Computer Communications 116 (2018) 118–131

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

A user-oriented performance comparison of video hosting services

Alessio Botta[*,a,b], Aniello Avallone[a], Mauro Garofalo[a], Giorgio Ventre[a,b]

[a] University of Napoli Federico II, Napoli, Italy
[b] NM2 srl, Napoli, Italy

ARTICLE INFO

Keywords:
Network neutrality
Internet video streaming
Network performance

ABSTRACT

Internet streaming is responsible for a significant fraction of Internet traffic. It has been reported to account up to 70% of peak traffic in North American fixed access networks, and this figure is expected to reach 80% by 2020 [1]. Regarding such a killer service of the Internet, much has been discussed regarding if and how video hosting providers violate or may violate neutrality principles, in order to give users a "better" service compared to their competitors or other services. In this paper, we provide a contribution to this discussion studying three video hosting providers (i.e. YouTube, Vimeo, and Dailymotion). Specifically, we analyze their delivery infrastructures, including where the servers that provide videos are located, and the performance from a user viewpoint. To assess the performance, we measure throughput and RTT as experienced by users watching real videos of different popularity from several locations around the world and at different day hours.

We uncover the performance differences of these providers as a function of the different variables under control and move a step forward to understand what causes such differences. We also study the changes in the infrastructures and related performance over time, performing different measurement campaigns over different years.

Our results allow to understand what are the real performance users currently get from these providers and if the performance differences observed can be due or considered as a violation of network neutrality principles.

1. Introduction

There is a long ongoing debate on the so-called "network neutrality". Several definitions exist for this term, and they all share the common idea that data on the Internet should be treated in the same way despite several its characteristics such as technology, device, application, service, user, provider, and the country they come from or go to. An early debate about network neutrality regarding Internet traffic management policies appeared in 2003 [2], but concerns about possible threats to the end-to-end nature of the Internet raised already in the late 1990s [3]. Nowadays the debate has gained momentum also because of recent events such as the one involving the provider Comcast, which was slowing down uploads from peer-to-peer file sharing applications [4]. The discussion on whether the Internet should be fully neutral, or rather providers should be allowed to use techniques to differentiate traffic does not concern only economic aspects, but also and increasingly both legal and regulatory ones. For example, a research work regarding legal aspects was presented in 2011 [5]. In our paper, we do not want to advocate a position pro or against network neutrality. We rather aim at providing a contribution to understanding the current situation from a user viewpoint.

Our work focuses on three Video Hosting Services, YouTube, Vimeo, and Dailymotion, for which we measured the performance achievable by end users depending on video popularity and user location (i.e. the country) and inferred the characteristics of the infrastructure used for video delivery. The choice of these providers was driven by two properties: the global diffusion of the service, in order to identify every possible country-specific violations of network neutrality (e.g. due to censorship issues), and the service model of the providers, not asking subscription fees to the users. We analyze traffic related to video streaming because this service accounts a very high share of Internet traffic (about 70% according to CISCO [6]). The highlights of our work could be summarized as follow:

- We introduce a provider-independent methodology that allows to capture, analyze, and compare the performance statistics of the video hosting services.
- We measure, analyze, and compare these statistics for YouTube, Vimeo, and Dailymotion, from several locations around the world.
- We compare the performance obtained by the different providers.
- We provide insights on the topology of the infrastructures and routing policies used by the video hosting services to deliver their

* Corresponding author.
E-mail addresses: a.botta@unina.it (A. Botta), anie.avallone@studenti.unina.it (A. Avallone), mauro.garofalo@unina.it (M. Garofalo), giorgio@unina.it (G. Ventre).

https://doi.org/10.1016/j.comcom.2017.11.016
Received 20 February 2017; Received in revised form 6 October 2017; Accepted 26 November 2017
Available online 02 December 2017
0140-3664/ © 2017 Elsevier B.V. All rights reserved.

content.
- We compare the results obtained over different years to investigate recent changes to the performance and the infrastructure.

The remainder of the paper is structured as follows. Related work and highlights of the novel aspects of this work are presented in the next section. Available information regarding the infrastructures of video hosting providers is given in Section 3. In Section 4, we present the way we collected the dataset and our analysis methodology. Section 5 describes in detail the results of two measurement campaigns and with a longitudinal comparison. This section is followed by a discussion about geographical location (Section 6). Finally, we conclude the paper in Section 7.

2. Related work

There are several interesting works related to OTT streaming video services. A comparison of the content delivery frameworks of YouTube, Dailymotion, and Metacafe was provided by Saxena et al. [7]. The performance measures were collected using PlanetLab [8] nodes deployed in 9 different countries to have a global perspective. This work covered three aspects: the measurement of QoS, the investigation of the service delay variation, and the analysis of QoE. These evaluations considered the different geographical locations of end-users and how video content storage and distribution were impacted by the meta information associated with videos, such as popularity (i.e. the number of views) and video ages. Finally, the authors inferred the content delivery frameworks of the providers, showing that all providers relied on one or more CDNs (Content Delivery Network) to deliver their contents. With respect to this work, our measurements comprise Dailymotion, for which several sources of information revealed a “centralized” infrastructure with all its servers deployed in France.

Often in literature, the analysis of video hosting services is focused only on YouTube, since it generates a significant share of Internet traffic. An extensive data-driven analysis of YouTube concerning users behavior, video popularity, and their evolution was presented in 2007 [9]. Firstly, the authors compared YouTube and Daum, two video providers of User Generated Content (UGC), with non-UGC video providers (such as Netflix and Yahoo! Movies). Secondly, the authors made an extensive analysis of meta-information of videos, to investigate user behavior and video popularity distribution patterns. The data collection was related to several years and involved video information both fixed (such as category and length), and time-varying (such as the number of views and ratings). A tool to measure QoS and QoE of YouTube was designed in 2012 by Plissonneau et al. [10]. Metrics collected by a hundred of volunteers were analyzed by the authors to infer the video delivery policies of YouTube and understand how these metrics were impacted by Internet Service Providers (ISPs). Finally, a comparison of YouTube policies in the US and Europe was presented. One of the earliest analysis of HTTP video streaming with a comparison between YouTube and Dailymotion was presented in 2012 [11]. The authors performed passive measurements from a residential ISP network (i.e. Orange) to infer video characteristics (such as duration and encoding rate) and TCP-level performance (such as RTT and packet loss rate). A recent work [12] studied Netflix and Hulu, two leading providers of subscription-based video streaming services in the US. Both providers use the same three CDNs (i.e. Akamai, Limelight, and Level3) for video content delivery. The authors performed passive and active measurements using both residential users and PlanetLab nodes as vantage points. The aim was to uncover the provider architectures and their different CDN selection strategies. Results show that neither Netflix nor Hulu used the network conditions to choose the CDN. Consequently, the authors proposed an alternative CDN selection strategy to improve QoE to end-users. Unlike our work, all the vantage points were located in the US and both services needed a subscription fee to be accessed. Furthermore, rather than studying the QoE perceived by end-users and how to improve it, we aim to understand the reasons behind the performance differences experienced by end-users when accessing videos of different providers from different geographical locations around the world.

To the best of our knowledge, only our previous work [13] evaluated the performance of Vimeo. At the time of that measurement campaign (from 2014 to early 2015), Vimeo used only Akamai CDN as the content delivery platform. In the first quarter of 2016, during our second measurement campaign, describe in this paper, Vimeo has switched to Fastly CDN. For this reason, we also relate our work with ones regarding the evaluation of these two CDNs. For a comprehensive description on CDN in general and more specifically on the Akamai infrastructure the reader is referred to [14] and [15]. The former work gave the first insights into the overall infrastructure of Akamai. The authors presented an overview of the Akamai network, describing the mechanisms used for redirect user requests and the approaches to optimize the content delivery to end-users. Moreover, this paper provided an “agenda” of technical issues encountered in the development of the Akamai CDN. Several aspects of Akamai infrastructure with an overview of all components composing its platform and their capabilities were reported by the latter. Moreover, a comprehensive description of how Akamai redirects client requests to the “nearest” available server (load balancing system) and how its servers deliver content including video streaming (delivery policy) was provided. An extensive measurement of the Akamai network was performed in 2009 [16]. The aim of this work was to infer information about network condition, measuring network paths and refresh frequency of Akamai DNS server. The probing phase relied on 140 PlanetLab vantage points. Measurements were performed sending DNS requests to Akamai customers and then gathering the IP address of the Akamai edge servers. The analysis showed that redirection depends on the latency between clients and edge servers. Using DipZoom, a peer-to-peer Internet measurement platform, authors of [17] exposed the distribution of Akamai edge servers and performed active measurements to estimate the performance of Akamai infrastructure. Furthermore, a performance comparison between the Akamai CDN configuration and a possible consolidated configuration was presented, where a high amount of servers were clustered.

More focused on the geographic location of the infrastructures, Padmanabhan and Subramanian [18] built a service that pairs the IP address of an Internet host with its geographical location. The authors proposed three techniques to infer the target host position. *GeoTrack*, based on information provided by DNS server about the target and its “neighbors”. *GeoPing*, exploiting the correlation between RTT and geographic distance between target and vantage points with a well-known location. Finally, *GeoCluster*, grouping IP addresses into clusters assuming that all hosts in a cluster are geographically near, and combining partial host-to-location mapping information and BGP prefix to infer the host location. A recent work gathered all servers of Google infrastructure in serving sites then localize them using a technique called Client-Centric Geolocation (CCG). The CCG is based on the hypothesis that clients that are directed to the server are likely to be topologically, and probably geographically, close to the server [19].

Summarizing, studies more relevant to our work investigated either the CDN infrastructure and performance measures or the geographic location of such infrastructure. Our work moves a step forward with respect to existing literature. To the best of our knowledge, our work is the first study that presents a comparative analysis of the performance of YouTube, Vimeo, and Dailymotion. We also present the results of a second measurement campaign made one year after the first one, to evaluate how the delivery infrastructure of each provider evolved over time.

Unlike works related to residential ISPs networks, which involve a large number of volunteers or needs measurements from the network of the ISPs, we perform active measurements using a globally distributed research infrastructure (i.e.PlanetLab). Performance indicators collected may be different from the ones of residential users. However, our primary aim of comparing the different services (to understand if and how they violate or may violate network neutrality principles) is not affected by this choice.

3. Infrastructures of providers

This section describes the infrastructure of the video hosting services we analyzed in this work. Although there are several studies focused on YouTube (Google) infrastructure (e.g. [19]), that try to identify the number, structure, and location of caches and servers, less information is available for the other providers. In general, information about these infrastructures is not publicly disclosed. We crossed several sources of information for obtaining the views that follow. Then we verify and confirm the accuracy of such views with the experiments in Section 6.

3.1. Dailymotion

Dailymotion was launched in France in 2005. Originally its infrastructure consisted of one homemade Linux cluster with limited connectivity able to serve only a few thousand users, [20,21]. Afterward, it moved to a more scalable architecture based on a network file system. Through this configuration, input/output bandwidth, caching, and latency was shared throughout the whole system and performance scales linearly with the numbers of nodes. There is no evidence of official agreements between Dailymotion and third-party CDN until June 2014, when it advertised that Orange Business Services and Akamai Technologies were chosen as technical partners to optimize the delivery of video content for users in Europe, North America, and Asia [22]. Thanks to this agreement, Dailymotion was able to exploit the Akamai "Media Delivery Solutions" that consists of an extensive network of servers to accelerate the distribution of video content over the Internet. The provider used this platform for the delivery of its "premium channels" dedicated to the live streaming of sports and gaming events [23]. Hence, Dailymotion differentiated its users, delivering contents through two different platforms. Contents for "non-paying" users were provided using data centers in France while videos of the live premium channels were delivered through a distributed platform. Our measurements of Dailymotion concerns the first distribution channel, consequently in the following we consider from our viewpoint its infrastructure as centralized.

Since February 2016, Dailymotion is using the Limelight (CDN) Orchestrate Delivery service to deliver broadcast quality videos in Asia Pacific (APAC) region to all devices connected to the Internet [24].

3.2. Vimeo

Vimeo was founded by a group of filmmakers in 2004. Since 2015, it used the Akamai CDN to distribute its contents [25]. Akamai has a very broadly deployed network of edge servers, consisting of tens of thousands of nodes globally deployed in over 120 countries, and within thousands of networks around the world [15]. Nowadays, it delivers a significant fraction (i.e. up to 20 percent [26]) of the Internet web traffic. Akamai's edge servers are located within thousands of ISPs networks. Through the partnerships with these providers, a huge audience of Internet users around the world was within a single "network hop" to an Akamai CDN server. Akamai provides a CDN service dedicated to the streaming media content deployment, named Adaptive Media Delivery [23]. This service allows the transmission of video streams with Adaptive Bit Rate and backs up in Akamai NetStorage, for later viewing [27]. Akamai has also developed a custom version of TCP/IP to optimize the transmission speed. This protocol called Fast TCP uses the delay as a measure to control network congestion and improve the throughput. The CDN exploited it for the acceleration of both video distribution, and download [28]. During our second measurement campaign, Vimeo turned from Akamai to Fastly CDN [25]. Unlike Akamai, it consists of fewer powerful PoPs (about 30 location in only 11 countries) with large amounts of memory (i.e. a cache server configuration with 768 GB of RAM and 18 TB of SDDs) to provide a higher cache hit ratio.

3.3. YouTube

YouTube, born in 2005, is the most popular service on which users can share and watch video content. After Google had acquired it in 2006, the content distribution of YouTube became entirely provided within the Google network infrastructure. In this infrastructure, the following elements can be distinguished.

- *Core Data Centers*: A set of high-efficiency backend servers deployed in data centers in Americas, Europe, and Asia. They are used for both computation and storage operation, and a private WAN connects them together. This network is based on Software Defined Networking (SDN) principles and OpenFlow to manage switches [29].
- *Edge Points of Presence (PoPs)*: Cache servers distributed worldwide [30]. PoPs represent the terminal nodes of Google network and are connected via peering with ISPs to deliver services to users. In particular, YouTube has a three-tier caching infrastructure (Primary, Secondary, and Tertiary Cache Cluster) that comprises four different *logical* namespaces (lscache, tccache, altcache, and cache). The logical naming structure of these cache clusters is reported in literature [31].
  - *Primary cache*: Each primary cache cluster has 192 logical caches corresponding to the *lcache* namespace. a_c.v[1–24].lscache[1–8].c.youtube.com where "a_c" matches the IATA airport code [32].
  - *Secondary cache*: Each has 192 logical caches corresponding to the *tccache* namespace. tc.v[1–24].cache[1–8].c.youtube.com
  - *Tertiary cache*: Each has 64 logical caches corresponding to the *cache* or *altcache* namespaces.
- *Backbone*: a global fiber network to interconnect data centers and deliver traffic to Edge PoPs.
- *Edge nodes*: also called *Google Global Cache (GGC)*, include Google servers deployed inside the network of ISPs. Therefore, Edge nodes represent the tier of infrastructure closest to users. Static, popular content, such as YouTube video, is stored on edge nodes. Thus, the user requests are redirected by Google DNS to the edge node that will provide the best QoE [19].

Expected performance

Summarizing, analyzing the aforementioned knowledge about these providers we have inferred for each one its delivery infrastructure. Considering these assumptions, we can hypothesize that providers using their own (i.e. YouTube) or third-party (i.e. Vimeo) CDNs will have higher performance than those which have centralized infrastructures, such as Dailymotion. Specifically, we foresee that the performance gap, both in terms of RTT and throughput, will result much wider in the countries geographically farthest from France (the country where the Dailymotion infrastructure is deployed). Subsequently, we will study more in depth the network paths from clients to servers that physically contain video files. In fact, any Server Caches or Peering agreements may have the effect of placing servers in the same Autonomous Systems of the client. In these cases, we expect very high throughput (in the order of tens of Mbps, as in a local network) and very low RTT values (in the order of few milliseconds, as in a local network), differently from providers that do not have these kinds of agreements, for which the performance parameter values will be strongly dependent on the geographical location of the client (i.e. the distance from the servers).

4. Methodology and tools

In this section, firstly we describe the meta information of the videos gathered in the data collection phase. Then we briefly introduce the infrastructure used to perform the measurements campaign. Finally, we outline the procedure for downloading the videos, measuring the performance, and discovering the path from clients to providers.

4.1. Scope

The aim of our work is to evaluate the performance of video hosting services to understand whether performance differences are noticed and could impact network neutrality. To do that, we evaluate the performance of video hosting service providers for different kinds of content and different geographical locations.

The aim of our measurements is to discover:

- Differences in performance of infrastructures depending on the countries from which the measurement is made;
- If there are *cache-servers* deployed inside ISP or Internet Exchange Point (IXP) infrastructures;
- If discriminating routing policies exist, and how and when they are implemented.

4.2. Collection phases

The data collection phase is divided into two distinct measurement campaigns. The first campaign covers several months between the end of 2014 and the beginning of 2015. The second campaign lasted from March to June 2016.

The results reported in the following sections refer to a single day of experimental measurements. The choice was guided by the fact that neither of the two measurement campaigns had pattern behavior dependent on the particular day in which the measurements were made. Therefore, the results reported are representative of what we observed in all the other days. In the collection phases, four videos are downloaded from each provider. Each video falls into one of four video categories. We define these categories depending on video popularity (i.e. the number of views): less than 500 views, between 10K and 120K views, between 120K and 1M views, and over a million views. Furthermore, the providers had made available the same video at different resolutions. We used videos with a resolution of 720 p for all the categories. Each measurement run is performed over a 24 h period. In each run, the batch task of video download is started every two hours.

4.3. PlanetLab

A set of geographically distributed vantage points (VP) has been used in order to measure the performance of providers from all around the world. We performed our experiment campaigns using PlanetLab nodes as VP. VP of the first campaign are deployed in 32 countries, for a total of 200 PlanetLab nodes, The second campaign relies on a fewer number of VP in 21 countries. For both campaigns, the measurements involved the countries of 5 continents: North America, South America, Europe, Asia, and Australia. PlanetLab uses high-speed networks inside Research Centers and Universities. Therefore, our analysis cannot describe the behavior of the video providers as seen by residential users. However, our aim is to compare the performance of such providers, therefore, the use of PlanetLab nodes as VPs does not affect our evaluation.

4.4. Measurements

For each PlanetLab node, we perform a set of simple operations in batch to acquire the data, as described in the following.

- Running in background *youtube-dl* to download the video.
  In our data-collection workflow, Youtube-dl software is used to measure the throughput. It provides a periodic throughput estimation (with a period of 1 s) throughout the video download phase. Each of these estimations is then considered as a throughput sample. The standard deviation can therefore be computed also for each video download.
- Using *lsof* to detect the IPv4 address of the server to which the client has been redirected through DNS or other means. It guarantees to identify the server that physically contains the video.
- Given the IPv4 address of the server that physically contains the video, running *ping* to evaluate RTT and TTL.
- Finally, using *traceroute* tool to discover the path from the client to server and mapping the name of routers in the path and the Autonomous System where the IP address of the server is mapped.

To give a good approximation of the network status experienced by the user watching videos, RTT, and TTL measurements, as well as the path evaluation with *traceroute*, have also been assessed during the video download.

4.5. Architectures and techniques for video streaming optimization

Video hosting providers can adopt various approaches to cope the fragmentation of terminals and network connection issues. On the one hand, in the last decades, researchers have introduced several Information-Centric Network (ICN) [33] paradigms to define a more efficient architecture for the Internet. Among them, Named Data Networking (NDN) [34] is one of the promising candidates to deal with the data-centric nature of many existing Internet services, such as video streaming [35]. In the literature, there are several papers dealing with the problem of streaming video exploiting the concepts of the NDN. Authors of [36] investigated the performance of the principal DAS techniques in NDN context. The paper has exposed how the streaming performance could be influenced by the cache sizes and the caching policies. An application of live video streaming (i.e. NDNVideo) was implemented by Kulinski and Burke [37] exploiting the NDN architecture paradigms. To the best of our knowledge, none of the providers under test uses the NDN solution for its infrastructure. On the other hand, the state of the art for on-demand and real-time video streaming services is exploring solutions based on Dynamic Adaptive Streaming (DAS). With DAS we refer to a set of techniques to adapt the bitrate in response to changing bandwidth conditions. Among the providers under test, only Vimeo seems not to use a web player that supports a DAS technique for video delivery. Therefore, even if *youtube-dl* supports two spread DAS standard implementation, namely DASH (Dynamic Adaptive Streaming over HTTP) [38] and HLS (HTTP Live Streaming) [39], we have performed video download operations without taking advantage of these techniques. Only in this way was it possible to ensure a fair performance comparison of the providers.

5. Evaluation

In this section, we present the performance of the three video hosting services obtained from two measurement campaigns (Sections 5.1 and 5.2). Both campaigns aimed at evaluating throughput, RTT, and distance in hops, related to users downloading videos of different popularity from several locations around the world. In the following, we first report the results of each campaign. Then, a comparison between the campaigns is provided to assess the evolution of the infrastructures over time. The following are the results of the two campaign presented with various statistical metrics with the aim of taking into account also the impact of the high standard deviations, where present. Hence, the conclusions consider not only the averages and standard deviations, but also the other calculated statistics, such as quartiles, minimums and maximums, and medians.

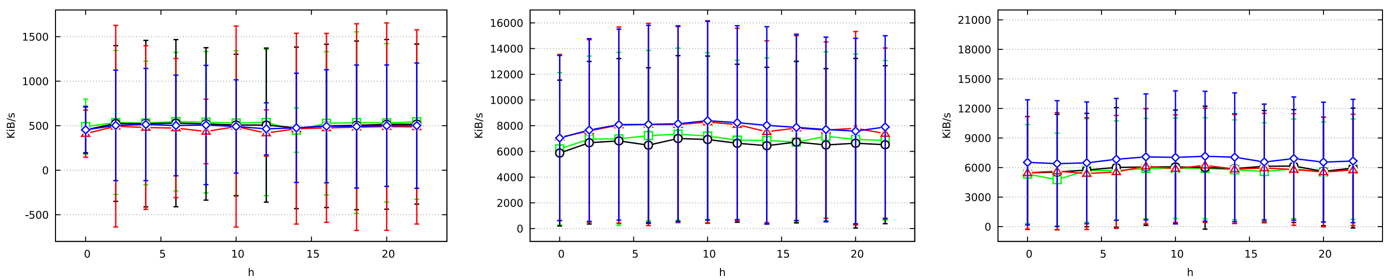

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 1. Average throughput over 24h - 2015.

Figs. 1 and 6 show high standard deviation values of the throughput for all the providers under measurement regardless the video category considered. These high values are mainly due to two factors: the data aggregation by country and the methodology for calculating the throughput statistics (i.e. average and standard deviation). In the first case, the throughput values are measured by nodes deployed in different countries with differences between the measures that can reach up to 1 order of magnitude. These values do not depend on the infrastructure of the providers but on the different networks where our PlanetLab nodes are deployed. In fact, such variations are observed for all the providers. In the second case, as anticipated in Section 4.4, we used youtube-dl in order to measure the throughput statistics for each provider. Table 1 shows an example of the throughput measurement for a single video download.

During these measurements, the trend is not constant throughout the download period, but the values decrease until they converge to a constant value, after a transient phase in which very high throughput is measured. This trend is probably related to buffering policies of providers, which slow down the bitrate of the download in case such rate is much larger that the video data rate.

5.1. First measurement campaign

5.1.1. Throughput

Table 2 shows different statistical indicators of throughput (i.e. minimum, first quartile, median, average, third quartile, and maximum) calculated over all the clients for each provider and each video category. This table allows to have a first, gross-grain comparison of the providers. Vimeo and YouTube show similar metrics while Dailymotion obtained an average throughput value that is one order of magnitude smaller. The median values of the video different categories for Dailymotion are about 400 KiB/s and most of the clients do not reach more than 500 KiB/s. The average value of throughput is influenced by specific cases that we will discuss more in depth in the following. Both Vimeo and YouTube show similar trends, although Vimeo has larger values of throughput for each video category, as shown in Table 2. Also in these cases, the average values are affected by spikes which significantly deviate from the median. For example, regarding Vimeo, we measured throughput values of about 39 MiB/s and RTT values of about of 0.13 ms, as shown in the following paragraph. This high performance is related to PlanetLab nodes which are only one “hop” away from Vimeo servers. A brief discussion about these values will be presented in Section 6.

Table 1
Throughput measurement example.

Time (sec)	Download (%)	Throughput
1	0.1%	1.57 MiB/s
2	0.5%	2.32 MiB/s
3	1.9%	4.60 MiB/s
4	7.4%	5.95 MiB/s
5	29.7%	1.19 MiB/s
6	40.9%	697.45 KiB/s
7	45.4%	614.91 KiB/s
8	50.0%	570.85 KiB/s
9	54.9%	529.90 KiB/s
10	59.4%	505.70 KiB/s
11	63.9%	484.10 KiB/s
12	68.6%	468.55 KiB/s
13	73.4%	453.73 KiB/s
14	77.7%	442.69 KiB/s
15	82.7%	432.25 KiB/s
16	87.4%	425.79 KiB/s
17	92.0%	418.95 KiB/s
18	96.7%	411.62 KiB/s
19	100%	407.91 KiB/s

Table 2
Throughput (KiB/s)–Statistical indicators 2015.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	26	355	416	525	505	15049
dm-120K	27	353	403	506	476	13797
dm-10K	26	344	379	467	420	16964
dm-500	19	355	403	492	472	9804
vi-1M	24	1788	5368	6935	9763	35013
vi-120K	19	1770	4986	6597	8946	34789
vi-10K	17	1976	6099	7781	10264	38999
vi-500	29	2132	6212	7889	10295	38833
yt-1M	12	1192	4387	5672	8215	23303
yt-120K	34	1125	4145	5901	8453	34686
yt-10K	32	1113	4008	5783	8583	34897
yt-500	7	1535	5013	6788	8997	31763

5.1.2. Round Trip Time

The performance values related to the delay (i.e. RTT) are shown in Table 3. Regarding Dailymotion, the interquartile range (IQR) of RTT varies between 45 and 150 ms, with median values of about 102 ms. These values are significantly larger than ones recorded by the other providers and could be related to the centralized nature of Dailymotion

Table 3
RTT(ms)–Statistical indicators 2015.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	3.38	45.25	103.06	116.57	151.97	1013.37
dm-120K	3.57	44.93	101.62	114.70	147.66	1546.63
dm-10K	3.52	43.87	101.85	113.25	149.50	697.50
dm-500	3.42	45.40	101.68	114.57	148.61	473.36
vi-1M	0.14	2.43	6.72	16.83	16.79	421.08
vi-120K	0.14	2.57	7.25	17.27	18.44	405.12
vi-10K	0.14	1.95	5.68	14.09	15.23	473.73
vi-500	0.14	2.02	5.85	14.30	15.03	358.91
yt-1M	0.22	2.60	7.28	31.15	23.51	510.76
yt-120K	0.21	2.90	7.62	32.41	27.96	492.32
yt-10K	0.24	3.25	7.78	32.40	29.86	420.90
yt-500	0.21	2.45	6.93	28.68	19.33	488.53

Table 4
Jitter(ms)–Statistical indicators 2015.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	0.09	0.18	0.51	2.73	1.56	850.00
dm-120K	0.01	0.22	0.69	3.21	1.99	1863.66
dm-10K	0.01	0.15	0.42	2.14	1.36	213.85
dm-500	0.01	0.16	0.51	2.52	1.61	399.19
vi-1M	0.00	0.08	0.20	1.60	0.67	138.44
vi-120K	0.00	0.08	0.20	1.71	0.63	249.12
vi-10K	0.00	0.08	0.18	1.42	0.62	220.89
vi-500	0.00	0.09	0.21	2.01	0.72	222.57
yt-1M	0.00	0.05	0.17	1.93	0.78	241.94
yt-120K	0.03	0.05	0.17	1.78	0.77	293.25
yt-10K	0.03	0.06	0.17	1.85	0.84	348.39
yt-500	0.00	0.06	0.18	1.85	0.88	184.31

infrastructure (see Section 3.1). Regarding Vimeo, there is a smaller IQR, from 2 ms to 17 ms, with minimum values that are smaller than 0.142 ms, independently of the video category. The smallest RTT value, about 0.13 ms, has been recorded by a PlanetLab node (*planetlab1.arizona-gigapop.net*) located in the United States. YouTube shows a behavior similar to the one of Vimeo. Its servers show RTT values comparable with those of Vimeo, although on a wider range (from 0.2 ms to 25 ms). The small RTT values reported for Vimeo and YouTube are basically due to the wide distribution of their delivery infrastructures, respectively Akamai CDN and Google (see Section 3). These infrastructures allow to obtain better performance regarding both throughput and RTT, even when averaged on the several countries from which the client requests the content.

5.1.3. Jitter

The performance values related to the jitter are shown in Table 4. Regarding Dailymotion, the interquartile range (IQR) of jitter varies between 0.15 and 1.99 ms, with median values of about 0.5 ms. As for RTT values, Dailymotion jitter values are significantly larger than ones recorded by the other providers. The worst values of jitter have been recorded by a PlanetLab node located in the Canada. Regarding Vimeo, there is a smaller value of IQR, from nearly 0 ms to 2 ms, with minimum values that are zero for all the video category. YouTube shows a behavior similar to Vimeo. Its servers show jitter values comparable with those of Vimeo, on almost the same range (from 0 ms to 1.9 ms). The small values reported for Vimeo and YouTube are due to the wide distribution of their delivery infrastructures, as mentioned for RTT.

5.1.4. Temporal behavior

In the following sections, we will point out how performance evolve over the day. The aim of this kind of analysis was to identify behavioral patterns in traffic at different day hours.

Fig. 1 shows the throughput achieved by the three providers during a whole day. Regarding Dailymotion, Fig. 1(a) shows that the average value is almost constant over the day. On the other hands, values highlight peculiar behavior at 8h and 12h, which refer to servers deployed in South Korea and Singapore. These values, as opposed to the assumptions in Section 3, suggest the presence of cache servers near these countries. A further investigation to understand the behavior of these PlanetLab nodes is presented in Section 5.1.5. Fig. 1(b) and (c) show the temporal evolution of the throughput for Vimeo and YouTube respectively. The mean values of this parameter for both providers are constantly above 5000 KB/s, an order of magnitude larger than Dailymotion. The high variability observed for each hour of the day is mainly due to the aggregation of very different values obtained from all over the world and to the buffering policies of the providers, as explained at the beginning of Section 5.

5.1.5. Performance evaluation by country

In this section, the evolution of the throughput and RTT as a function of the country is presented. Data acquisition is performed over different one-day (24 hours) periods, with tests carried out at two-hour intervals, from each of the countries reported below. Fig. 3(a) shows for Dailymotion the average throughput of each video category in each country. This figure supports our hypothesis about the centralized location of all Dailymotion servers: the performance is better in France than in all the other countries, and all European countries have better performance than non-European ones. Servers located in Singapore and South Korea achieve a maximum throughput of about 12 MiB/s that is ten times higher than that of the other countries. Values related to all countries but Singapore and Korea are shown in Fig. 2, where we have excluded the values of Singapore and South Korea to make the graph more readable. Fig. 3(b) depicts the average throughput measured for Vimeo in each country. Due to the wide spreading of Akamai, a plausible explanation for why we observe the high variability of throughput is the different network quality in each country.

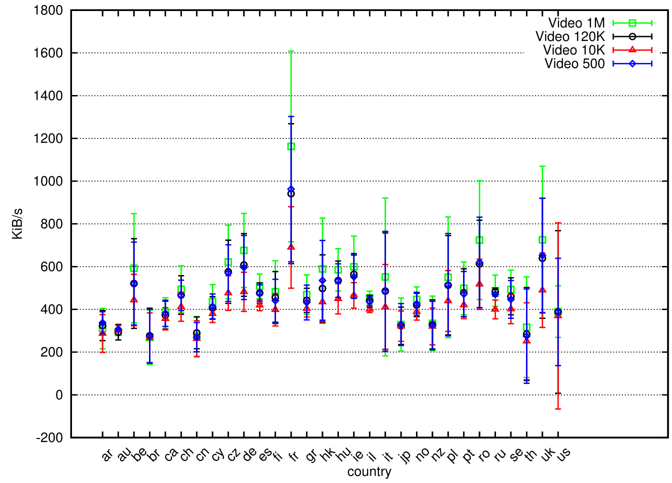

Fig. 2. Average throughput in each country excluded Singapore and Korea.

Regarding the analysis of various video categories, there is an overlap of the values of the average throughput in almost all the countries, with higher values for the videos having less than 500 views, denoting the presence of different treatments according to the video category. Finally, Fig. 3(c) shows mean and standard deviation values of the throughput of YouTube servers. Google distribution network is based on *peering*, anticipated in above mentioned Section 3. In this case, network infrastructure of the given country strongly influences the performance. The management through *peering* has led the CDN to define supply agreements with third-party companies, in order to obtain the widest possible capillarity. The comparison of mean values of different countries denotes variability related to the country. As with other providers, video with fewer views have performance slightly better than others categories.

5.1.6. RTT by country

The average and standard deviation values of RTT, regarding Dailymotion, are shown in Fig. 4(a). The overlap of the average values of RTT means that no treatments to the different categories of videos are applied. The European countries have smaller RTT values than other countries. In particular, the average value is in most cases smaller than 50 ms, while extra-European countries have values always larger than 100 ms. Specifically, RTT values in South Korea and Singapore are respectively 300 ms and 350 ms. Using traceroute, we made an analysis of the paths followed by packets traveling from the clients in South Korea towards the servers. Firstly, packets pass through hops inside the network of the Kookmin University (where the PlanetLab nodes are deployed), then they travel towards Level 3 US and Level 3 Paris, and finally, they arrive at the Dailymotion server. The same analysis was repeated for the clients in Singapore and results confirm the hypothesis

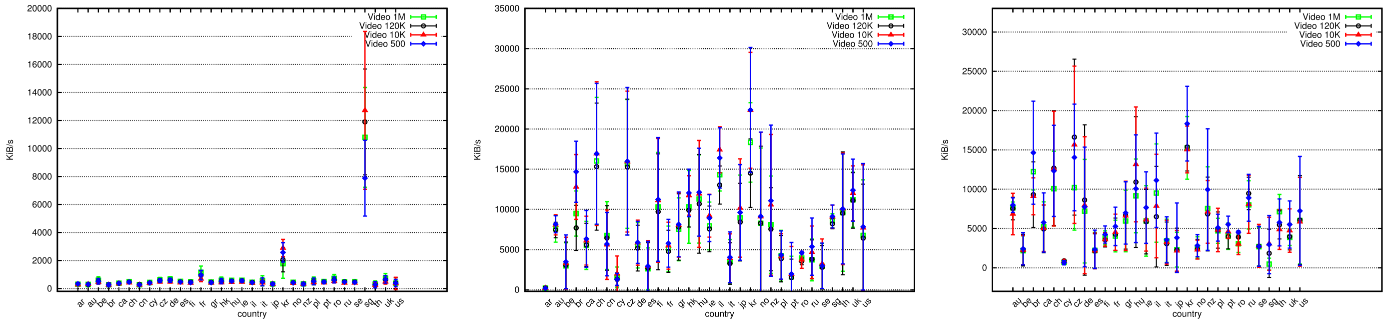

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 3. Average throughput in each country 2015.

of a centralized infrastructure for Dailymotion, as anticipated in Section 3. Concerning the RTT values of Vimeo, shown in Fig. 4(b), the majority of countries record values smaller than 20 ms. The average is similar for each video category. Therefore, we can imagine that every client is routed to the same servers that contain all videos. It indicates how effectively distributed is the Akamai infrastructure. Observing the average values of the RTTs shown in Fig. 4(c), we can assert that YouTube, thanks to the CDN created by Google, is very globally distributed, bringing the contents as close as possible to the clients. However, a high RTT value for Asian countries can be spotted. The results from China can probably be related to the censorship operations implemented by the Government.

5.1.7. Jitter by country

Fig. 5 shows the average and standard deviation of the Jitter aggregated by country. More specifically, Fig. 4(a) shows the values regarding the measurements for Dailymotion. The average values are small regardless of the video category and the country involved in the measurement. The countries with highest Jitter values are Canada, Italy, and Korea. The average Jitter in these countries is higher than 10 ms, with Canada that shows an average jitter value larger than 30 ms for 120K views video category. Fig. 5(b) shows the Jitter values regarding Vimeo. In this provider, the majority of countries record Jitter near to 0 ms. The average is similar for each video category, but even in this provider Canada, Italy, and Korea depict Jitter average higher than 10ms. Fig. 5(c) shows how YouTube have a behavior quite similar to the previous providers in the same countries.

5.2. Second measurement campaign

We performed another measurement campaign to analyze the evolution over time of these infrastructures and their related performance.

5.2.1. Throughput

As in Section 5.1.1, Table 5 describes the statistical indicators of the throughput achieved by the PlanetLab nodes for each provider and video category. In the second campaign, Dailymotion has median values that do not reach more than 532 KiB/s. This provider shows similar values of throughput for every video category, except that of 10K views, which is the category achieving the worst results. Vimeo shows nearly the same statistical indicators for each video category. YouTube has similar behavior with the best performance recorded by 1M-views video category. This provider has obtained the best performance compared to its competitors. More in depth, YouTube achieves a maximum throughput value of about 65 MiB/s and a minimum RTT of 0.235 ms. This value was recorded by a PlanetLab node in the US where the YouTube cache server is the same Autonomous System.

5.2.2. Round Trip Time

Table 6 shows the statistical values related to the RTT. Regarding Dailymotion, RTT varies from 2.7 ms to 343 ms, with median values of about 48 ms. These values are always larger than the ones from the competitors. Nodes deployed in France achieve minimum values for all video categories, which confirms the centralized nature of Dailymotion infrastructure, Section 3.1. Regarding Vimeo, there is a variation from 1 ms to 289 ms, with comparable minimum values independent of the video category. All the values have been recorded by PlanetLab nodes located at the University of Denver at only one hop distance from the nodes of Fastly CDN used by the PlanetLab node. Finally, YouTube presents better performance than the other providers. Its servers show RTT values from 0.235 ms to 1214 ms. The minimum RTT values are reached by nodes that are in the same AS of YouTube cache servers that physically contain the videos.

5.2.3. Jitter

Table 7 shows the statistical values related to the Jitter. Jitter values measured for Dailymotion vary from 0.01 ms to 38 ms, with median values of about 50 ms. These values are better than the ones from the competitors. Regarding Vimeo, there is a variation from 0.02 ms to 119 ms. Finally, as for RTT values, YouTube presents very similar to

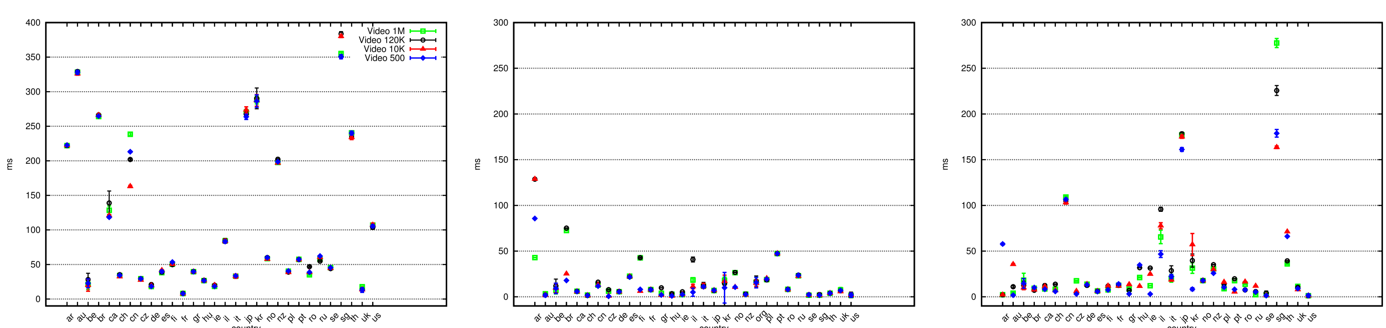

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 4. Average RTT in each country 2015.

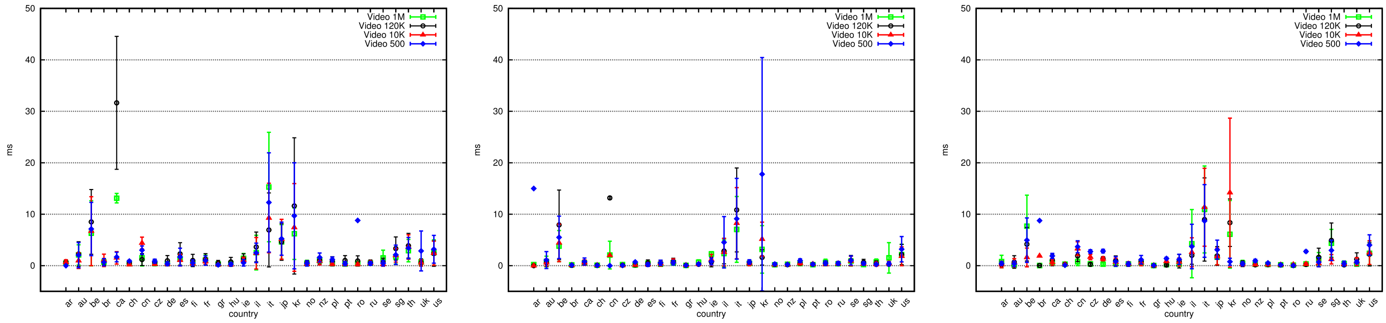

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 5. Average Jitter in each country 2015.

Vimeo with jitter values from 0.01 ms to 83 ms. These results confirm the trend observed for the other performance parameters.

5.2.4. *Temporal behavior*

In the following sections, the evolution of performance over time is presented. Fig. 6 shows the performance achieved by the three providers in a whole day of experimentation. Regarding Dailymotion, Fig. 6(a), we can notice small standard deviations and almost the same mean value of throughput independently of the hour of the day for all the video categories. The temporal evolution of Vimeo and YouTube, in Fig. 6(b) and (c), show better performance but high variability compared to Dailymotion. These providers do not show an average almost constant as the former one but they achieve the worst performance between 4 pm and 10 pm for all video categories. More precisely, Vimeo presents average values between 2 and 3 MiB/s, and Youtube values between 5 and 7.5 MiB/s. The high variability observed for each hour of the day is mainly due to the aggregation of very different values obtained from all over the world and to the buffering policies of the providers, as explained at the beginning of Section 5.

5.2.5. *Performance evaluation by country*

This section presents the evolution of throughput as a function of the country where nodes are deployed. Dailymotion has recorded almost the same throughput for each video category, but three countries (i.e. Australia, Brazil, Hong Kong, and New Zealand) show lower performance (see Fig. 7(a)) than others. Fig. 7(b) depicts the throughput achieved by Vimeo in each country. In this case, better performance is recorded where clients are deployed “near” Fastly PoPs. In fact, in countries with the highest performance, such as Germany, United States, and Sweden, Fastly has its servers deployed. Fig. 7(c) shows throughput per country for YouTube. It depicts a high variability (high values of standard deviation) in almost all the countries involved in the measurements. Fig. 7(c) also shows that like previous providers, YouTube has almost the same performance regardless video categories in all countries. However, in Poland and Brazil, the videos with a larger number of views have better performance.

Finally, for all providers under test, particular traffic management policy depending on the video category did not emerge from the results.

Table 5
Throughput (KiB/s)–Statistical indicators 2016.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	28	482	485	459	487	517
dm-120K	25	438	440	422	440	532
dm-10K	31	375	375	366	375	390
dm-500	27	430	431	415	432	436
vi-1M	64	1108	2666	3204	4259	24114
vi-120K	64	1085	2657	3089	4229	21929
vi-10K	57	1123	2647	3126	4231	24171
vi-500	54	1136	2598	2963	4064	21514
yt-1M	72	3354	6385	8844	10788	66711
yt-120K	73	3413	5729	8599	10649	43812
yt-10K	74	3287	5815	8180	10496	47921
yt-500	66	3112	6094	8591	10685	49608

Table 6
RTT(ms)–Statistical indicators 2016.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	2.84	26.28	48.00	64.57	87.21	329.91
dm-120K	2.70	26.19	48.71	64.56	87.06	342.82
dm-10K	2.86	26.27	48.36	65.50	86.33	315.04
dm-500	2.70	26.09	48.05	63.65	78.69	331.79
vi-1M	1.11	8.11	12.58	35.55	30.98	277.62
vi-120K	1.08	8.16	12.46	35.45	31.00	289.06
vi-10K	1.09	7.97	12.49	35.59	31.11	277.54
vi-500	1.07	8.13	12.40	35.38	30.99	289.06
yt-1M	0.25	1.60	6.01	9.71	11.61	77.80
yt-120K	0.24	1.66	6.35	9.52	10.89	70.71
yt-10K	0.24	1.74	6.37	10.57	10.87	1213.97
yt-500	0.26	1.55	5.92	12.61	11.60	148.35

5.2.6. *RTT by country*

Average and standard deviation values of RTT, regarding Dailymotion, are shown in Fig. 8(a). In all countries, these values are

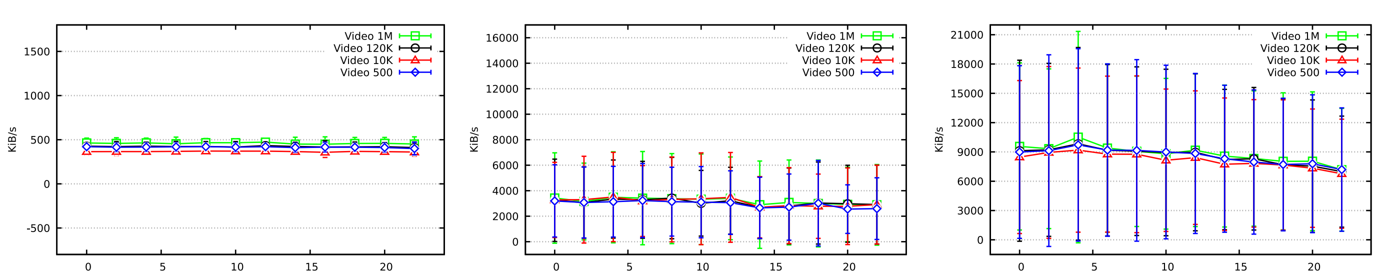

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 6. Average throughput over 24h - 2016.

Table 7
Jitter(ms)–Statistical Indicators 2016.

Provider-Class	Min	I Quart.	Median	Average	III Quart.	Max
dm-1M	0.01	0.19	0.49	1.50	1.36	22.58
dm-120K	0.01	0.22	0.51	1.63	1.44	32.04
dm-10K	0.01	0.22	0.52	1.58	1.50	19.98
dm-500	0.01	0.21	0.50	1.56	1.25	38.25
vi-1M	0.02	0.13	0.36	1.48	1.16	75.90
vi-120K	0.02	0.14	0.34	1.41	1.13	108.90
vi-10K	0.02	0.12	0.32	1.31	1.12	76.10
vi-500	0.02	0.12	0.32	1.45	1.13	119.10
yt-1M	0.02	0.08	0.14	0.69	0.62	19.53
yt-120K	0.02	0.08	0.13	0.80	0.64	30.69
yt-10K	0.01	0.08	0.14	0.74	0.60	82.69
yt-500	0.01	0.07	0.14	0.87	0.65	14.85

almost overlapped, suggesting that no traffic treatments are applied depending on video category. Nodes located in European countries show better performance than ones in other continents. The average values of RTT for the latter are always larger than 50 ms, up to 200 ms. Specifically, the worst RTT is recorded in Australia, Brazil, and New Zeland (more than 150 ms) while the client deployed in France record the best value (about 0.2 ms). Results confirm the hypothesis of a centralized infrastructure for Dailymotion, as anticipated in Section 3. Some exceptions are PlanetLab nodes located in the United States and Asia. These clients connect to servers with nameservers that suggest the presence of a CDN localized respectively in New York and Singapore (i.e. **.nyc.dailymotion.com* and **.sg.dailymotion.com*). In the following section, a further analysis using traceroute is provided to confirm the geographic location of these servers. Concerning the RTT values of Vimeo, shown in Fig. 8(b), in each country the value of RTT recorded are almost the same regardless of video category. Therefore, it can be assumed that every client in a county is routed to the same PoP that contains the videos. As for throughput measurements, the worst RTT values are recorded by nodes in Australia, Brazil, Hong Kong, and New Zealand. The low value of RTT in Germany, the United States, and Sweden confirm the presence of Fastly PoPs in these countries, as suggested the previous section. Except for Brazil (Video 500), Spain (Video 10K), and Mexico, Fig. 8(c) shows that clients register RTT smaller than 25 ms in all countries when connecting to YouTube, confirming that using Google as CDN, it exploits a more globally distributed infrastructure than its competitors.

5.2.7. Jitter by country

Fig. 9 shows for each provider the jitter values aggregated for country. Jitter values do not differ between providers and high values generally repeat in the same countries for all providers. Therefore, we can suppose that the high jitter values are not due to the behaviors of the different video hosting providers but simply to the state of the network of the country where the PlanetLab node is deployed.

5.3. Campaigns comparison

In this section, it is provided a comparison between the two campaigns of measurement. Fig. 10 summarizes the actual throughput by video category, showing the results of first measurement campaign in Fig. 10(a) and the ones of the second campaign in Fig. 10(b)). We report four box plots for each provider: From videos with more than one million of views to ones with less than 500 views. The boxplot represents the minimum, 1st quartile, mean, median, 3rd quartile and the maximum of the throughput. All providers show median values smaller than mean ones. The comparison indicates that Dailymotion has in both campaigns comparable minimum, median, and mean values. Compared to the first campaign, the campaign of 2016 depicts smaller maximum values bounded around 500KiB/s; Vimeo records similar behaviors over the years but overall worst performance. The opposite is true for YouTube, which performance slightly improves for each video category.

Regarding the RTT, the same kind of comparison is shown in Fig. 11. All providers has statistical indicators comparable among the video categories as for the previous comparison (Fig. 11(a) and (b)). On the one hand, the second campaign shows better RTT value over the years for Dailymotion and YouTube regardless the video category. That can be due to improvements in the delivery infrastructure and new peering agreements with ISP (and research institutions for PlanetLab). On the other hand, even if Vimeo presents improvements of the minimum values, its overall performance is worse for all video category. This may be caused by the fact that this provided moved from Akamai to Fastly CDN. In fact, though Fastly has PoPs more powerful than Akamai, it has less peering with ISPs and research entity than the Akamai (see Section 3).

More in depth, Tables 8–10 report respectively the percentage differences of throughput, RTT, and jitter statistical indicators, compared to that of the first campaign. To estimate the performance evolution from the first to the second campaign, we have calculated the percentage differences of all three performance parameters. These results are reported in –10 respectively. In Table 8, the first four rows reveal that Dailymotion has maximum values of throughput decreased 98%, but its median values slightly deviate from the ones recorded in the second campaign. The second block of rows shows that Vimeo has higher minimum values, with an increase up to 245% compared to the first campaign. Its median, average, and maximum values are worse in the first campaign than in the second one. The rows related to YouTube presents improvements in the second campaign of all estimators of throughput. Table 9 shows an improvement of RTT for Dailymotion and YouTube. Both providers have a smaller median and average values that mean ones, which is probably due to an enhancement of their delivery infrastructures. Lower values of RTT can be due to an improvement of the network or by the addition of new PoPs closer to clients. On the other hand, Vimeo presents smaller maximum values compared to the first campaign, and performance probably deteriorated because of the switch from Akamai to Fastly CDN. Finally, Table 10

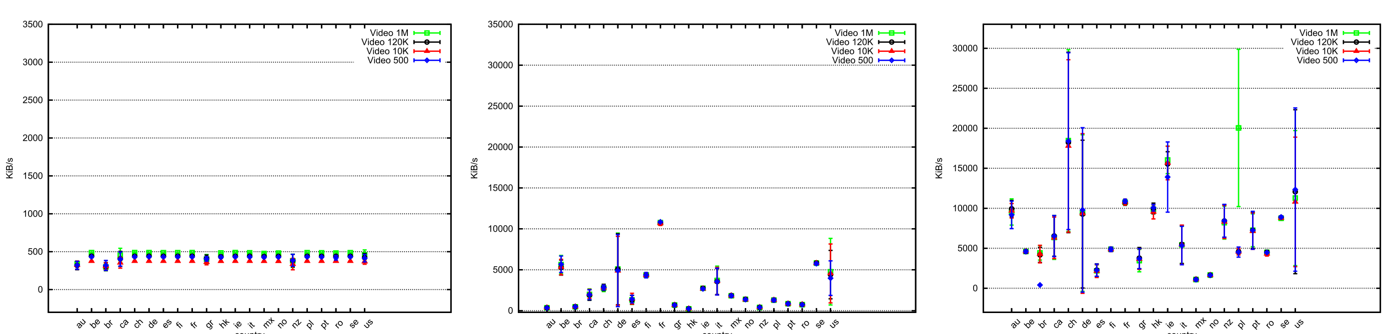

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 7. Average throughput in each country 2016.

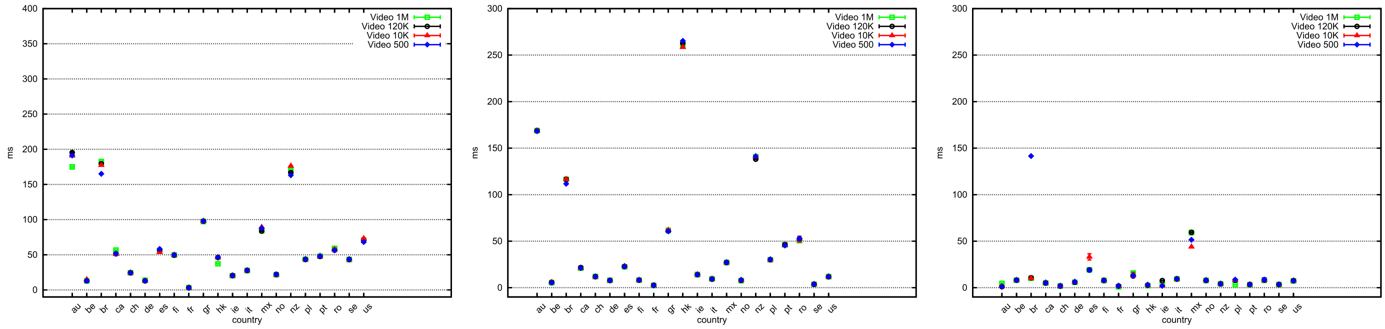

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 8. Average RTT in each country 2016.

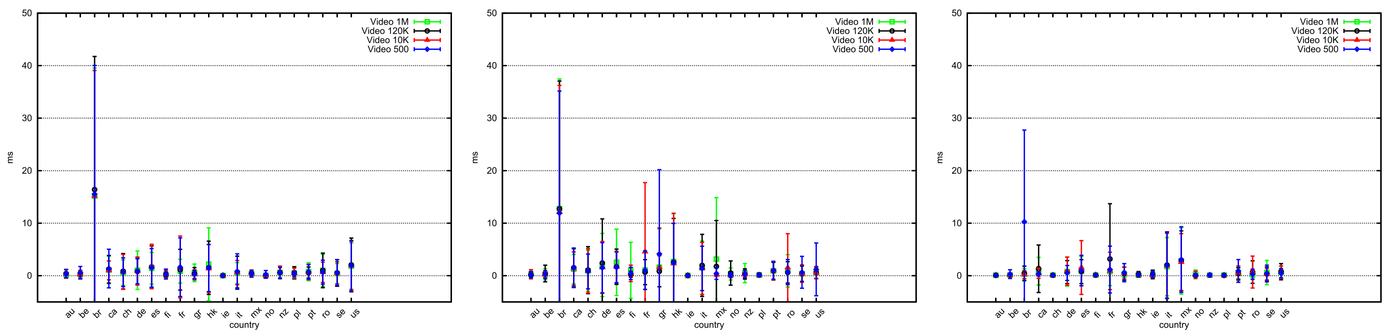

(a) Dailymotion (b) Vimeo (c) YouTube

Fig. 9. Average Jitter in each country 2016.

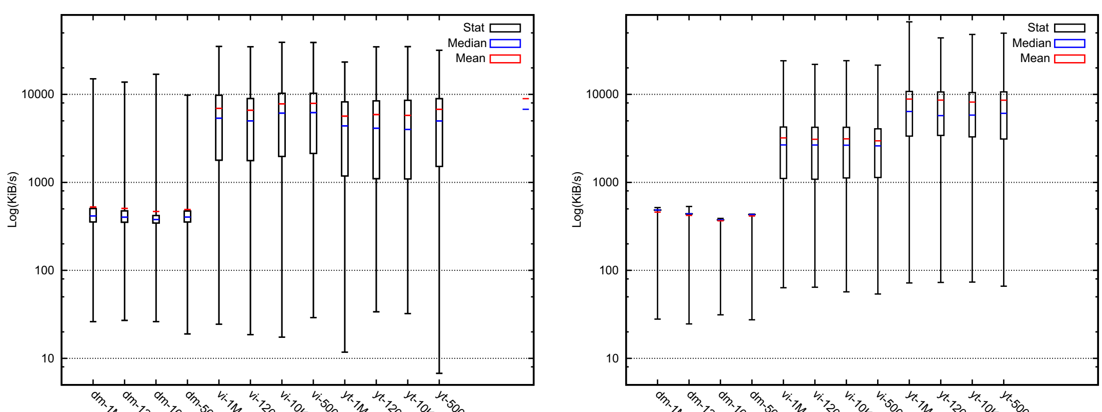

(a) Throughput - Campaign 2015 (b) Throughput - Campaign 2016

Fig. 10. Throughput–comparison among Dailymotion, Vimeo, and YouTube.

shows an improvement of jitter for Dailymotion and YouTube. These value corroborating the hypothesis that the two providers have enhanced their infrastructure. On the other hand, Vimeo shows a performance deterioration regarding the jitter, because of the change of the CDN infrastructure provider.

6. Discussion

In this section we report the results obtained applying the techniques described in Section 2 to geographically locate the servers of video hosting services. For each provider, the list of IPv4 addresses was obtained through the methodology described in Section 4. We also used *traceroute* and *ping* from different vantage points with known location. Notice that *traceroute* does not always correctly provide all the hops for the entire path. The whole operation is affected by issues well known in literature, for example: load balancing [40], anonymous routers [41], hidden routers [42], misleading intermediate delay [43], and third-party addresses [44].

However, we checked that these issues were not affecting the considerations we made out of the obtained results, which are reported below.

6.1. Geolocation IPv4 servers

We have compared the results obtained using two different techniques: *Geoping*, based on the values of RTT measured in every country during each video download, and *Geotrack*, which uses *traceroute*, and provides the names and addresses of the routers through which the data flow travels from client to the server.

The results for three significant use cases are reported in the following.

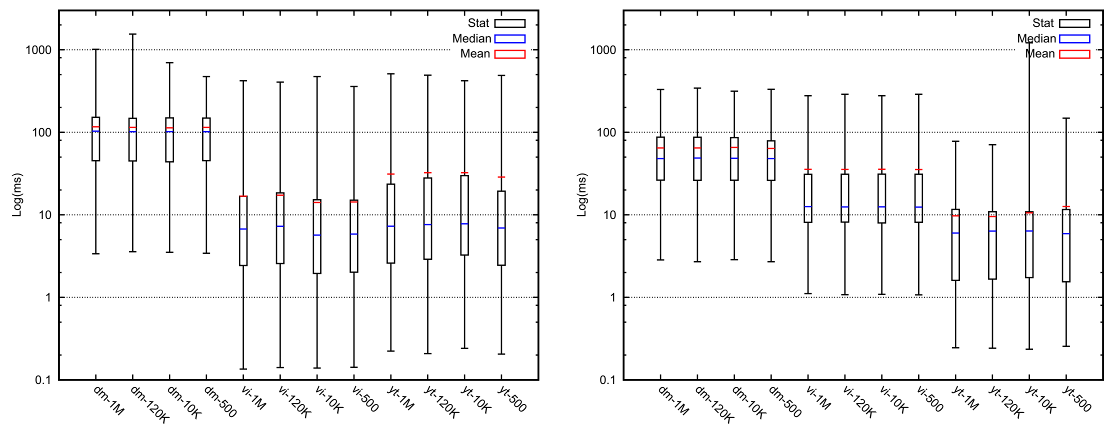

(a) RTT 2015 (b) RTT 2016

Fig. 11. RTT–Campaign comparison.

Table 8
Throughput–Statistical comparison 2015 vs 2016.

Provider-Class	Min diff	Median diff	Average diff	Max diff
dm-1M	7%	17%	− 12%	− 97%
dm-120K	− 9%	9%	− 17%	− 96%
dm-10K	19%	− 1%	− 22%	− 98%
dm-500	45%	7%	− 16%	− 96%
vi-1M	160%	− 50%	− 54%	− 31%
vi-120K	245%	− 47%	− 53%	− 37%
vi-10K	227%	− 57%	− 60%	− 38%
vi-500	86%	− 58%	− 62%	− 45%
yt-1M	511%	46%	56%	186%
yt-120K	116%	38%	46%	26%
yt-10K	128%	45%	41%	37%
yt-500	880%	22%	27%	56%

Table 9
RTT–Statistical comparison 2015 vs 2016.

Provider-Class	Min diff	Median diff	Average diff	Max diff
dm-1M	− 16%	− 53%	− 45%	− 67%
dm-120K	− 24%	− 52%	− 44%	− 78%
dm-10K	− 19%	− 53%	− 42%	− 55%
dm-500	− 21%	− 53%	− 44%	− 30%
vi-1M	724%	87%	111%	− 34%
vi-120K	664%	72%	105%	− 29%
vi-10K	683%	120%	153%	− 41%
vi-500	655%	112%	147%	− 19%
yt-1M	10%	− 17%	− 69%	− 85%
yt-120K	16%	− 17%	− 71%	− 86%
yt-10K	− 2%	− 18%	− 67%	188%
yt-500	24%	− 15%	− 56%	− 70%

Table 10
Jitter–Statistical comparison 2015 vs 2016.

Provider-Class	Min diff	Median diff	Average diff	Max diff
dm-1M	− 89%	− 4%	− 45%	− 98%
dm-120K	0%	− 26%	− 49%	− 98%
dm-10K	0%	24%	− 26%	− 91%
dm-500	0%	− 2%	− 38%	− 90%
vi-1M	*NaN*	80%	− 8%	− 45%
vi-120K	*NaN*	70%	− 18%	− 56%
vi-10K	*NaN*	78%	− 8%	− 66%
vi-500	*NaN*	52%	− 28%	− 46%
yt-1M	*NaN*	− 18%	− 64%	− 92%
yt-120K	− 33%	− 24%	− 55%	− 90%
yt-10K	− 67%	− 18%	− 60%	− 76%
yt-500	*NaN*	− 22%	− 53%	− 92%

- Switzerland: Fig. 12 shows the paths from a user in Switzerland (i.e. a PlanetLab node in Swiss Federal Institute of Technology Lausanne) to the servers containing the videos of the three providers (Dailymotion, Vimeo, and YouTube). The Autonomous System Numbers was obtained by *traceroute* with the option to perform AS path lookups. The user IP address was in AS599 (SWITCH Information Technology Services). The red line shows the path to the Google server, the blue line shows the path to the Dailymotion server, and the sky-blue line shows the path to the Vimeo server. Both Google and Vimeo put its cache inside the same AS of the user. The server containing the videos was hosted in the AS41690 owned by Dailymotion itself.
- South Korea: Fig. 13 shows the path from the user in South Korea. The user IP address was in AS18176 (Kookmin University). Even in this case, the server that physically contains the video of Dailymotion was hosted on its own AS in France, in agreement with the initial assumption that Dailymotion does not use a globally distributed CDN.
- France: This last use case refers to France users. Fig. 14 shows that Dailymotion appears more distant, in terms of network hops than its competitors, even in the case when requests come from the same country where its entire infrastructure was deployed (i.e. France).

For each provider, it is possible to summarize the following results. Dailymotion deploys its entire infrastructure in France (i.e. in Paris). Although there are abnormal activities by some nodes as we described in the previous section, the presence of other caches distributed elsewhere in the world has not been detected. Vimeo uses Akamai (first campaign) and Fastly (second campaign) infrastructures. Both providers have a worldwide distributed servers network. At every video request, the clients are redirected to the "closer" server to maximize the performance. It is highlighted by the lower values of the RTT and by the names of the servers containing the video. Customers requests are redirected to the same server, irrespective of the day time or the network overload. Only in case the content is not present in the cache-server, the clients are redirected to the back-end servers. YouTube (i.e. Google) infrastructure presence is wide. Google Global Cache is located inside the networks of *Internet Service Providers* and *Internet Exchange Points* to serve regional users. The presence of these servers covers almost all the countries in which PlanetLab nodes involved in experimental measurements are deployed. Unlike Akamai, there is evidence of delivery strategies that assess both the "distance" between client and server and the "load" of the network.

6.2. IPv4 Identification and name-server of the providers

We used the reverse DNS to determine the *name-servers* associated to

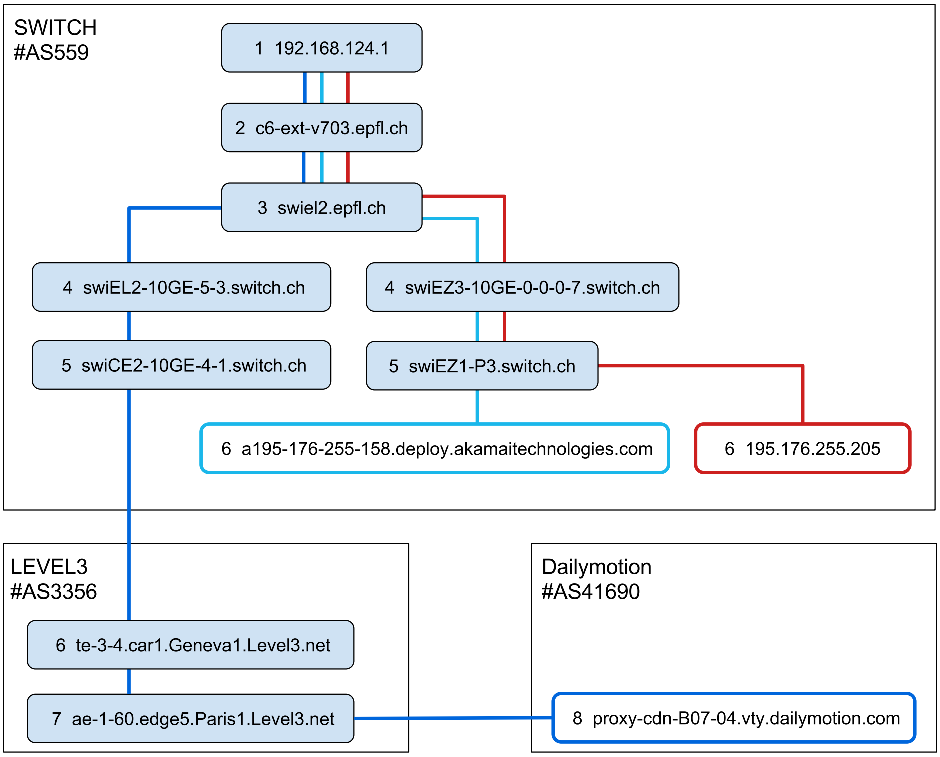

Fig. 12. Example of routing related to users from Switzerland.

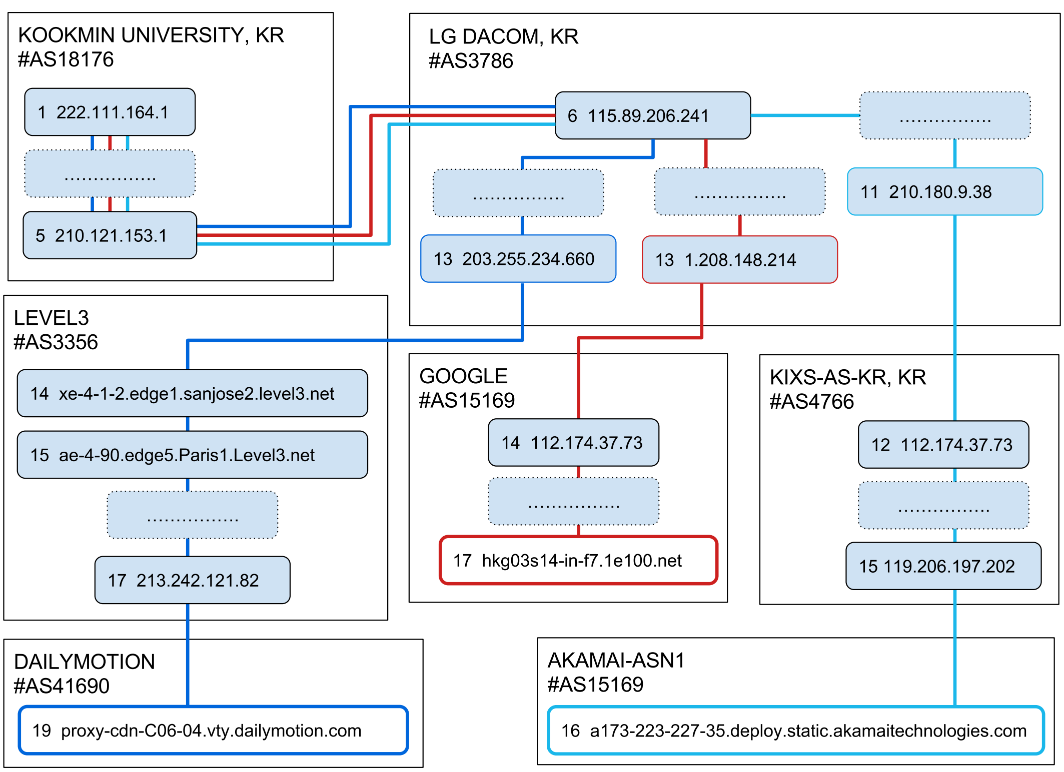

Fig. 13. An example of routing related to measurements in South Korea.

the IPv4 of all the servers contacted. From the list of nameservers consequently obtained, we can say that:

- For Dailymotion, there are only 8 servers from which clients, from all over the world, download videos;
- For Vimeo, the servers from which the downloads are predominantly made are part of the network of Akamai or Fastly CDN;
- For YouTube, the servers are globally distributed, but not always they belong to Google. Sometimes *name-server* identified them as part of a telecommunication or a hosting company, such as Tiscali, Asianet Web, or Oneandone.

Referring to the discussions on the Network Neutrality and the study of addresses and nameservers, there is no evidence of preferential

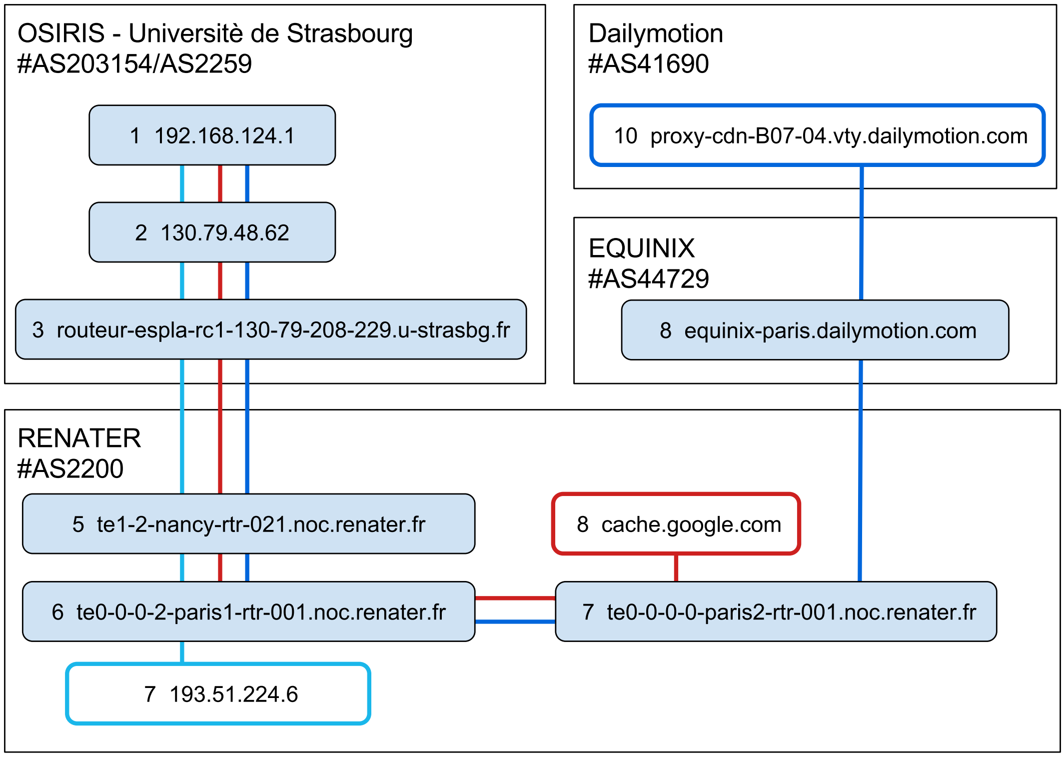

Fig. 14. An example of routing related to measurements in France.

treatments related to the video categories by any provider. On the other hand, we have evidence of the use of cache servers in ISPs, IXPs and even in the same Autonomous Systems of the clients. For this reason, the different performance observed is due to different infrastructures used by providers.

7. Conclusion

The aim of our study was to compare the performance indicators of video hosting services, to understand whether the performance differences could impact network neutrality and users privacy. It is worth noting that we did not want to determine whether neutrality is good or bad, but we rather wanted to evaluate the performance differences from the user point of view. We proposed a methodology that, regardless of the provider, allows to acquire and analyze performance data and topology information about the infrastructure of the providers. To validate the methodology, a comparison of the three video hosting services (Dailymotion, Vimeo, and YouTube) was performed on the basis of performance indicators (i.e. throughput, RTT, Jitter, and TTL), geographical location of the infrastructures, and routing policies used by the video hosting services. The data collection phase was repeated after one year in order to evaluate how the providers' infrastructure has evolved over time. In agreement with the analysis of the information available in the news, regarding the infrastructures of the providers we performed, our experimental results confirm that Dailymotion has a centralized infrastructure. Moreover, its performance decays with the clients' distance from infrastructure location. Vimeo and YouTube use CDNs to deliver their contents, where the first showed the best performance indicators compared to its competitors. For both infrastructures, there is evidence of cache servers in the same Autonomous Systems of the PlanetLab nodes used as clients. We clearly showed how providers using distributed infrastructure could reach better performance, up to one order of magnitude larger throughput and smaller RTT values. Regarding network neutrality, no evidence of special treatment based on video category has been collected. The highlighted performance differences may still be regarded as a lack of neutrality because all providers should be able to benefit from the same conditions of distribution and spread of their contents. However, such differences are not due to different treatments of traffic, but rather to different delivery infrastructures. Deciding on whether this is or not a neutrality violation is out of the scope of this paper. We rather aimed at providing a clear picture of the current situation and performance of video hosting services over the Internet. We are currently working on the continuation of this work considering other important parameters that can be an indication of mechanisms put in place by providers to gain better performance as well as other kinds of access network besides the academic Internet.

Acknowledgment

This work is partially funded by the MIUR project art. 11 DM 593/2000 for NM2 srl.

References

[1] Sandvine, What are the latest internet trends in North America and Latin America?, 2016.
[2] T. Wu, Network neutrality, broadband discrimination, J. Telecommun. High Technol.Law 2 (2001) (2003) 141, http://dx.doi.org/10.2139/ssrn.388863.
[3] M.A. Lemley, L. Lessig, The end of end-to-end: preserving the architecture of the internet in the broadband era, UCLA Law Rev. 48 (2000).
[4] Svensson, Comcast blocks some subscriber internet traffic, ap testing shows, 2007, Accessed: 2016-04-21.
[5] B.-J. Koops, J.P. Sluijs, Network neutrality and privacy according to art. 8 echr, SSRN Electron. J. (08) (2011) 0–20, http://dx.doi.org/10.2139/ssrn.1920734.
[6] Cisco, Cisco visual networking index: Forecast and methodology, 20152020, 2015, Accessed: 2016-04-21.
[7] M. Saxena, U. Sharan, S. Fahmy, Analyzing video services in Web 2.0, Proceedings of the 18th International Workshop on Network and Operating Systems Support for Digital Audio and Video - NOSSDAV '08, ACM Press, New York, New York, USA, 2008, p. 39, http://dx.doi.org/10.1145/1496046.1496056.
[8] D. Chun B.and Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, M. Bowman, Planetlab: an overlay testbed for broad-coverage services, SIGCOMM Comput. Commun. Rev. 33 (3) (2003) 3–12, http://dx.doi.org/10.1145/956993.956995.
[9] M. Cha, H. Kwak, P. Rodriguez, Y.Y. Ahn, S. Moon, I tube, you tube, everybody tubes: analyzing the world's largest user generated content video system, Proc ACM

Internet Measurement Conference IMC San Diego CA October 2007, New York, (2007), pp. 1–14, http://dx.doi.org/10.1145/1298306.1298309.
[10] L. Plissonneau, E. Biersack, P. Juluri, Analyzing the impact of youtube delivery policies on user experience, Proceedings of the 24th International Teletraffic Congress, ITC '12, International Teletraffic Congress, 2012, pp. 1–28.
[11] L. Plissonneau, E. Biersack, A longitudinal view of http video streaming performance, Proceedings of the 3rd Multimedia Systems Conference, MMSys '12, ACM, New York, NY, USA, 2012, pp. 203–214, http://dx.doi.org/10.1145/2155555.2155588.
[12] V.K. Adhikari, Y. Guo, F. Hao, V. Hilt, Z.-L. Zhang, M. Varvello, M. Steiner, Measurement study of netflix, hulu, and a tale of three CDNs, IEEE/ACM Trans. Netw. 23 (6) (2015) 1984–1997, http://dx.doi.org/10.1109/TNET.2014.2354262.
[13] A. Botta, A. Avallone, M. Garofalo, G. Ventre, Internet streaming and network neutrality: comparing the performance of video hosting services, Proceedings of the 2nd International Conference on Information Systems Security and Privacy, (2016), pp. 514–521, http://dx.doi.org/10.5220/0005798705140521.
[14] J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, B. Weihl, Globally distributed content delivery, IEEE Internet Comput. 6 (5) (2002) 50–58, http://dx.doi.org/10.1109/MIC.2002.1036038.
[15] E. Nygren, R.K. Sitaraman, J. Sun, The akamai network: a platform for high-performance internet applications. ACM SIGOPS Operating Syst. Rev. 44 (3) (2010) 2, http://dx.doi.org/10.1145/1842733.1842736.
[16] A.J. Su, D.R. Choffnes, A. Kuzmanovic, F.E. Bustamante, Drafting behind akamai: inferring network conditions based on CDN redirections, IEEE/ACM Trans. Netw. 17 (6) (2009) 1752–1765, http://dx.doi.org/10.1109/TNET.2009.2022157.
[17] S. Triukose, Z. Wen, M. Rabinovich, Measuring a commercial content delivery network, Proceedings of the 20th International Conference on World Wide Web–WWW '11, ACM Press, 2011, p. 467, http://dx.doi.org/10.1145/1963405.1963472.
[18] V.N. Padmanabhan, L. Subramanian, An investigation of geographic mapping techniques for internet hosts, SIGCOMM Comput. Commun. Rev. 31 (4) (2001) 173–185, http://dx.doi.org/10.1145/964723.383073.
[19] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, R. Govindan, Mapping the expansion of google's serving infrastructure, Proceedings of the 2013 Conference on Internet Measurement Conference, IMC '13, ACM, New York, NY, USA, 2013, pp. 313–326, http://dx.doi.org/10.1145/2504730.2504754.
[20] P. Pelaprat, Dailymotion construit son infrastructure sur un cluster, 2007, Accessed: 2016-04-21.
[21] EMC2, A breakthrough in web-based video content sharing, 2010, Accessed: 2016-04-21.
[22] Orange, Orange business services helps dailymotion optimize live streaming for its partners, 2014, Accessed: 2016-04-21.
[23] Akamai, Adaptive media delivery, 2014, Accessed: 2016-04-21.
[24] Limelight networks helps dailymotion to deliver video as peak activity increases in the asia pacific region, 2016, Accessed:2016-04-21.
[25] Vimeo, Cdn usage, 2013, Accessed: 2016-04-21.
[26] Akamai, Akamai: facts & figures, 2016, Accessed: 2016-04-21.
[27] Akamai, Download delivery, 2014, Accessed: 2016-04-21.
[28] Akamai, Improving online video quality and accelerating downloads, 2015, Accessed: 2016-04-21.
[29] S. Jain, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, A. Vahdat, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, B4: Experience with a globally-deployed software defined wan, Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM - SIGCOMM '13, ACM Press, 2013, p. 3, http://dx.doi.org/10.1145/2486001.2486019.
[30] Google, Google peering & content delivery - content delivery ecosystem, 2016, Accessed: 2016-04-21.
[31] T. Hoßfeld, R. Schatz, E. Biersack, L. Plissonneau, Data Traffic Monitoring and Analysis: From Measurement, Classification, and Anomaly Detection to Quality of Experience, ch. Internet Video Delivery in YouTube: From Traffic Measurements to Quality of Experience, Springer Berlin Heidelberg, Berlin, Heidelberg, 2013, pp. 264–301.
[32] IATA, Iata airport code, 2010, Accessed: 2016-04-21.
[33] B. Ahlgren, C. Dannewitz, C. Imbrenda, D. Kutscher, B. Ohlman, A survey of information-centric networking, IEEE Commun. Mag. 50 (7) (2012).
[34] L. Zhang, A. Afanasyev, J. Burke, V. Jacobson, P. Crowley, C. Papadopoulos, L. Wang, B. Zhang, et al., Named data networking, ACM SIGCOMM Comput. Commun. Rev. 44 (3) (2014) 66–73.
[35] D. Saxena, V. Raychoudhury, N. Suri, C. Becker, J. Cao, Named data networking: a survey, Comput. Sci. Rev. 19 (2016) 15–55.
[36] B. Rainer, D. Posch, H. Hellwagner, Investigating the performance of pull-based dynamic adaptive streaming in ndn, IEEE J. Sel. Areas Commun. 34 (8) (2016) 2130–2140.
[37] D. Kulinski, J. Burke, Ndnvideo: random-access live and prerecorded streaming using ndn, Technical Report NDN-0007, (2012).
[38] S. Lederer, C. Müller, C. Timmerer, Dynamic adaptive streaming over http dataset, Proceedings of the 3rd Multimedia Systems Conference on - MMSys '12, ACM Press, New York, New York, USA, 2012, p. 89, http://dx.doi.org/10.1145/2155555.2155570.
[39] IETF, Http live streaming, 2016, Accessed: 2016-04-21.
[40] B. Augustin, T. Friedman, R. Teixeira, Measuring multipath routing in the internet, Networking, IEEE/ACM Trans. 19 (3) (2011) 830–840, http://dx.doi.org/10.1109/TNET.2010.2096232.
[41] M.H. Gunes, K. Sarac, Resolving anonymous routers in internet topology measurement studies, 2008 Proceedings IEEE INFOCOM - The 27th Conference on Computer Communications, IEEE, 2008, pp. 1076–1084, http://dx.doi.org/10.1109/INFOCOM.2008.162.
[42] P. Marchetta, A. Pescape, Drago: Detecting, quantifying and locating hidden routers in traceroute ip paths, INFOCOM, 2013 Proceedings IEEE, IEEE, 2013, pp. 3237–3242, http://dx.doi.org/10.1109/INFCOM.2013.6567144.
[43] P. Marchetta, A. Botta, E. Katz-Bassett, A. Pescapé, Dissecting Round Trip Time on the Slow Path with a Single Packet, in: M. Faloutsos, A. Kuzmanovic (Eds.), Passive and Active Measurement, Lecture Notes in Computer Science, 8362 Springer International Publishing, 2014, pp. 88–97, , http://dx.doi.org/10.1007/978-3-319-04918-2_9.
[44] P. Marchetta, W. de Donato, A. Pescapé, Detecting Third-party Addresses in Traceroute Traces with Ip Timestamp Option, in: M. Roughan, R. Chang (Eds.), Passive and Active Measurement, Lecture Notes in Computer Science, 7799 Springer Berlin Heidelberg, 2013, pp. 21–30, , http://dx.doi.org/10.1007/978-3-642-36516-4_3.