14], Staehle *et al.* presented a Firefox plugin that generated the QoE metrics to analyze YouTube videos; however, this approach required users to watch YouTube videos.

III. PYTOMO: THE TOOL

Pytomo is an open-source, active measurement tool developed using Python. The code is available for download from <http://code.google.com/p/pytomo/> for all platforms. It is designed to progressively download and playback a YouTube video thus emulating the progressive video playback as experienced by a user. This tool can be used to study both the end users' QoE and the YouTube CDN servers involved in the video delivery.

When a user accesses a YouTube video, there are three major components that are downloaded: (1) the basic HTML page of the video, (2) the YouTube API that is embedded in the webpage, and (3) the progressive download of the video itself. We collect video server information and QoE metrics by performing the following actions on multiple videos:

- **Select initial set of videos:** *Pytomo* begins by selecting an initial set of videos (*VideoList*) to start a crawl. It selects the most popular videos world-wide within a configurable

time-frame (day, week, month, or all-time) at the time of accessing YouTube. These videos are used as the initial seed for each crawl.

- **Obtain video URL:** For each video in the list, *Pytomo* parses the HTML page to obtain the URL of the video that identifies the video server hosting the video file. The IP address of the video server is resolved from the URL.
- **Obtain latency metrics:** The latency to the video server is measured using the RTT metrics from ping.
- **Obtain meta-data of video:** Using the URI of the video, *Pytomo* downloads it from the video server. We observed that a request may be redirected to another server. In this case, the tool goes back to the beginning and retraces all the earlier steps to ensure that the latency metrics correspond to the new server. The meta-data consists of the following information: *VideoDuration*, *VideoLength*, and *EncodingRate*.
- **Obtain playback statistics of video:** Using an emulation mechanism (to be described next), *Pytomo* progressively downloads the video for a certain duration, that we refer to as *VideoDurationDownloaded*, which is a configurable parameter. The emulation allows us to collect data on two different time trackers. The first time tracker, *Download-TimeTracker*, keeps track of the download of the video and follows the absolute download time. The second one, *PlaybackTimeTracker*, keeps track of the playback time of the video to emulate the progressive download. Note that the video playback rate is not dependent on the encoding rate of the video that varies with every frame. Thus, we parse audio and video tags of the frame to precisely calculate the playback rate. The playback of the video goes through the following stages:

- 1) **Initial buffering stage:** *Progressive download* allows initial buffering of a video before it starts the video playback to minimize the effect of jitter(interruptions) in viewing a video. Using reverse engineering, we measured the initial buffer size to be the amount of data needed to play 2 sec. of a video.¹ This value is stored as *InitialData* and the data rate at which this buffer is filled is saved as *InitialRate*. From *InitialData* and *InitialRate*, we calculate *StartTime*, which is the time taken for the initial buffer of the video to be filled and the video to start playing.
- 2) **Playback stage:** Once the initial buffer is filled, the video playback begins. *PlaybackTimeTracker* keeps track of the playback time of the video. If the number of bytes of a video in the buffer falls below the *BufferingThreshold*, we interpret this as an interruption in the playback and *PlaybackTimeTracker* is paused. The number of times a video is interrupted during playback is registered as *DownloadInterruptions* and the total time the video is interrupted is stored as *BufferingDuration*. After the video is

¹This is of variable data length and varies for every video

progressively downloaded for *VideoDurationDownloaded*, we measure *PlaybackDuration* and the total bytes of the video downloaded (*DownloadBytes*).

- **Prepare for next round of crawls:** From the related video of the current video, *Pytomo* then selects a random subset of videos that have not already been crawled and adds them to the *VideoList*. By using a small number of related videos (default=2), we ensure that the crawl is made broader and not limited to only the popular videos.

The above steps are continued until the tool has crawled the predefined maximum number of videos (*MaxNumberURLS*).

IV. DATA COLLECTION

We identified volunteers located throughout the Kansas City metropolitan area to run the *Pytomo* tool on their local machines through their residential ISPs. Based on our preliminary assessment, we found that the local time between 8:00 p.m. and 10:00 p.m. is when residential use of the Internet is most common. Hence, volunteers were asked to simultaneously run the tool during this time on specific evenings so that all of them were accessing YouTube at the same time.

The first collection was performed by 19 users on December 8, 2011. A total of 1,260 videos were downloaded from 442 distinct video servers. The second collection was performed on the evenings of March 14 and 15, 2012 by 32 different users. This collection crawled 2,390 videos from 988 distinct video servers. For both sets of data collection, *VideoDurationDownloaded* for each video was configured to 30 sec. *Pytomo* crawled on average over 60 videos from each volunteer's computer. The volunteers were connected to three different ISP's, with 5, 9, 4 users in the December crawl and 8, 19, 5 users in the March crawl who were connected to ISP-1, ISP-2, and ISP-3 respectively. Based on preliminary tests, a small time gap was introduced before moving from viewing one video to the next so that YouTube's servers did not think that it was a bot crawling their servers and did not block access for any requests from each user's computer. Users were located in three residential ISPs of which ISP-1 and ISP-3 were cable-based providers while ISP-2 was a DSL-based provider.

While our sample size is relatively small, we present confidence interval on the measures and conducted two-sample *t*-tests to determine the significance of our observation. Based on the *p*-values computed from *t*-tests, we found the difference between the December 2011 and the March 2012 crawls to be highly statistically significant.

V. RESULTS

In this section, we present our analysis and observations on the QoE of the end users and how the video server selection dynamics is related to ISPs. We analyze how the perceived QoE of the users varies depending on the residential ISP that the users are connected to.

We first explain the video server selection and the server naming convention used by YouTube to serve videos as this impacts QoE. When a user accesses a YouTube video, the browser first communicates with the YouTube server




Fig. 1. StartTime (December)




Fig. 2. StartTime (March)

to download the HTML page. The HTML page consists of an embedded YouTube player and a link to the video server hosting the actual video. The YouTube API connects to the video server in order to download the video. These steps are followed by our tool, *Pytomo*. On checking the data collected by *Pytomo*, we find that the actual video server URL was typically of the form: `http://o-o.preferred.SERVER_CODE.v[1-24].lscache[1-8].c.youtube.com`, which identifies the domain name of the actual server to be in the form `o-o.preferred.SERVER_CODE.v[1-24].lscache[1-8].c.youtube.com`. Here, `SERVER_CODE` identifies a server ID, behind which, is likely, a cluster of video servers in a data center environment. The naming convention of `SERVER_CODE` seems to indicate the geographic location of video server clusters, that are commonly named by including IATA 3-letter airport codes usually at the beginning; for example, the IATA airport code `dfw` refers to Dallas-Forth Worth, `iad` refers to Washington, DC, `ord` refers to Chicago, and so on (see Table I); this was also noted in [6]. While these are certainly logical domain names, the naming convention seems to indicate that they are located in the geographic region identified via the IATA airport code. The `SERVER_CODE` also gives us an indication of where or how far the video servers are from the users.

A. Analysis of Users' QoE

The factors that we considered to compare the QoE of the users were: *EncodingRate*, *StartTime*, *InitialRate*, *Download-Interruptions*, and *VideoDurationDownloaded*.

EncodingRate: The encoding rate represents the quality of the video and is the average amount of data required to playback one second of the video. We found that the quality of videos downloaded by users of different ISPs were similar. Since the video format selection is dependent on the type of the machine and YouTube's estimation of the data rate to the user, we can say that the machines and the networks used by the users are homogenous as perceived by YouTube. This confirms that any variation in the QoE dependent on the network or video server performance and not the video quality.

StartTime and InitialRate: Table I shows the average start time (with confidence interval) observed for users accessing YouTube from different ISPs when video servers are located

TABLE I
ACCESS PATTERNS FOR THE VIDEO-SERVER FROM 3 ISPs WITH QoE METRICS (\pm DENOTES CONFIDENCE INTERVAL)

Server code	December 2011					March 2012				
	No. of Samples	Playback Start Time (sec)	% Videos Interrupted	RTT (ms)	Download Rate (Mbps)	No. of Samples	Playback Start Time (sec)	% Videos Interrupted	RTT (ms)	Download Rate (Mbps)
ISP-1										
all (ISP-1)	654	1.18 \pm 0.09	10.24	160.39 \pm 9.05	1.35 \pm 0.08	1468	0.49 \pm 0.05	1.90	68.89 \pm 5.61	2.60 \pm 0.20
dfw06g01	41	1.45 \pm 0.71	9.75	100.63 \pm 19.67	1.30 \pm 1.16	-	-	-	-	-
iad09g05	613	1.15 \pm 0.09	10.29	164.43 \pm 9.48	1.35 \pm 0.08	213	0.65 \pm 0.09	2.81	104.92 \pm 22.40	1.82 \pm 0.22
xo-ord1	-	-	-	-	-	1213	0.44 \pm 0.05	0.99	62.83 \pm 4.68	2.82 \pm 0.24
lax04s12	-	-	-	-	-	20	0.97 \pm 0.16	0.00	80.44 \pm 1.52	1.69 \pm 0.46
ord12s05	-	-	-	-	-	22	1.44 \pm 0.55	45.45	43.84 \pm 1.10	0.55 \pm 0.13
ISP-2										
all (ISP-2)	347	0.77 \pm 0.07	1.80	85.33 \pm 2.64	1.50 \pm 0.08	602	1.23 \pm 0.07	16.28	65.24 \pm 2.52	1.78 \pm 0.24
dfw06g01	175	0.89 \pm 0.09	2.85	90.59 \pm 6.26	1.35 \pm 0.08	-	-	-	-	-
iad09s12	6	1.50 \pm 0.31	0.00	77.15 \pm 20.78	1.14 \pm 0.21	-	-	-	-	-
lga15s20	36	0.62 \pm 0.45	5.56	72.58 \pm 9.46	1.76 \pm 0.54	-	-	-	-	-
dfw06s10	3	0.85 \pm 0.74	0.00	35.33 \pm 5.74	1.58 \pm 2.37	-	-	-	-	-
dfw06s08	-	-	-	-	-	46	1.42 \pm 0.20	4.34	63.46 \pm 6.13	0.96 \pm 0.06
mia05s05	65	0.42 \pm 0.06	0.00	79.07 \pm 15.48	1.78 \pm 0.23	93	0.67 \pm 0.67	1.07	69.61 \pm 3.41	2.52 \pm 0.55
sjc07s11	62	0.83 \pm 0.10	0.00	87.65 \pm 6.90	1.50 \pm 0.18	6	2.71 \pm 2.63	33.33	122.15 \pm 34.63	1.06 \pm 0.56
sjc07s15	-	-	-	-	-	147	2.91 \pm 0.54	43.54	99.29 \pm 7.08	0.67 \pm 0.07
atl14s01	-	-	-	-	-	88	0.34 \pm 0.05	0.00	57.65 \pm 2.39	2.77 \pm 0.64
atl05s01	-	-	-	-	-	28	0.34 \pm 0.05	0.00	57.84 \pm 1.25	2.42 \pm 1.06
ord12s05	-	-	-	-	-	44	1.30 \pm 0.34	54.54	44.46 \pm 1.06	0.50 \pm 0.04
ord12s06	-	-	-	-	-	109	0.52 \pm 0.12	4.35	45.18 \pm 1.57	1.82 \pm 0.26
xo-ord1	-	-	-	-	-	41	0.34 \pm 0.05	0.00	66.54 \pm 16.37	3.86 \pm 1.36
ISP-3										
all (ISP-3)	259	0.78 \pm 0.07	0.77	78.27 \pm 2.64	1.68 \pm 0.17	320	0.63 \pm 0.05	0.63	109.657.05	1.84 \pm 0.20
ISP3-dfw1	259	0.78 \pm 0.06	0.77	78.27 \pm 2.65	1.68 \pm 0.17	-	-	-	-	-
ISP3-mial	-	-	-	-	-	132	0.54 \pm 0.05	1.53	95.07 \pm 11.72	2.02 \pm 0.35
ISP3-sjc1	-	-	-	-	-	188	0.68 \pm 0.07	0.00	119.90 \pm 8.67	1.73 \pm 0.25

in different geographic areas (more on server locations will be discussed in Section V-C). The CDF of *StartTime*(Figs. 1 & 2) reflects the overall distributions for the two periods. While in most cases, the average start time was less than 1 sec., we found that for ISP-2, there were at least two server locations that resulted in experiencing an average start time of 2 sec or more in the March 2012 crawl. Although the average *StartTimes* for all servers for ISP-1 and ISP-3 had reduced from December 2011 to March 2012, for ISP-2 it had increased. For the most popular servers of these three ISPs, servers used for ISP-3 had significantly faster *StartTimes*.

Interruptions and DownloadRate: From the data collected, we also calculated the percentage of video interruptions experienced, which is shown in Table I. During the March 2012 crawl, videos served for ISP-3 customers had the least number of interruptions, while at the same time, the number of interruptions was high for ISP-2 customers. From the table, we can also see how each of the video servers contributed to the interruptions. For example, the new server ID *sjc07s15* (i.e., San Jose, California) that served most videos in the March 2012 crawl for ISP-2 had an unusually high number of interruptions at 43.54%. Not surprisingly, this affected the average download rate from this server ID, which was among the lowest. Also, there was a strong correlation between the percentage of interruptions and the download rate.

While we do not have additional data to correlate whether this anomaly is due to the network or servers, we can make a preliminary inference in the case when multiple server IDs are located in the same city. We note two server IDs located in Chicago: *ord12s05* and *ord12s06*; the former




Fig. 3. Download Interruptions (December)




Fig. 4. Download Interruptions (March)

one is used by both ISP-1 and ISP-2, while the latter one is used only by ISP-2. Since accessing *ord12s05* shows high interruptions (and a low download rate) for both providers, but low interruptions in case of *ord12s06*, it is very likely that *ord12s05* is experiencing high load.

In Figs. 3 and 4, we present severity on interruptions by categorizing the number of interruptions (shown in percentage) into three groups: 1 to 2 (“low”), 3 to 4 (“moderate”), and more than 4 (“high”). We observe that the severity of interruptions was often low, while for a notable number of cases, severity was moderate for ISP-1 in the December crawl and for ISP-2 in the March crawl.

B. Two-Sample t-test

In order to determine if the difference in the QoE between December 2011 and March 2012 crawls was significant, we

TABLE II
 p -VALUE FOR TWO-SAMPLE t -TEST

p -value	ISP-1	ISP-2	ISP-3
StartTime	2.2×10^{-16}	1.18×10^{-8}	0.44
DownloadRate	2.2×10^{-16}	0.52	0.22
Average. RTT	2.2×10^{-16}	1.35×10^{-08}	4.19×10^{-13}

$$H_0: \mu_{Dec} = \mu_{Mar} \text{ vs. } H_{alt}: \mu_{Dec} \neq \mu_{Mar}$$

performed the two-sample t -test between the two datasets to determine the p -values (see Table II). For the t -test the most commonly used significance levels are 0.1, 0.05 and 0.01. We select the significance level as 0.01 because it is the most conservative one. We found that for ISP-1 and ISP-2, p -value $\ll 0.01$ for StartTime. Thus, we can say that StartTime in the March 2012 crawl was highly, statistically and significantly different than in the December 2011 crawl. While for ISP-1, StartTime was less in March 2012, it was higher for ISP-2. There was no statistical difference for ISP-3. These results are consistent with Fig.2.


In regards to the download rate, for ISP-1, the March 2012 crawl was highly statistically significantly different (actually higher) than in the December 2011 crawl. However, there was no statistical difference for the download rate for the other two providers. In regard to the Average RTT, the March 2012 crawl was highly statistically significantly different (actually higher) than in the December 2011 crawl for all the providers.

C. Video Server Selection Dynamics


We first briefly comment on whom the video servers belong to. We analyzed the mapping from the IP addresses of the servers and the autonomous systems (AS) that they were associated with. We found that for the December 2011 data set, there were three ASes 15169, 36040, 43515 – all are owned by YouTube or its parent company, Google. For the March 2012 data set, the same three ASes showed up, except that we found a non-YouTube autonomous system (AS 2828) to respond to about 2% of the requests. This verifies that CDN is largely owned by YouTube as reported in [8], except that YouTube still seems to have some partnerships with others for some content distribution.

1) *Video Server Selection based on ISP:* We observed that video servers that were accessed varied significantly between users of different ISPs located in the same geographical location. We will discuss our results by dividing the set of videos into two different groups: the top ten most popular videos for the week and the rest. The top ten videos were accessed by *all* users irrespective of all the providers.

Fig. 5(a) shows access patterns for the ten most popular videos of the week in the December 2011 crawl. The nodes on the left indicate the ISPs from where users are gaining access while each small dot on the right indicates a server cluster that is organized by server IDs. Here we see that except for the server ID dfw06g01 that was accessed by users in ISP-1 and ISP-2, no other video servers were accessed by users of more than one ISP. The observations for the March 2012 crawl are shown in Fig. 5(b). We noted a similar pattern while the common set of server IDs was found to be different and the access spread out to a larger number of video servers.



(a) December 2011 Crawl



(b) March 2012 Crawl

Fig. 5. Server access pattern for top ten videos

Based on these observations, we can infer that even when accessed from the same geographic area, the selection of the video servers for the most popular videos was dependent on the ISP and not on the geographic location of the users. If we consider the access patterns for the top ten videos in March 2012, ISP-1 and ISP-2 were served from two common video servers (`xo-ord1`, `ord12s05`) for some videos while users in ISP-3 did not access any common video servers. Furthermore, by the presence of ISP-3's name in the `SERVER_CODE` of the ISP (not shown here to maintain anonymity) and distinct IP addresses for these servers, which were not observed to be associated with other ISPs, we inferred that these video servers are dedicated to serve requests from users located in this ISP.

When all videos were considered, we found that the server allocation patterns were similar to the ones found for the top ten videos. From Table. I (March 2012 crawl), we see that out of all video servers that were accessed by users of ISP-1 and ISP-2, only two are common (`xo-ord1`, `ord12s05`) (similarly, for the December 2011 crawl). Users of ISP-3 are served by a dedicated set of video servers.

From the above observations, it is safe to conclude that video servers at YouTube are divided into cluster groups with each cluster dedicated to different ISPs, or shared in some cases between a group of ISPs. While we can come to this conclusion, the question remains on how providers and YouTube worked together to invoke preferred video server selection. We believe that this was done using DNS. For example, Bind version 9 of DNS introduced a feature called Views [15]. This can be used to channelize a request from a source IP address or a group of IP prefixes (that resides in a particular ISP) to resolve to a particular domain name and IP address. In other words, there seems to be an agreement between YouTube and large ISPs so that requests for videos from end users from a particular ISP can be directed to

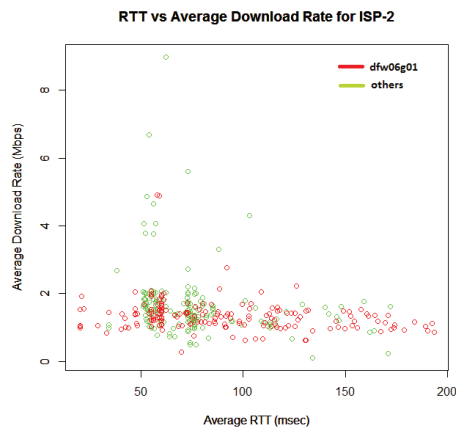


Fig. 6. RTT (msec) and Download Rates (Mbps) Comparison for the December 2012 Crawls for ISP-2

a particular server or a set of specific servers. While [8] discussed that selections are based on DNS, their work does not identify that there are potential policies in place between YouTube and ISPs to make this happen. To our knowledge, we are the first to point out that policies may be in place between YouTube and large ISPs to ensure that certain servers are preferred based on the source IP addresses of an ISP.

2) *RTT and the Location of Video Servers*: We observed that the value of the average RTT does not play a significant role in the selection of the video server. For the three different ISPs, servers that were serving the most number of videos are not necessarily the ones that had the lowest RTT. For ISP-1 and ISP-2 in the December crawl, we see that the most popular server clusters `iad09s12` and `dfw06g01` had the highest average RTT values. Also, servers that had the lowest average RTT did not have the best average data rates.

In Fig. 6, we show the scatter plot of RTT and the download rates for all the videos downloaded from the most popular server `dfw06g01`, and other servers by all users in the December 2012 crawl for ISP-2. From this plot, we see that the RTT to the most popular server was not always the lowest. In the March crawl, a similar observation was made for ISP-1, ISP-2, and ISP-3 where the popular clusters were the ones with a relatively higher RTT. Based on this observation, we believe that users were predominantly directed to a specific server cluster depending on the originating provider, which was not necessarily based on the average RTT.

Our observation contradicts the observation in [8], where the authors observed that the server with the lowest RTT values was the preferred one. The previous work considered RTT from the edge of the network, *not* from the end users' computers. Our measurement was conducted from the end users' perspective and confirms the findings in [16]. Thus, it is possible that the network segment from the end users to the edge of the network contributed to the overall RTT. Secondly, the edge of the network to the video server is not a good indicator for how end users would perceive RTTs.

3) *Change in Video Server Dynamics*: Since the two crawls were performed with a gap of three months, we were able to

study the change in the YouTube distribution architecture over this time span for all three ISPs considered (see Table I).

For ISP-1, there were two video server clusters in the December 2011 crawl; of these two, only one of them was found in the March 2012 crawl along with three new ones. Similarly, for ISP-2, there were six video server clusters in the December 2011 crawl; three of them were found in the March 2012 crawl along with six new ones. This suggests that video servers allocated to the same ISP also change over a short period of time (three months).

Note further that users of ISP-3 were always served by dedicated servers. Servers observed in March 2012 were found to be different compared to the ones in December 2011. Based on the location `SERVER_CODE` we can also infer that the locations of these servers were different. From this, we infer that YouTube has a dedicated set of servers in multiple locations for ISP-3 customers. These servers, at any given time, were exclusive to this provider irrespective of the popularity of the videos. This substantiates the observation we made in the previous section on policies being in place between YouTube and its providers.

D. Limitation of the Study

Recall that our study is limited to a particular geographic area with users accessing from residential providers in an evening time frame, and that this is an observational study and the data collected is a voluntary and convenient sample. Clearly, observations may differ from different geographic areas. Secondly, we limited our data collection to first 30 sec of each video; observations may differ when the entire video is considered. Our crawl mechanism limits the number of videos crawled, allowing us to observe only a representative set of server locations – it is not meant to capture the entire CDN architecture of YouTube.

VI. SUMMARY

In this work, we studied the perceived QoE of viewing YouTube as would be experienced by end users in the Kansas City metropolitan area from three different residential providers. Metrics were collected as volunteer users invoked our tool *Pytomo* from their personal computers. By analyzing the data, we observed that there is a noticeable difference in the perceived QoE depending on the residential ISPs, which can also vary over a short window of time (three months).

We observed that video server selection is not based on geographic location. However, YouTube video distribution policies differ depending on residential providers, which seem to be based on agreements between providers and YouTube. This causes users to have different experiences depending on their residential ISPs. With the availability and selection of multiple server locations, the end users' experience can also be server-dependent. We observed that the set of servers varied noticeably within the span of three months for all providers. The perceived QoE with ISP-3 remains the same over the span of three months while it improved with ISP-1.

ACKNOWLEDGEMENT

We thank the student volunteers at UMKC who took part in collecting the data for this study.

REFERENCES

- [1] Cisco, "Cisco visual networking index: Forecast and methodology, 2010-2015." http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360_ns827_Networking_Solutions_White_Paper.html. Accessed March 2, 2012.
- [2] P. Juluri, L. Plissonneau, and D. Medhi, "Pytomo: A Tool for Analyzing Playback Quality of YouTube Videos," in *Proc. of the 23rd International Teletraffic Congress*, pp. 304–305, September 2011. (poster paper).
- [3] M. Cha, H. Kwak, P. Rodriguez, Y.-Y. Ahn, and S. Moon, "I tube, you tube, everybody tubes: Analyzing the world's largest user generated content video system," in *Proc. of the 2007 ACM Internet Measurement Conference*, pp. 1–14, 2007.
- [4] X. Cheng, C. Dale, and J. Liu, "Understanding the characteristics of Internet short video sharing: YouTube as a case study," *CoRR*, 2007.
- [5] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "YouTube traffic characterization: A view from the edge," in *Proc. of the 2007 ACM Internet Measurement Conference*, pp. 15–28, 2007.
- [6] V. K. Adhikari, S. Jain, Y. Chen, and Z.-L. Zhang, "Reverse engineering the youtube video delivery cloud," in *IEEE HotMD 2011*, 2011.
- [7] V. K. Adhikari, S. Jain, Y. Chen, and Z.-L. Zhang, "Vivisectioning YouTube: An active measurement study," in *IEEE INFOCOM 2012 Miniconference*, 2012.
- [8] R. Torres, A. Finamore, J. R. Kim, M. Mellia, M. M. Munafò, and S. Rao, "Dissecting video server selection strategies in the youtube cdn," in *Proceedings of the 2011 31st International Conference on Distributed Computing Systems (ICDCS'11)*, pp. 248–257, 2011.
- [9] L. Plissonneau and E. Biersack, "A longitudinal view of HTTP video streaming performance," in *MMSys*, 2012.
- [10] A. Finamore, M. Mellia, M. M. Munafò, R. Torres, and S. G. Rao, "YouTube everywhere: Impact of device and infrastructure synergies on user experience," in *Proc. of 2011 ACM Internet Measurement conference*, pp. 345–360, 2011.
- [11] T. Hossfeld, M. Seufert, M. Hirth, T. Zinner, P. Tran-Gia, and R. Schatz, "Quantification of YouTube QoE via Crowdsourcing," in *Proc. of the 2011 IEEE International Symposium on Multimedia*, pp. 494–499, December 2011.
- [12] M. Saxena, U. Sharan, and S. Fahmy, "Analyzing video services in Web 2.0: A global perspective," in *Proc. of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, pp. 39–44, 2008.
- [13] F. Dobrian, V. Sekar, A. Awan, I. Stoica, D. Joseph, A. Ganjam, J. Zhan, and H. Zhang, "Understanding the impact of video quality on user engagement," in *Proc. of ACM SIGCOMM'2011*, vol. 41, pp. 362–373, August 2011.
- [14] B. Staehle, M. Hirth, F. Wamser, R. Pries, and D. Staehle, "YoMo: A YouTube application comfort monitoring tool," Technical Report 467, Institut für Informatik, March 2010.
- [15] C. Liu, "Views in BIND 9." May 2000, http://www.oreillynet.com/pub/a/oreilly/networking/news/views_0501.html.
- [16] L. Plissonneau, E. Biersack, and P. Juluri, "Analyzing the impact of YouTube delivery policies on the user experience," in *Proc. of 2012 International Teletraffic Congress (ITC'2012)*, (Krakow, Poland), 2012.