




Revealing the Load-Balancing Behavior of YouTube Traffic on Interdomain Links

Ricky K. P. Mok¹ , Vaibhav Bajpai², Amogh Dhamdhere¹,
and K. C. Claffy¹

¹ CAIDA, UCSD, San Diego, USA

{cskpmok, amogh, kc}@caida.org

² Technische Universität München, Munich, Germany

bajpai@in.tum.de

Abstract. For the last decade, YouTube has consistently been a dominant source of traffic on the Internet. To improve the quality of experience (QoE) for YouTube users, broadband access providers and Google apply techniques to load balance the extraordinary volume of web requests and traffic. We use traceroute-based measurement methods to infer these techniques for assigning YouTube requests to specific Google video content caches, including the interconnection links between the access providers and Google. We then use a year of measurements (mid-2016 to mid-2017) collected from SamKnows probes hosted by broadband customers spanning a major ISP in the U.S. and three ISPs in Europe. We investigate two possible causes of different interdomain link usage behavior. We also compare the YouTube video cache hostnames and IPs observed by the probes, and find that the selection of video cache has little impact on BGP selection of interdomain links.

1 Introduction

Over a billion users collectively watch billions of hours of videos every day [25], making Google's YouTube the most popular video streaming web service on the Internet. The tremendous growth in volume of users and video content has occurred in parallel with – and as a key driver of – the development and improvement of broadband infrastructure around the world. Indeed, many consumers have canceled their cable television subscriptions in favor of media services such as YouTube or Netflix available over the Internet. Accompanying this evolution are growing performance expectations of users – that the streaming video quality of experience should match that of cable television, a service historically provided over a dedicated private network infrastructure. In parallel, the evolution of video technologies, such as 8K resolution, 60 frame per second (fps), and High Dynamic Range (HDR), has increased the network bandwidth requirement and has further challenged network provisioning economics.

ISPs can coordinate (contracts) with Google to install Google Global Caches (GGCs) inside their networks, and can also rely on their peering relationships

with Google (AS 15169/AS 36040) to connect users to Google/YouTube front-end servers and video caches inside Google’s Points of Presence (PoPs). Many of these interdomain links are of significant and growing capacity, but they can still experience congestion during peak hours [17] that may induce inflated round-trip delay and packet losses, and thus degrade user QoE.

We report the results of a study that combines interdomain topology measurement and YouTube-specific probing measurements to investigate performance-relevant traffic dynamics of ISPs that do not deploy GGCs. We inferred interdomain router-level topology by executing the `bdrmap` [18] tool on ~50 of CAIDA’s Archipelago (Ark) probes [7]. We used a recently developed end-to-end YouTube performance test [3] that streams a video clip similar to a normal client, and reports information including the hostname and IP address of the YouTube video cache (GGC) streaming the video. The test then immediately performs a `paris-traceroute` [4] toward that IP to capture the forward path information. The test ran on ~100 SamKnows probes [6] for about a year (May 2016 to July 2017) [5]. We selected the SamKnows probes connected to ISPs that did not deploy GGCs internally, but whose interdomain topology to Google was captured by our `bdrmap` measurements. This constraint limited our study to 15 SamKnows probes connected to four major ISPs: 1 in the U.S., 3 in Europe.

Our study had two major goals. The first one was to investigate factors that influence ISP strategies for distributing YouTube traffic flows across different interdomain links. We studied two possible factors – geographic location and time of day. We developed a metric of *link usage probability* to characterize the link usage behavior observed by our probes. Our results revealed that geographic location appeared to influence interdomain link assignment for Comcast users, i.e., proximate users were more likely to use the same set of links to reach a cache. We also found that a German ISP (Kabel Deutschland) showed different link usage behavior during peak vs. off-peak hours; other ISPs did not show such a significant difference. By analyzing the interdomain topology, we also discovered three European ISPs that relied on the YouTube AS (AS 36040) rather than the primary Google AS (AS 15169) to reach YouTube content. Our second goal was to study whether YouTube’s cache selection approach could also determine the choice of interdomain links due to the topological location of the cache. We did not observe such a correspondence; more than half of the video caches we observed used at least two interdomain links. We also discovered that the DNS namespace for YouTube video caches (`*.googlevideo.com`) had a more static hostname-IP mapping than front-end hostnames (e.g., `youtube.com` and `google.com`), which used DNS-based redirection [8]. 90% of video cache hostnames were reported (by the probes) to have the same IP, even if they were resolved by different probes.

Section 2 presents related work on YouTube measurement. Sections 3, 4, and 5 describes our datasets and methodology, reports our findings, and offers conclusions, respectively.

2 Related Work

Previous studies have evaluated the architecture or characteristics of YouTube by actively sending video requests. Pytomo [22] crawled YouTube video clips from residential broadband (volunteer) hosts, and collected YouTube server information including hostname and network throughput. They found that the YouTube cache selection depended on user’s ISP rather than geographical proximity. Adhikari *et al.* [2] dissected the architecture of YouTube by requesting video clips from PlanetLab nodes. To increase the coverage, they exploited various geographically distributed public DNS servers to trigger DNS-based redirection in YouTube front-end servers. Recent studies [8, 10] used the EDNS extension to geolocate Google’s CDN infrastructure. A closely related work by Windisch [24] deployed five monitors in a German ISP and parsed YouTube responses to analyze selection of video caches. These studies did not investigate interdomain link structure, which could impact latency and streaming performance. Our study fills this gap by integrating interdomain topology and end-to-end measurement to understand the ISP’s role in load balancing YouTube traffic.

Others have used passive measurement to study YouTube traffic, including analyzing traffic characteristics of video flows [12, 13] and cache selection mechanisms [23]. Casas *et al.* [9] used a 90-h Tstat trace to contrast YouTube traffic characteristics between fixed-line and mobile users. YouLighter [14] used passive monitoring to learn the structure of YouTube’s CDN and automatically detect changes. Because passive measurement relies on user traffic, it is hard to perform a longitudinal study from the same set of clients to observe changes in load balancing across interdomain links over time.

3 Methodology

We deployed the YouTube test [3] on ~ 100 SamKnows probes connected to dual-stacked networks representing 66 different origin ASes [5]. The probes were mostly within the RIPE (60 probes) and ARIN (29) region, and hosted in home networks (78). The YouTube test ran once per hour for IPv4 and then for IPv6. Each test streamed a popular video from YouTube, and reported the streaming information and performance, including start-up delay, YouTube cache hostname and IP. We then ran `paris-traceroute` [4] with `scamper` [16] toward the cache IP reported by the YouTube test, obtaining forward path and latency measurements. Details of the YouTube tests and SamKnows probe measurements are in [3] and [5], respectively.

To identify which interdomain links (if any) were traversed on the paths from our SamKnows probes to YouTube servers, we first compiled the set of interdomain interconnections of the access network visible from a vantage point (VP) in that network. We used `bdrmap` [18], an algorithm that infers interdomain interconnections of a VP network visible from that VP. In the collection phase, `bdrmap` issues traceroutes from the VP toward every routed BGP prefix, and performs alias resolution from the VP on IP addresses seen in these traceroutes.

In the analysis phase, **bdrmap** uses the collected topology data along with AS-relationship inferences from CAIDA’s AS relationship algorithm [19], and a list of address blocks belonging to IXPs obtained from PeeringDB [21] and PCH [20] to infer interdomain links at the router level. The **bdrmap** algorithm then uses constraints from traceroute paths to infer ownership of each observed router, and identifies the routers on the *near* and *far* side (from the perspective of the VP) of every observed router-level interdomain link. We could not run **bdrmap** from the SamKnows probes, so we used the results of **bdrmap** running on Ark VPs located in the same ASes as the SamKnows probes.

3.1 Identifying Interdomain Links from YouTube Dataset

The first step of identifying interdomain links seen in our YouTube traceroutes is to extract all the interdomain links to the Google ASes (AS 15169/AS 36040) observed by Ark VPs. Each link is represented by a pair of IP addresses indicating the interfaces of the near and far side routers. We used these pairs to match consecutive hops in the traceroutes to YouTube video caches. This approach avoids false inference of links, but could miss some links with the same far side IP but a near side IP that **bdrmap** did not observe, because **bdrmap** and the YouTube traceroutes run from different VPs. Section 4.1 describes why we consider our coverage of interdomain links to be satisfactory.

The next step is to aggregate pairs with the same far side IP, because different VPs in the same network may take different paths before exiting via the same interdomain link; in such cases, they likely observe different addresses (aliases) on the near router. Even though **bdrmap** has performed some IP alias resolution, multiple links may connect the same near and far side routers. We resolve this ambiguity by conducting additional IP alias resolution with MIDAR [15] on these far side IPs. Table 1 shows the number of inferred interconnection links at each stage.

Table 1. Number of identified interdomain links at each stage.

Stages	Number of links
Interdomain links to Google inferred by bdrmap	1,268
Links identified in YouTube traceroutes	468
Aggregated with far side IPs	61
IP alias resolution with MIDAR	45

3.2 Descriptive Statistics

We analyzed data collected from May 17, 2016 to July 4, 2017, which included a gap in data between January 4, 2017 and February 15, 2017 for all probes due to technical problems. The data includes more than 74,000 experiment sessions/traceroute records, collected from 15 SamKnows probes connected to 4

broadband ISPs in the United States and Europe. We only used a small subset of the entire YouTube traceroute dataset in our study, constrained by our needs for: (1) co-located Ark VPs in the same ISP to obtain `bdrmap` coverage, and (2) ISPs without GGC deployment internal to their network. The YouTube test collected more than 3,000 distinct video cache hostnames and IPs. Table 2 shows the details of the combined dataset. We adopt the notation `#XX` to represent SamKnows probes. The number (XX) matches the probe ID in the metadata of the SamKnows probes listed in (<https://goo.gl/E2m22J>).

Table 2. Summary of the combined dataset.

ISP		Comcast	Kabel ^a	Italia ^b	Free
Country		US	DE	IT	FR
No. of SamKnows probes		12	1	1	1
No. of interdomain links with Google		26	5	10	4
No. of observed video caches by	Hostname	2,918	303	183	176
	IP	2,983	300	185	185

The full company name: ^aVodafone Kabel Deutschland;
^bTelecom Italia Sparkle S.p.A.

4 Results


We analyzed load balancing behavior on both the ISP and server side, by characterizing the use of interdomain links and the video cache assignment. These choices are interdependent, since ISPs route YouTube requests according to the IP address of the video cache assigned by YouTube. We attempted to isolate these two behaviors and investigate them separately. We investigated the impact of two factors – geographic location and time of day. We also used hostnames and IP addresses of YouTube caches to estimate the influence of YouTube’s video cache selection mechanism on interdomain paths traversed by YouTube requests.

4.1 Interconnection Between ISPs and Google


Consistent with public data [21], we observed multiple interdomain links connecting ISPs to Google in various locations. Figure 1(a) and (b) are two heatmaps showing the interdomain links used by probes in Comcast and the three European ISPs, respectively. Each row represents a SamKnows probe; changing colors on a row represent changing interdomain links. The YouTube tests and traceroutes execute once per hour, so the time resolution of each cell in a row is 1 h. Gray color indicates no data available. Apart from the blackout period, some probes began probing after the measurement period starts (e.g., #89 and #96) or went offline. White color indicates the probe was online, but we could not identify an interdomain link discovered by `bdrmap` measurement

from the traceroute. For Comcast, which hosts multiple Ark VPs, we identified an interdomain link in 83.4% of traceroutes. For ISP Free (#71) and Italia (#43), we identified an interdomain link in only 40.2% and 77.7% of traceroutes, respectively. The large white portion in #02 after February 2017 was caused by relocation of the probe from a Kabel user to a M-net (another German ISP) user. Ark did not have any VP in the M-net network.

We found that each probe used at least 2 interdomain links throughout the measurement period. Some probes (e.g., #78, #43) observed more than 6 links. Load balancing among links was frequent, reflected by a change in color over time. Although not clearly visible in the heatmap, we observed some monitors cease using a link that other monitors continued to use, suggesting another reason for the switch than a link outage. We observed only one link (light blue color) captured by five monitors (#27, #67, #44, #60, #32) that switched entirely to



(a) Probes in Comcast.



(b) Probes in Kabel(#02), Italia(#43), and Free(#71).

Fig. 1. Observed interdomain links against time. Changing colors represents the switching of interdomain links. (Color figure online)

another link (darker blue) after mid-February 2017. The set of links used by two different monitors differed widely, even in the same ISP. For example, there was no intersection of links between #61 and #44, unlike #44 and #32.

We systematically studied the assignment of interdomain links to probes by computing the probability of observing each link by the probes. We define the *link usage probability*, P_l^b as

$$P_l^b = \frac{n_l^b}{\sum_{i \in \mathbb{L}} n_i^b}, \tag{1}$$

where \mathbb{L} is the set of all 45 interdomain links observed in our data, n_l^b is the number of observations of link l by probe b . Higher values indicate higher probability for the probe to use that link.

Due to space limitation we show results of only six representative probes (including 38 links and all 4 ISPs) in Fig. 2. The x -axis of the figure shows different interdomain links, while the y -axis indicates the link usage probability (log scale). Different color bars distinguish results of the six probes. The gray dotted vertical lines separate links of different ISPs. Four probes in Comcast (#61, #38, #78, #44) showed slight overlap in interdomain link use (e.g., Link ID 2 and 8). Three probes in Comcast (#38, #78, #44) showed comparable probability of using at least 2 links, indicating load balancing behavior. Probes #02, #43, and #71 distribute requests to at most 10 links. To demystify the assignment of links, we examined two possible factors: geographical location and time of day.

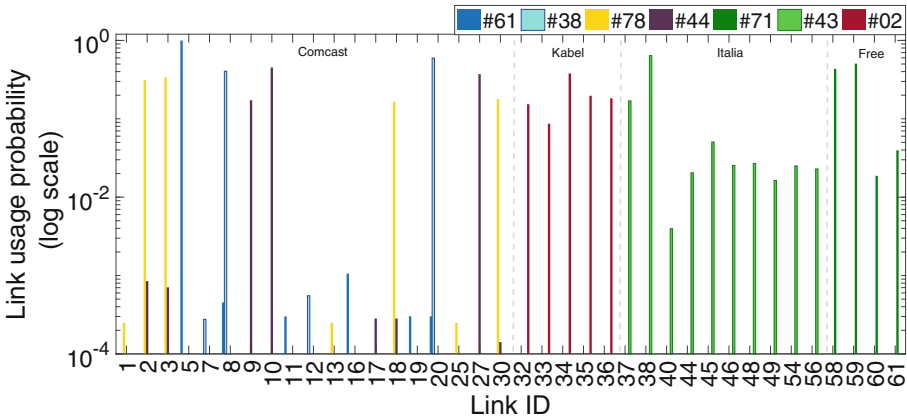


Fig. 2. Link usage probability of 6 probes. Each ISP employs at least two interdomain links to load balance the traffic to video caches. (Color figure online)

Geographic Location. The first factor to study is the relationship between geographic location of probes and the use of interdomain links. Figure 1(a) shows that different probes in Comcast showed similar/dissimilar behavior in terms of link use. We investigated this sharing of interdomain links among probes.

We characterize this behavior by computing a *link usage probability vector*, $\mathbf{P}^b = \langle P_1^b, P_2^b, \dots, P_i^b \rangle, \forall i \in \mathbb{L}$, for each probe. We then performed agglomerative hierarchical clustering in Matlab, and used squared Euclidean distance as a similarity measure between two vectors. We considered only Comcast monitors, because interdomain links will not overlap across ISPs. Figure 3 shows the dendrogram of the resulting five clusters, which reflect the locations of the probes.

The leftmost cluster (red) consists of 6 monitors in the Northeastern U.S. The second cluster (#30) is in the Southeastern U.S. The remaining three clusters are in northern central, southwest, and central areas of the U.S., respectively. This clustering is consistent with the goal of reducing latency of requests by routing them across the nearest interconnection.

Time of Day. Another important factor is time of day, because ISPs or YouTube can employ different load balancing strategies during peak hours. We adopted the “7 p.m. to 11 p.m.” definition of the peak usage hour from the FCC Broadband America Report [11], and recomputed the link usage probability for peak and off-peak hours. The German ISP (Kabel) showed a significant difference in terms of link usage probability in the two time periods. Figure 4 shows the five interdomain links observed by probe #02. During the off-peak hours, the five links were somewhat evenly utilized. In the peak hours, only three of the five links were significantly used. The link usage probability of the three links increased 5% to 15% relative to off-peak hours. For the other ISPs, we did not find significant differences in link usage (not to be confused with utilization!) between peak and off peak hours.

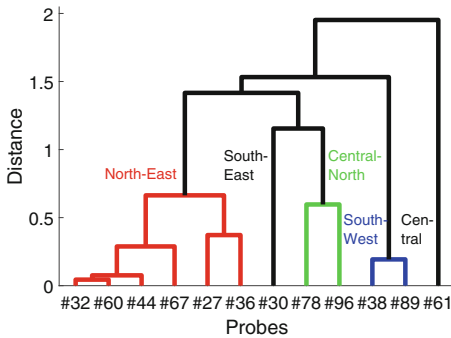


Fig. 3. Dendrogram for hierarchical clustering of Comcast probes. (Color figure online)

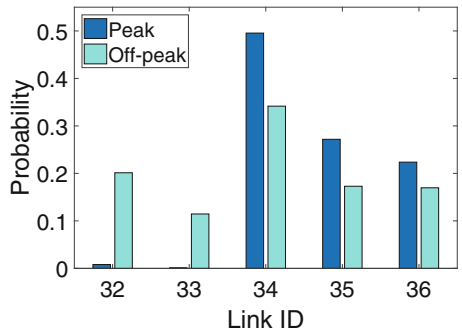


Fig. 4. The link usage probability of Kabel (#02) during peak/off-peak hours.

4.2 Destination Google AS

According to [1], ISPs can establish peering with Google on two ASes—AS 15169 and AS 36040. The former is the most common option and can access all Google services and content, while the latter provides only the most popular content and is not available at all IXPs [1]. Table 3 shows the link usage probability according to the destination AS of the links. Values in brackets are the number of links in the respective categories.

Table 3. Link usage probability (number of interdomain links) to Google ASes.

Google AS	Comcast	Kabel	Italia	Free
15169	0.99 (25)	0.76 (3)	0.16 (6)	0.04 (1)
36040	0.0001 (1)	0.24 (2)	0.84 (4)	0.94 (2)
43515	0	0	0	0.02 (1)

Comcast mostly connects users to Google with AS 15169. For the other three ISPs in Europe, load balancing with AS 36040 is more common. ISP Italia has more interdomain links peering with AS 15169, but accesses YouTube caches mainly using AS 36040. This arrangement could be for historical reasons, because AS 36040 was assigned to YouTube before the merger (Today, the AS name is still ‘YouTube’). For the German ISP Kabel, we found that the links (Link ID 32 and 33) mostly used in the off-peak hours (see Fig. 4) were peering with AS 36040, while the remaining three links were peering with AS 15169.

Interestingly, we found that ISP Free connected users to Google with AS 43515 between Jun 1, 2016 and Aug 17, 2016. Google currently manages this AS for its core network but not for peering purposes [1]. These YouTube test sessions were assigned to video caches with IP prefix (208.117.224.0/19), announced by AS 43515. We believe that the purpose of this AS recently changed. Some video caches were still assigned to AS 43515 during that time period, but now no longer responded to ICMP ping, as other caches did. This example illustrates that ISPs may have different preferences in engineering traffic to and from Google ASes.

4.3 Video Cache Assignment

YouTube mainly employs two techniques to load balance requests, namely DNS-based redirection and HTTP-based redirection. DNS-based redirection assigns users to a front-end server according to the DNS server making the querying [8, 10]. These front-end servers, apart from serving static web elements on youtube.com, are responsible for assigning users to video caches hosted under the domain `*.googlevideo.com`. In some cases, the video content is not available in the assigned video cache (cache miss), Google uses HTTP-based redirection to redirect users to another cache using HTTP response status code 302.

We investigated whether the video caches selected by the front-end server considered the use of interdomain links. Our YouTube measurements captured more than 3,000 hostnames and IPs of video caches. The SamKnows probes resolved these hostnames with their default list of DNS servers, during each YouTube measurement. We found that around 90% of the hostnames mapped to a single IP address, except a special hostname (`redirector.googlevideo.com`) designed for handling cache misses. This result indicated that DNS-based redirection is not common for hostnames of Google’s video caches.

To study the mechanism of video cache selection method, we compared video cache hostnames and IPs between any two probes. In Sect. 4.1 we described how user geographic location appears to influence selection of interdomain link. If Google uses video cache selection to engineer the desired use of specific interdomain links, the front-end servers will likely direct nearby probes to a similar set of caches. Figure 5 depicts the overlapping in video cache hostname/IP mappings for any two monitors, with probes (rows) sorted according to the clustering results in Fig. 3. The lower/upper triangular part of the matrix compares the hostnames/IPs collected by the two probes, respectively. The triangular symmetry is a reflection of the largely one-to-one mapping between IPs and hostnames. From the similarity of the use of interdomain links, we expect nearby probes (e.g., #32, #60, and #44) should share a similar set of video caches (i.e., many over-

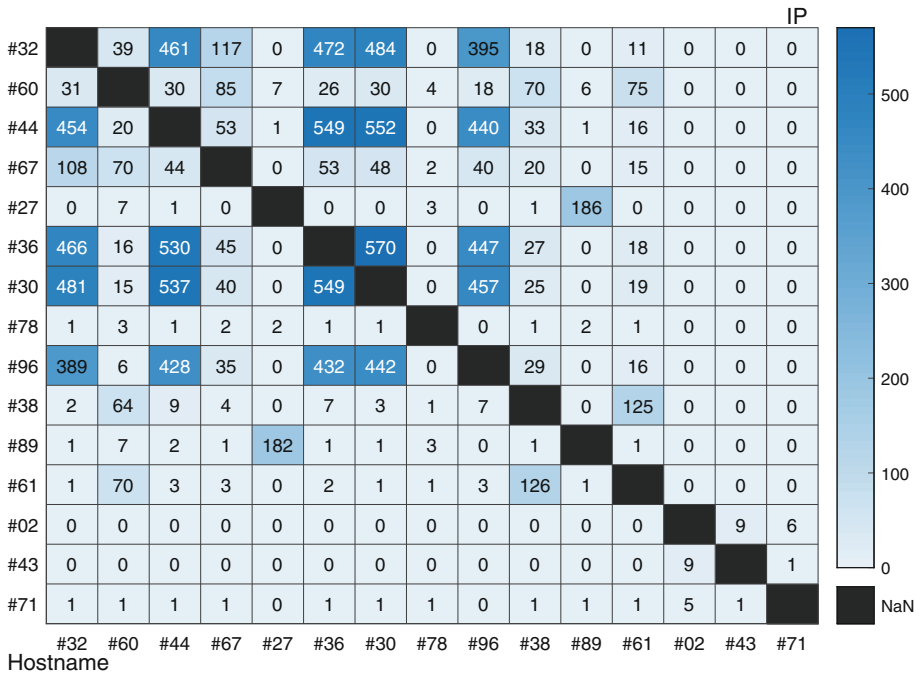


Fig. 5. Number of overlapping Hostnames and IPs across all probes. Probes showed similar interdomain link usage behavior did not show the same degree of similarity in video cache selection.

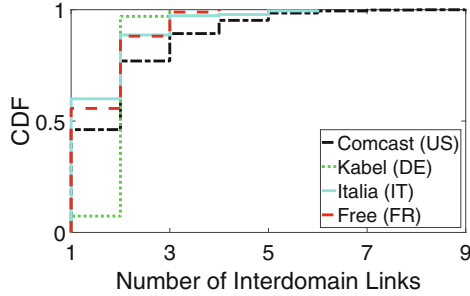


Fig. 6. The CDFs of the number of links used to reach the same video cache IP. Multiple links are used to access the same video caches.

lapping IPs or hostnames). However, the two probes with the highest similarity (#32 and #60) had fewer than 40 overlapping IP/hostname pairs. Surprisingly, probes #32 and #30 had the most such IP/hostname pairs. These two Comcast probes were around 1,500 km apart. There was no overlapping interdomain links among ISPs, but we observed 16 cross-ISP overlapping video cache IPs between Italia (#43) and Free (#71). Given these dissimilar patterns with the use of interdomain links we presented in previous sections, we believe that video cache selection did incorporate any interdomain link preference.

ISPs can also balance YouTube workload by distributing traffic to the same video cache via different interdomain links. In our measurements, around half of the YouTube video cache IPs were accessed with more than one interdomain link (Fig. 6). For Kabel, about 90% of the video caches were reached with at least two different links, suggesting that access ISPs are heavily involved in load balancing traffic to/from YouTube.

5 Conclusion

We used topological measurement and inference and YouTube-specific end-to-end measurement to explore how Google and ISPs perform load balancing of YouTube traffic. By incorporating interdomain link information, we discovered that ISPs play an important role in distributing YouTube traffic across multiple interdomain links that connect ISPs to Google infrastructure. Unsurprisingly, location and time-of-day influence load balancing behaviors. For the server side, our analysis of DNS bindings between hostnames and IPs of video caches suggests that YouTube front-end servers select video caches by controlling hostnames, rather than DNS-redirectation. We further observed that the same video cache can be accessed with multiple interdomain links, and the varied patterns of such links across different access ISPs suggests that ISPs, rather than Google, play a primary role in balancing YouTube request load across their interdomain links toward Google. In the future, we plan to investigate the impact of load balancing behavior on video streaming performance and its correlation to user-reported QoE.

Acknowledgment. This work was partly funded by the European Union’s Horizon 2020 research and innovation programme 2014–2018 under grant agreement No. 644866, Scalable and Secure Infrastructures for Cloud Operations (SSICLOPS), and by U.S. National Science Foundation CNS-1414177. This work represents only the position of the authors, and not of funding agencies.

References

1. Google edge network. <https://peering.google.com/>. Accessed 11 Oct 2017
2. Adhikari, V.K., Jain, S., Chen, Y., Zhang, Z.-L.: Vivisecting YouTube: an active measurement study. In: IEEE INFOCOM (2012)
3. Ahsan, S., Bajpai, V., Ott, J., Schönwälder, J.: Measuring YouTube from dual-stacked hosts. In: PAM (2015)
4. Augustin, B., Teixeira, R., Friedman, T.: Measuring load-balanced paths in the Internet. In: ACM IMC (2007)
5. Bajpai, V., Ahsan, S., Schönwälder, J., Ott, J.: Measuring YouTube content delivery over IPv6. In: ACM SIGCOMM CCR (2017)
6. Bajpai, V., Schönwälder, J.: A survey on internet performance measurement platforms and related standardization efforts. *IEEE Commun. Surv. Tutor.* **17**(3), 1313–1341 (2015)
7. CAIDA: Archipelago (Ark) measurement infrastructure. <http://www.caida.org/projects/ark/>
8. Calder, M., Fan, X., Hu, Z., Katz-Bassett, E., Heidemann, J., Govindan, R.: Mapping the expansion of Google’s serving infrastructure. In: ACM IMC (2013)
9. Casas, P., Fiadino, P., Bar, A., D’Alconzo, A., Finamore, A., Mellia, M.: YouTube all around: characterizing YouTube from mobile and fixed-line network vantage points. In: EuCNC (2014)
10. Fan, X., Katz-Bassett, E., Heidemann, J.: Assessing affinity between users and CDN sites. In: IFIP TMA (2015)
11. FCC: Measuring broadband America fixed broadband report, December 2016. <https://www.fcc.gov/reports-research/reports/measuring-broadband-america/measuring-fixed-broadband-report-2016>. Accessed 15 Oct 2017
12. Finamore, A., Mellia, M., Munafò, M.M., Torres, R., Rao, S.G.: YouTube everywhere: impact of device and infrastructure synergies on user experience. In: ACM IMC (2011)
13. Gill, P., Arlitt, M., Li, Z., Mahanti, A.: Youtube traffic characterization: a view from the edge. In: ACM IMC (2007)
14. Giordano, D., Traverso, S., Grimaudo, L., Mellia, M., Baralis, E., Tongaonkar, A., Saha, S.: YouLighter: a cognitive approach to unveil YouTube CDN and changes. *IEEE Trans. Cognit. Commun. Netw.* **1**(2), 161–174 (2015)
15. Keys, K., Hyun, Y., Luckie, M., Claffy, K.: Internet-scale IPv4 alias resolution with MIDAR. *IEEE/ACM Trans. Netw.* **21**(2), 383–399 (2012)
16. Luckie, M.: Scamper: a scalable and extensible packet prober for active measurement of the Internet. In: ACM IMC (2010)
17. Luckie, M., Dhamdhere, A., Clark, D., Huffaker, B., Claffy, K.C.: Challenges in inferring Internet interdomain congestion. In: ACM IMC (2014)
18. Luckie, M., Dhamdhere, A., Huffaker, B., Clark, D., Claffy, K.C.: bdrmap: inference of borders between IP networks. In: ACM IMC (2016)
19. Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., Claffy, K.C.: AS relationships, customer cones, and validation. In: Proceedings of ACM IMC (2013)

20. Packet Clearing House: Full exchange point dataset (2017). <https://prefix.pch.net/applications/ixpdir/menu.download.php>
21. PeeringDB (2017). <http://www.peeringdb.com>
22. Plissonneau, L., Biersack, E., Juluri, P.: Analyzing the impact of YouTube delivery policies on user experience. In: ITC (2012)
23. Torres, R., Finamore, A., Kim, J.R., Mellia, M., Munafo, M.M., Rao, S.: Dissecting video server selection strategies in the YouTube CDN. In: IEEE ICDCS (2011)
24. Windisch, G.: Analysis of the YouTube server selection behavior observed in a large German ISP network. In: EUNICE (2014)
25. YouTube: Youtube for press. <https://www.youtube.com/intl/en-GB/yt/about/press/>. Accessed 10 Oct 2017