

Analysis of the YouTube Server Selection Behavior Observed in a Large German ISP Network

Gerd Windisch^(✉)

TU Chemnitz, Chemnitz, Germany
gerd.windisch@etit.tu-chemnitz.de

Abstract. This paper presents a detailed study on the YouTube server selection mechanism with respect to its long term behavior. We conducted a measurement campaign from different measurement points within a single German ISP network over a long period of time. Based on the measurement traces, the long-term temporal behavior of the server selection mechanism is investigated. The results reveal that the server selection behavior ranges from changing on a daily basis to being stable for several weeks. Besides analyzing the long term behavior, we also determined the dependency of the selection behavior on the geographic location of the client node. It appears that the server selection behavior mainly depends on the IP address currently assigned to a user connection. Finally the impact of the selection decision on the round trip time (RTT) is analyzed.

1 Introduction

Multimedia content, especially video, is responsible for the majority of IP traffic in today's networks. According to two studies [1, 2] the multimedia platform YouTube alone accounts for almost 30 % of the traffic in European ISP networks. Multimedia content mostly utilizes the HTTP protocol and is usually delivered by means of global content delivery networks (CDNs), like Akamai, Google, Limelight, etc. YouTube itself uses the CDN of Google.

CDNs usually deploy geographically disperse data centers to distribute content and to bring content closer to the user. Therefore, the same content is available at many servers in different locations. In order to utilize their distributed infrastructure, CDNs need an efficient way to select a server for a given content request. Most CDNs use a domain name system (DNS) based mechanism for this task. When a user requests a piece of content, the CDN resolves the host name of the content server to an IP address of a server, based on defined criteria. These criteria could be, for example, the load of the server or the round trip time (RTT). However, the load situation within the network is usually not taken into account. This may lead to increasing network congestion. Therefore, it is important for a network operator to better understand how the respective CDN server selection mechanism works and which parameters determine its behavior.

The main contribution of this paper is the analysis of the long term temporal behavior of the YouTube server selection mechanism and the parameters it depends on.

Since it is expected that the selection behavior depends on the geographic location of the requesting user, we emphasize on measurements in different regions within the ISP network.

The paper is organized as follows: Sect. 2 provides an overview of related work. In Sect. 3 our measurement approach is described and Sect. 4 presents the results of our study. Section 5 concludes the paper.

2 Related Work

Adhikari et al. published several papers (e.g. [3]) about reverse engineering the YouTube content delivery platform. Their focus is on analyzing the global footprint of YouTube. They utilized a distributed measurement approach using PlanetLab nodes located at 271 sites worldwide. Key findings are the discovery of how the Video ID is mapped to server host names, and the redirection hierarchy used by YouTube. A second study conducted by the same research group focused on the traffic exchanged between the YouTube infrastructure and a tier-1 ISP [4]. They discovered that load is balanced proportionally to the data center size and that it is not influenced by the proximity to the user. This work had been conducted in 2008, before YouTube was acquired by Google.

Torres et al. [5] analyzed the server selection strategies of YouTube using packet traces captured at five measurement points in Europe and in the US. They discovered that most of the time a user is served by the content server location with the lowest RTT. They identified that at least 10 % of the decisions are influenced by other factors. They claim that the content server load and/or diurnal effects have an impact on the decision as well.

Plissonneau et al. [6] analyzed the impact of the server selection mechanism on the video Quality of Experience (QoE) perceived by the user. They reported that the RTT has no impact and that the QoE solely depends on the content server load and the peering agreement between ISPs. The results most relevant to our work are that the geographic proximity does not matter in Europe, while ISP-dependent policies do, and that the selection mechanism exhibits a distinct diurnal behavior.

In a previous paper [7] we analyzed the short term behavior of the server selection mechanism of YouTube. We conducted a measurement study based on proxy servers located within several European ISP networks. The main observation was that the YouTube server selection mechanism behaves quite differently. We developed a classification scheme to group specific selection patterns based on the observed temporal behavior.

The work presented in this paper differs from the mentioned contributions in several ways. Even though all papers mentioned above analyze the YouTube server selection mechanism, they do not consider its long term behavior. Furthermore, we investigate the server selection behavior observed at several measurement points within a single ISP network, during a time period of at least three months. Measuring in an ISP network, rather than at university sites or using PlanetLab nodes, provides results closer to the behavior seen from real users. This is due to the fact that these sites usually have high-bit rate connections, which might influence the server selection behavior.

3 Measurement Approach and Methodology

We developed a measurement probe system and placed such probe devices at different measurement points within a large German ISP network. These measurement probes operate as follows: A probe requests a set of videos from the YouTube CDN. The response of a video request comprises the video web site including a configuration field for the YouTube video player. This field contains the name of the video content server hosting the requested video. The location of the video server can be derived from the hostname of the video server (see [7] for further details). The IP address of the video server is resolved and the RTT to this IP address is actively measured, using ICMP Echo Requests/Replies. These steps are repeated periodically for the set of videos. It is not necessary to download the video itself, since all required information can be derived from the web site referring to that video. Thus the load on the YouTube CDN caused by our measurement approach is minimized. We placed five measurement probes at DSL-connections within the ISP network. Since DSL-connections are changing the IP addresses every 24 h, we also recorded the IP address currently assigned to a DSL connection. A measurement interval of 15 min was applied and the video set consists of 20 popular videos. Popular videos are more likely available at all data centers and will therefore not limit the cache selection behavior by mere video availability constraints. Thus, the responses of 20 popular videos were recorded every 15 min. In the following the result of one video request at one time instance is called a “data point”. The associated RTT measurement of one data point contains four single measurements from which the minimum, maximum, average and standard deviation are derived. Note that in our measurement campaign we regard all YouTube content servers with the same location identifier in the hostname as being part of one data center location.

Before analyzing the traces, it is convenient to filter out the outliers in the server selection measurement results. For the filtering process we define an outlier as follows: if a data center location is only seen in one time instance and neither in the preceding nor in the next two successive time instance and its share of all video requests at this time instance is below 20 %, it is considered as an outlier. The removal of these outliers leads only to a minor reduction of data points. Across all measurement traces only 0.6 % of the data points were identified as outliers.

4 Results

4.1 Impact of the IP Prefix on the YouTube Server Selection

As mentioned earlier, a DSL-connection gets a new IP address every 24 h. However, we never observed the same IP prefix at different measurement points. For example IP prefix 5 was only observed at the measurement point in Munich and never at any other measurement point. We even found that two probes got IP addresses of two different IP prefixes. This change in the IP prefix can make a huge difference in the experienced YouTube server selection behavior. Figure 1 shows two server selection pattern examples recorded from the same probe, but with different IP prefixes. For IP prefix

1 the YouTube requests are mostly served by the data center in Amsterdam throughout the day. In case of the second IP prefix (prefix 2) the data center in Hamburg is the preferred one.

Using the IP address to determine the geographic location of a user is a common approach, since it is assumed that users with the same IP prefix are geographically close. However, as a DSL connection might be served with IP addresses of multiple IP prefixes, this hinders the derivation of the geographic location from IP addresses. Contrary to other CDNs which solve the location problem by means of DNS, YouTube is solely relying on the IP addresses (see [7] for further details). This finding also infers that an operator-wide YouTube server selection behavior analysis would require at least one measurement probe per IP prefix of an operator. Based on this knowledge, we decided to group the measurement traces by their IP prefix rather than by the physical location of the measurement probe.




Fig. 1. Comparison of different server selection patterns for different IP prefixes (observed by the same measurement probe)

4.2 Long-Term Temporal Behavior of the Server Selection Mechanism

Figure 2 shows a 90 day trace for three different IP prefixes: IP prefix 1, IP prefix 5 and IP prefix 7. The white spaces for IP prefixes 1 and 7 are due to the prefix change of these measurement probes. The white spaces in IP prefix 5 are caused by missing data. From Fig. 2 it can be seen that the server selection behavior observed for one IP prefix is not correlated to the behavior of the other IP prefixes. The server selection behavior at IP prefix 1 is quite stable using mainly the data center in Amsterdam. For the last 4 weeks the server selection behavior of IP prefix 1 changes and also the data center in Düsseldorf is used. For prefix 5 two main server selection patterns are visible: either the data center in Munich is used exclusively or it is used only during the night until afternoon (i.e. in the low load phase) whereas the data center in Frankfurt is used in the remaining time, from afternoon to midnight (i.e. in the high load phase). IP prefix 7 is served steadily for the first 5 weeks from Frankfurt, but then the server selection pattern is only stable for short periods of time and incorporates different data centers.

Since we are interested in the effects of the server selection on the traffic distribution within an ISP network, we focus our study on the data centers used in the busy hour of the ISP network. According to [1] the busy hour is chosen to be the hour between 20 h and 21 h. A given IP prefix is served from a subset out of all possible data

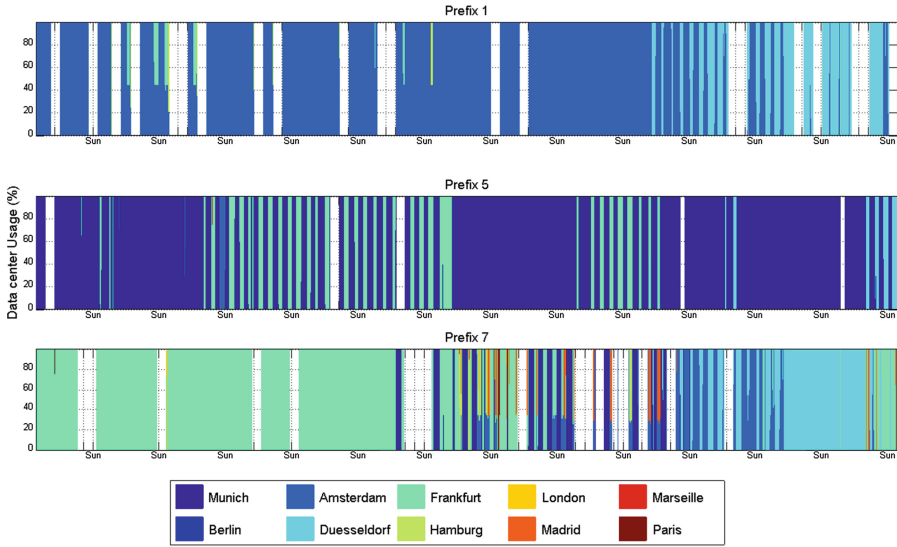


Fig. 2. 90 day server selection traces for IP prefix 1, 5 and 7

centers. Those subset elements contribute a certain share of served requests. If the share of an element changes by at least 25 % on busy hours of consecutive days we regard this as a YouTube server selection pattern change event.

In Fig. 3 the histogram of consecutive days with no server selection pattern change event is shown for IP prefixes 1, 5 and 7. The results confirm the observations from Fig. 2. Besides the 90 day trace documented in the figures of this paper, there are traces as long as 190 days available, which prove that server selection pattern change statistics also vary in the long run.

Table 1 shows how often a given data center is present in the busy and least-busy hour during the 90 day trace (relative frequency of data center occurrence). We defined the least-busy hour as the time from 5 h to 6 h. From the percentage values it can be seen that one or two dominant data centers exist per IP prefix. This holds true although the server selection behavior changes frequently for some IP prefixes.




Fig. 3. Histogram of consecutive days with no server selection pattern change event

Furthermore Table 1 shows that there is a dependency between the size of the data center and its usage during the busy hour. Since there is no information published by YouTube about the size of YouTube data centers, we follow the approach from [3] and estimate the data center size by the number of/24-prefixes assigned to the data center location. For IP prefix 1 the share of the data center in Amsterdam rises from 74 % to 99 % and the share of the data center in Düsseldorf drops from 27 % to 4 %. The same holds true for IP prefix 5 with respect to the data centers in Frankfurt and Munich. For IP prefix 7 a similar behavior is observed.

Table 1. Relative usage of data centers per IP prefix (busy and least-busy hour case)

	Number/24 prefixes	Prefix 1		Prefix 5		Prefix 7	
		Busy hour	Least-busy hour	Busy hour	Least-busy hour	Busy hour	Least-busy hour
Frankfurt	15	10 %	0 %	38 %	2 %	68 %	53 %
London	12	1 %	0 %	0 %	0 %	1 %	0 %
Amsterdam	6	99 %	74 %	1 %	0 %	32 %	1 %
Paris	6	0 %	0 %	0 %	0 %	3 %	0 %
Madrid	5	0 %	0 %	0 %	0 %	9 %	0 %
Munich	2	0 %	0 %	60 %	98 %	0 %	18 %
Hamburg	2	0 %	0 %	0 %	0 %	4 %	0 %
Düsseldorf	1	4 %	27 %	7 %	0 %	10 %	31 %

Even some time before the busy hour begins, a data center change which then remains in the busy hour can be observed. In the following, we have a closer look at this “pre-warning” time. For that, we processed exemplarily the trace of IP prefix 5 to determine the frequency of the hours of the day at which a data center change that then remains in the busy hour occur. The results of this analysis are depicted in Fig. 4. It shows for each hour of the day the probability that in this hour a data center change which then remains in the busy hour occurs. It is significant that in most cases the change can be observed already before 17 h, i.e. there is a pre-warning time of three hours. This analysis shows that there is no fixed time schedule for the data center changes, which leads to the conjecture of a data center load dependent change policy. Unfortunately we were not able to analyze this in more detail based on our measurement traces taken from a single ISP network as the load situation of the YouTube data centers result from YouTube requests coming from several ISPs.

4.3 Locality and RTT Characteristics of YouTube Data Centers

The YouTube data center locations were extracted out of the host names of the content servers (see [7]). From that we found that 76 % of all requests (conducted from all measurement probes including all IP prefixes within the whole measurement period) were served by data centers in Germany and 24 % by data centers in Europe. Only 0.01 % of the requests were served by data centers outside of Europe.

The prefix-dependent average RTT values are between 18 ms and 28 ms for the data center in Amsterdam and between 21 ms and 37 ms for the data center in London.

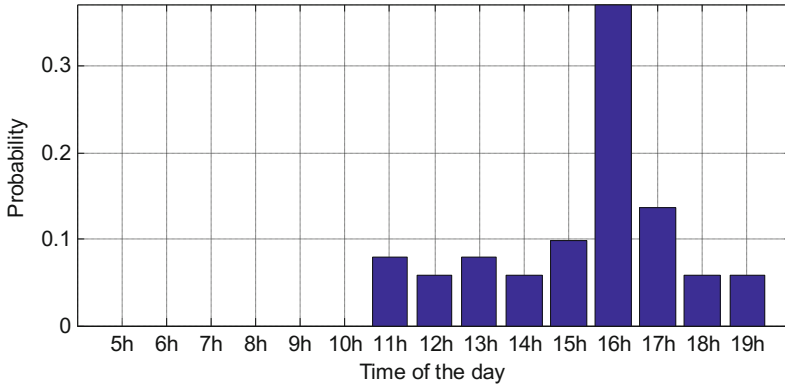


Fig. 4. Probability distribution of the hours of the day in which a server change occurs that remains in the busy hour (IP prefix 5 case)

Amsterdam and London count for 97 % of the requests served by European data centers in our traces. The RTT values for Amsterdam and London are similar to the values of the data centers located in Germany, as can be seen in Fig. 5. We found no explanation for the large RTT value measured for IP prefix 4 and the data center in Frankfurt. As for IP prefix 3, whose region is close to that of IP prefix 4, the RTT is significantly smaller (10 ms).

We also compared the RTT values in the busy hour, to the values in the least-busy hour. For this, we took the minimum RTT values of the data points, respectively. In Table 2 the mean, median and 95 % percentile of the (minimum) RTT values in the busy hour and in the least-busy hour are shown for each IP prefix. IP prefix 3 was excluded due to an insufficient number of RTT measurements. It can be seen that for IP prefixes 1, 2 and 4 the mean and the median values are almost identical in the busy

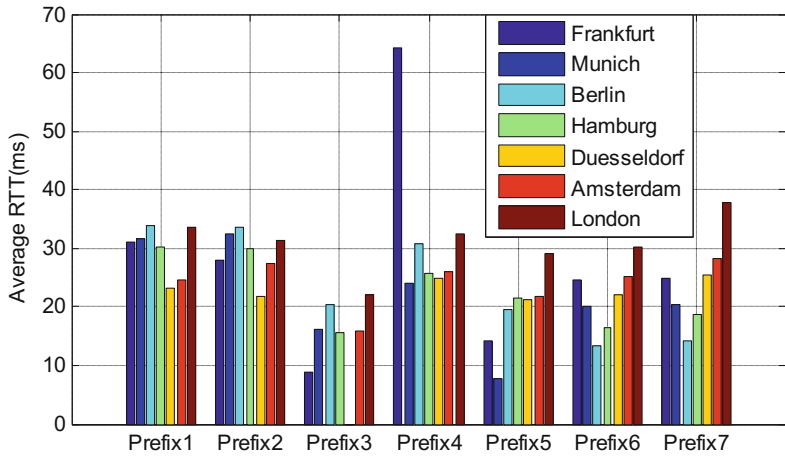


Fig. 5. Average RTTs related to selected data centers (for 7 IP prefixes)

Table 2. RTT values for 7 IP prefixes (busy and least-busy hour case)

RTT (ms)	Least-busy hour			Busy hour		
	Mean	Median	95 % per.	Mean	Median	95 % per.
Prefix 1	23.9	23.6	26.6	25.6	23.6	27.0
Prefix 2	28.2	29.6	30.4	30.1	30.0	31.4
Prefix 4	24.0	23.9	26.6	25.2	24.7	25.6
Prefix 5	8.2	6.9	13.8	10.9	9.4	15.7
Prefix 6	23.5	23.0	25.7	25.8	23.1	46.8
Prefix 7	22.9	22.5	25.3	25.9	23.0	44.0

hour and in the least-busy hour. For IP prefix 5 the RTT values increase by 35 % in the busy hour. For IP prefix 6 and 7 there is almost no difference regarding the median, but a significant increase in the 95 % percentile.

Table 3 shows the RTT values measured between selected IP prefixes and data centers, which mostly served these IP prefixes. It can be seen that there are only minor differences between the RTT values in the busy hour and in the least-busy hour. For example, the RTT related to IP prefix 5 and the data center in Munich increases only by 1 ms and for the data center in Frankfurt it does not increase at all. An explanation of this behavior might be that the load (congestion) in the ISP network is quite low even in the busy hour, so that the RTTs show no large variations. The RTT increase faced by some IP prefixes (5, 6 and 7) in the busy hour (depicted in Table 2) is caused by the selection of more distant data centers in the busy hour.

Table 3. RTT values related to selected data centers (busy and least-busy hour case)

RTT (ms)	Least-busy hour			Busy hour		
	Mean	Median	95 % per.	Mean	Median	95 % per.
Prefix 5 to Muc.	7.1	6.8	7.2	8.1	7.8	7.8
Prefix 5 to Fra.	13.4	13.3	14.2	13.4	13.4	14.2
Prefix 6 to Ams.	24.4	24.3	25.8	23.9	24.0	25.6
Prefix 6 to Fra.	22.8	22.8	24.1	22.8	22.8	24.2
Prefix 7 to Ams.	25.2	24.2	24.2	26.2	24.2	25.7
Prefix 7 to Fra.	22.9	22.7	24.2	24.5	22.8	32.4

The lower part of Fig. 6 shows the RTT values over the duration of 2 weeks for IP prefix 5. For comparison, the server selection behavior faced by IP prefix 5 is shown in the upper part of Fig. 6 (cf. Fig. 2). During the time, when requests from IP prefix 5 are only served by the data center in Munich, the RTT remains constant at 9 ms. When the data center in Frankfurt is used (in the busy hours of the second week) the RTT increases to 13 ms. In case of IP prefixes 6 and 7 the RTT increases are caused by selecting the even more distant European data centers in Madrid and Paris during the busy hour.




Fig. 6. Server selection behavior and corresponding RTT values for IP prefix 5 (2 week trace)

5 Summary

In this paper we report the results of a comprehensive analysis of the YouTube data center selection mechanism observed within a German ISP network. We investigated the dependency of the server selection decision on the IP prefix from which the request is sent, the long term temporal server selection pattern and the influences on the RTT. The results were obtained from a measurement campaign, where several measurement probes were placed in different regions within the ISP network and where traces were recorded over a 90-day period.

We discovered that the server selection behavior experienced by a user depends on the prefix of the IP address currently assigned to the user. No correlation between the server selection patterns faced by the different IP prefixes has been observed. Typically, only one or two dominant YouTube data centers are used per IP prefix independent of the server selection pattern. Furthermore, it has been observed that during the busy hours the IP prefixes are more likely served by larger (and more distant) data centers. This is the main reason for the increase in RTT in the busy hours. Contrary, the traffic load in the ISP network has no significant impact on the RTT as it is quite low even in the busy hours. These results go along with the findings of [5] and [6], which already reported the general behavior patterns, but not as detailed and precise as it is documented in this paper.

References

1. Sandvine.: Global internet phenomena report 2H 2013. <https://www.sandvine.com/trends/global-internet-phenomena/> (2013)
2. Finamore, A., Gehlen, V., Mellia, M., Munafò and Nicolini, M.M.: The need for an intelligent measurement plane: the example of time-variant CDN policies. In: 15th International Telecommunications Network Strategy and Planning Symposium, pp. 2–7 (2012)

3. Adhikari, V.K., Jain, S., Chen, Y., Zhang, Z.-L.: Vivisecting YouTube: an active measurement study. In: IEEE International Conference on Computer Communications: Mini-Conference, pp. 2521–2525 (2012)
4. Adhikari, V.K., Jain, S., Zhang, Z.: YouTube traffic dynamics and its interplay with a tier-1 ISP: an ISP perspective. In: Proceedings of the 10th ACM SIGCOMM Conference Internet Measurement, pp. 431–443 (2010)
5. Torres, R., Finamore, A., Kim, J.R., Mellia, M., Munafo, M.M., Rao, S.: Dissecting video server selection strategies in the YouTube CDN. In: 31st International Conference on Distributed Computing Systems, pp. 248–257, June 2011
6. Plissoneau, L., Biersack, E., Juluri, P.: Analyzing the impact of YouTube delivery policies on user experience. In: International Teletraffic Congress ITC'24 (2012)
7. Windisch, G., Knoll, T.M., Bauschert, T.: Analysis of the video server selection in the YouTube CDN. In: 2nd European Teletraffic Seminar (2013)